

Outline

- Specifications for High Intensity Running
 - Motivation, goals
- L3 trigger review
 - Charge
 - Extrapolated rates from Spring 2016 data
 - Total data rate
 - Data rate of individual crates
 - Actions to reduce event size
 - Word mapping (what takes space in EVIO file)
 - Time cuts, thresholds, TAGH
 - Estimate of rates with reduced event size
 - Level - 3 implementation
 - Data flow
 - BDT
 - Review recommendations
- Outlook

High Intensity Running in Hall-D




David Lawrence JLab

Nov. 10, 2016

From PR12-13-003 (GlueX strangeness proposal)

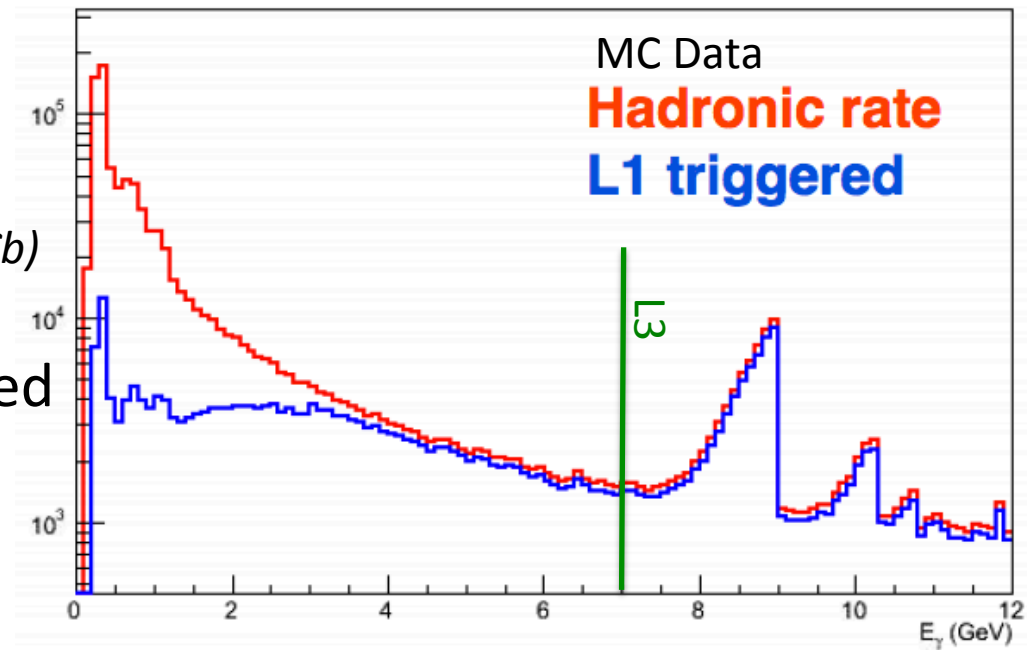
... we propose a gradual increase in the photon flux towards the GlueX design of 10^8 γ /s in the peak of the coherent bremsstrahlung spectrum ($8.4 \text{ GeV} < E_\gamma < 9.0 \text{ GeV}$). Yield estimates, assuming an average flux of 5×10^7 γ /s, are presented.



- 10^8 γ /s on LH_2 target -> ~ 400 kHz hadronic rate
- L1 trigger goal is to cut away $\sim 50\%$ leaving 200 kHz
- L3 trigger goal is to reduce by $\sim 90\%$ leaving 20 kHz **actual: $\sim 75\%$**
- Early simulation suggested ~ 15 kB/event **actual: $16 \text{ kB} + 0.02 \text{ kB}/nA^*$**
 - 15 kB/event @ 200 kHz = 3000 MB/s (front end)
 - L3 reduction by factor of 10 = 300 MB/s to RAID disk

L3 Algorithm Strategy

- Use Multivariate Analysis such as BDT or ANN to classify events (*similar to LHCb*)
- Multiple levels may be used with each level requiring more expensive input variables



- Use fully reconstructed, real data to provide training samples (signal and background)
- Simultaneously pursue with simulated data

L3 Trigger mini-Review

- Quickly organized for July 22nd 2016
- Standard review procedure of critical system for experiment that will run in 2018

From: Patrizia Rossi rossi@jlab.org
Subject: L3 mini-review
Date: July 18, 2016 at 4:54 PM
To: David Lawrence davidl@jlab.org, Eugene Chudakov gen@jlab.org
Cc: Rolf Ent ent@jlab.org, Gomez gomez@jlab.org, patrizia.rossi@jlab.org

Hi Eugene, Dave,

we confirm you that we will have the Hall D L3 min-review on Friday July 22.

We will start at 1pm and plan to finish by 3pm.

Amber Boehnlein and Steve Wood kindly agreed to be the subject matter experts for this review.

The charge is to assess the status of the L3 trigger and plan to completion.

Please prepare a ~1/2 hour presentation to answer this charge.

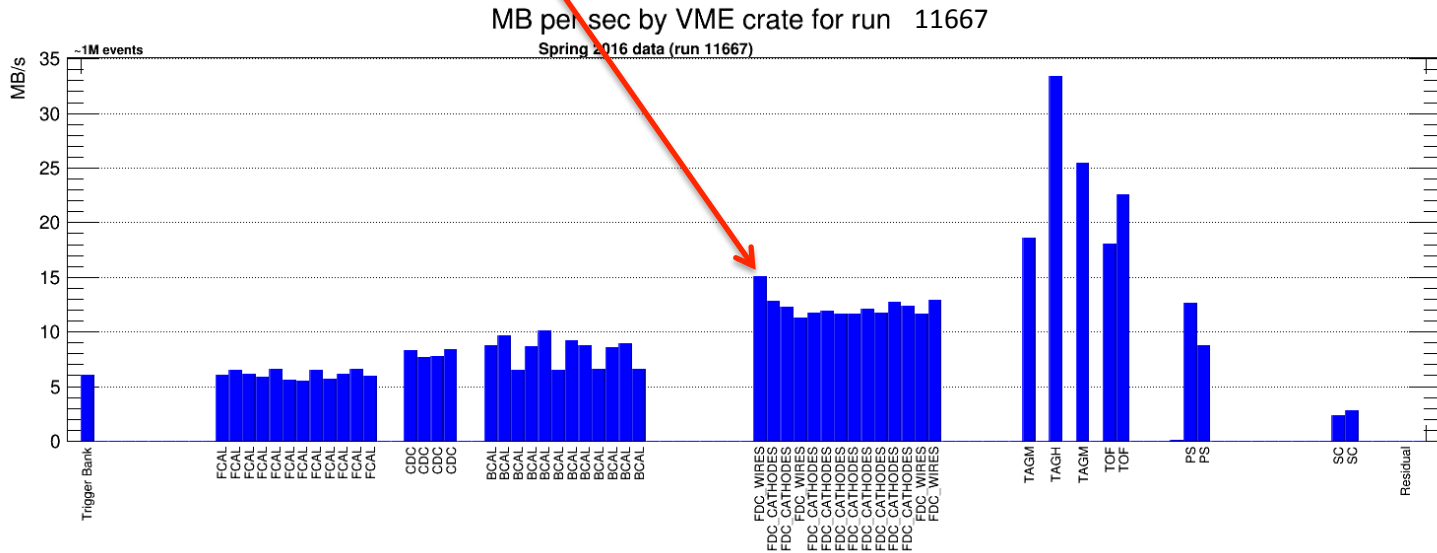

Also may you please schedule the meeting on Amber's and Steve's calendar?

Thank you

patrizia

Shown at July 22nd L3 Review


July 21, 2016 DL
git revision #BDC1004



Improvements since L3 Review

- Alternate Spring 2016 data extrapolated to “High Intensity” ($5 \times 10^7 \gamma/s$)
- Better understanding of photon beam intensity
- Timing windows reduced
- Thresholds increased

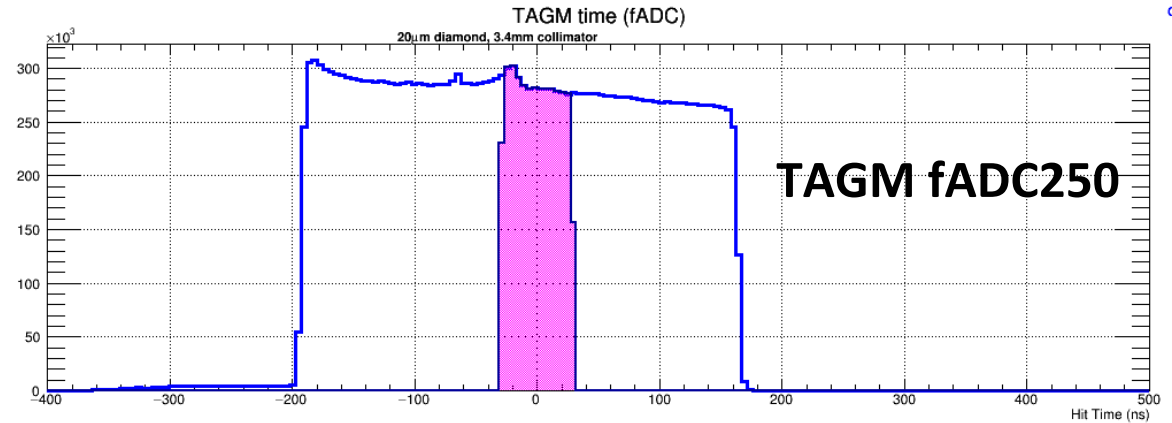
From raw data EVIO file



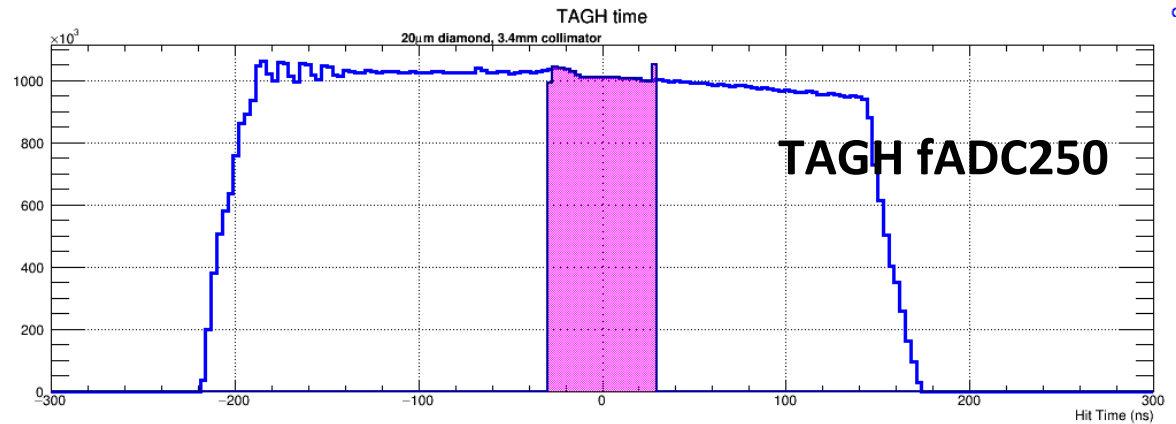
Timing cut examples

Hits outside of the shaded areas were dropped before writing out the reconstituted EVIO file

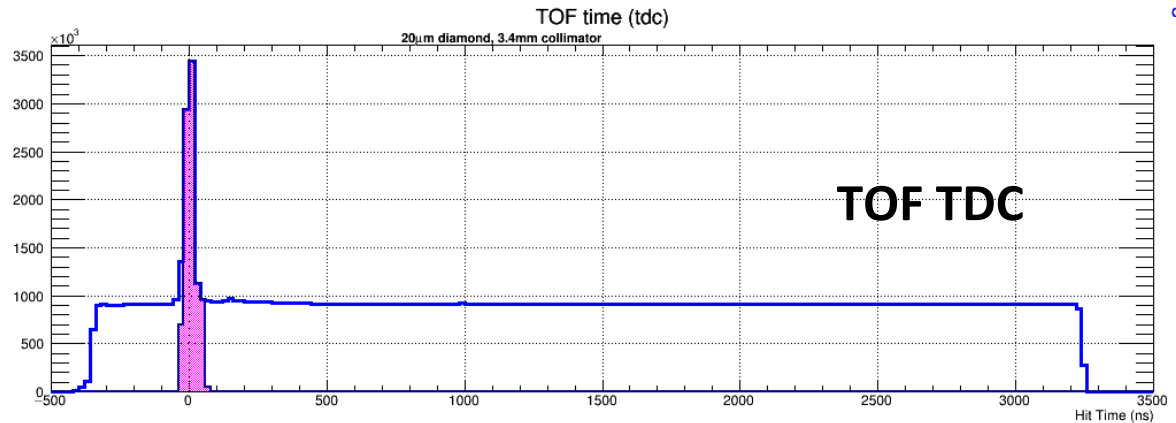
Cuts made on fully calibrated times. Cuts done in modules will be on uncalibrated times



October 1, 2016 DL
previews 25:00:08
run 11363




October 1, 2016 DL
previews 25:00:08
run 11363




October 1, 2016 DL
previews 25:00:08
runs 11363


TAGM pulse_peak




TOF pulse_peak





TAGH pulse_peak



TAGH slot



Event Size vs. Beam Current



$E_e = 12.113\text{GeV}$
 20 μ m diamond
 3.4mm collimator
 200nA = $0.93 \times 10^7 \gamma/s$

Event Size: $\sim 14\text{kB}$ (low intensity)

extrapolate to $I_{\text{beam}} = 0 \rightarrow 11.5\text{kB/event}$
(size of clean event with no accidentals)


Accidental data fraction (by volume):

$$\frac{(0.01\text{kB/nA})(I_{\text{beam}} \text{ nA})}{(0.01\text{kB/nA})(I_{\text{beam}} \text{ nA}) + (11.7\text{kB})}$$

0.9 γ/s : 15% of data is due to accidentals


1.8 γ/s : 26% of data is due to accidentals

5.0 γ/s : 50% of data is due to accidentals



- Single crate rates can be reduced to levels below VME bandwidth limit and the 1Gbit Ethernet limit
- With no changes to L1, the extrapolated rate for high luminosity is 170kHz
- Estimated reduction in L1 rate is 1.6 in addition to what is shown below

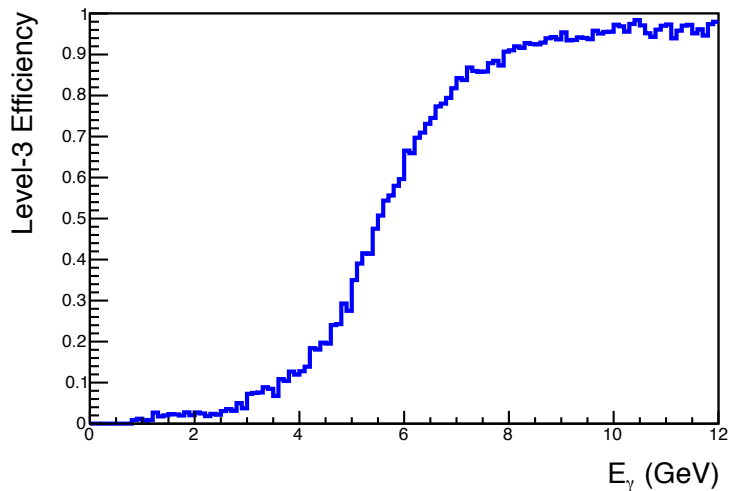
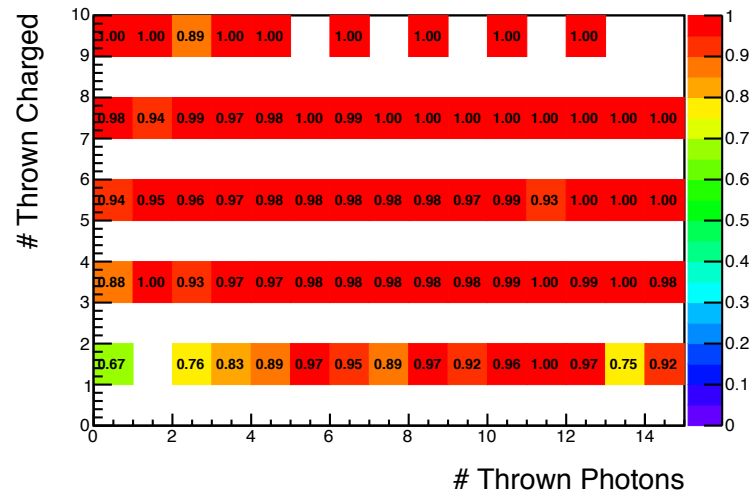
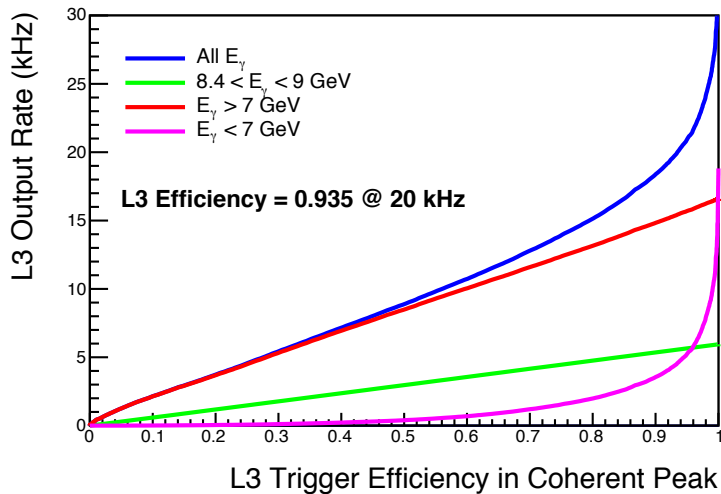
Estimated Data Rate vs. Beam Intensity



L3 Mini-review Recommendations

- fADC125 (performance issues)
 - **New firmware has been deployed that uses increased clock speed for VME transactions. Busy signals installed and testing underway**
- L1 trigger (optimize to reduce rate)
 - **Study by trigger expert suggests rate can be lowered by a factor of 1.6**
- Crate level readout (rate limited by 1Gbit link)
 - **Installed 10Gbit card in one crate and tested link to 8Gbyte/s. Will test with CODA during Fall run**
- System level readout (architecture and offline resources e.g. tape)
 - **Expect pre-L3 rate to be closer to 3GB/s design goal. If L3 reduction is factor of 4, then rate is only 750MB/s to disk.**
- L3 strategies (reconstruction rates with increased accidentals)
 - **High intensity test data will be obtained during Fall 2016 run (with low livetime). L3 algorithms will be benchmarked using this.**
- L3 system design (e.g. farm in counting house vs. Computer Center)
 - **Discussion deferred pending high intensity testing planned for Fall 2016**

Level-3 BDT Evaluation



- * Machinery resurrected for training and evaluating BDT for simulation
- * Similar performance as seen in studies ~2 years ago
- * Lower efficiency for low multiplicity final states

Summary and Outlook


- Spring 2016 data indicates data rates can be achieved that fit within existing hardware capabilities
- Additional high intensity testing has been incorporated into the Fall 2016 run plan
- L3 software:
 - Couples with DAQ system and offline reconstruction
 - Expect to use BDT(s)
 - Beam testing started in Spring 2016, will continue in Fall

Backup Slides




Estimated number of L3 nodes

- From Spring 2016 running: $\sim 32\text{kHz}/0.93 \times 10^7 \gamma/\text{s}$
- For $5 \times 10^7 \gamma/\text{s}$: $\sim 170\text{kHz}$
- Anticipated L1 reduction factor: 1.6
 - L1 trigger rate: $\sim 100\text{kHz}$
- 2013 Ivy Bridge nodes
 - 2.5-13kHz parsing only
 - 4kHz parsing+neutrals recon
 - 1kHz final algorithm (*conservative estimate*)
- Newer nodes more than x2 faster
 - 2kHz/node
- Total estimated number of nodes required:
 - $100\text{kHz}/4\text{kHz} = \mathbf{50}$




Hits outside of the shaded areas were dropped before writing out the reconstituted EVIO file


Cuts made on fully calibrated times. Cuts done in modules will be on uncalibrated times



October 1, 2016 DL
of physics @ SLAC
run 11363




October 1, 2016 DL
of physics @ SLAC
run 11363




October 1, 2016 DL
of physics @ SLAC
run 11363


BCAL pulse_peak




TOF pulse_peak



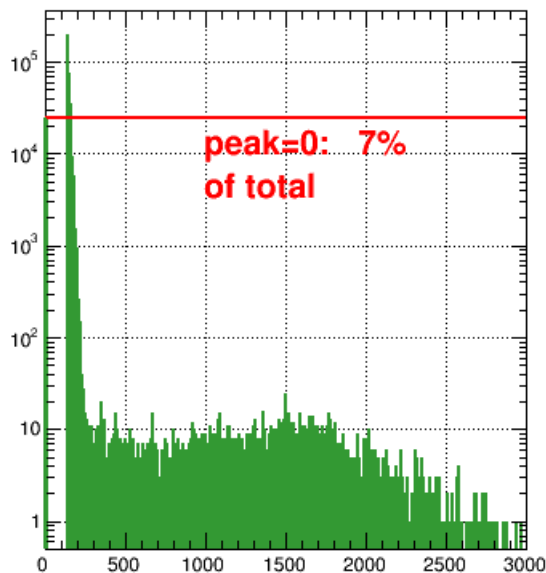
TAGM pulse_peak




FCAL pulse_peak



PS pulse_peak



TAGH pulse_peak



Reconstruction times survey

Time is divided by
#calls and #threads

J1A50 50 um radiator, PERP, 7 mode, 105 nA beam current, 27 kHz event rate, live time ~70% , LH2 fill, 5 mm collimator, 83 M total events



input file: hd_rawdata_010913_060.
phys_skim.evio

events	ALGORITHM	INPUT OBJECT	RECO TIME [s]	RECO TIME/event [ms]
10000	nominal reco	DNeutralShower	4.410	0.028
DATA	approx reco	DBCALShower	2.750	0.017
	approx reco	DFCALShower	8.010	0.050
	full tracking	DTrackTimeBased	18669.810	116.69
	approx tracking	DTrackWireBased	7397.300	46.23

hdgeant_smeared_14980

events	ALGORITHM	INPUT OBJECT	RECO TIME [s]	RECO TIME/event [ms]
10000	nominal reco	DNeutralShower	5.730	0.036
MC	approx reco	DBCALShower	3.430	0.021
	approx reco	DFCALShower	7.010	0.044
	full tracking	DTrackTimeBased	23878.840	149.243
	approx tracking	DTrackWireBased	12778.340	79.865

n.b. parsing of evio data takes 0.080 – 0.400 ms/event



```

--- BDT      : Ranking result (top variable is best ranked)
--- BDT      : -----
--- BDT      : Rank : Variable      : Variable Importance
--- BDT      : -----
--- BDT      : 1 : Efccl_clusters : 1.917e-01
--- BDT      : 2 : Ntrack_candidates : 1.710e-01
--- BDT      : 3 : Nfccl_clusters   : 1.279e-01
--- BDT      : 4 : Nbccl_points     : 1.258e-01
--- BDT      : 5 : Npshits          : 8.291e-02
--- BDT      : 6 : Ebccl_points     : 7.186e-02
--- BDT      : 7 : Ebccl_clusters   : 6.445e-02
--- BDT      : 8 : Ntof             : 6.424e-02
--- BDT      : 9 : Nstart_counter   : 5.138e-02
--- BDT      : 10 : Nbccl_clusters  : 4.873e-02
--- BDT      : 11 : Ptot_candidates  : 0.000e+00
--- BDT      : 12 : Npschits        : 0.000e+00

```


Spring 2016 data

\$SHALLD_HOME/src/plugins/Utilities/I3bdt

- “signal” events had **>4GeV** of fully reconstructed energy*


L3 running in pass-through with beam

- Events are read in blocks of 20
- Events must be disentangled and reconstituted as single events before writing to disk
- Redundant headers may be dropped to reduce event size




bytes/event

Date Rate by Crate




bytes/event

Date Rate by Crate




Hdl3 data




Number of words in EVIO file by type



RAW data




hdl3 data



RAW data



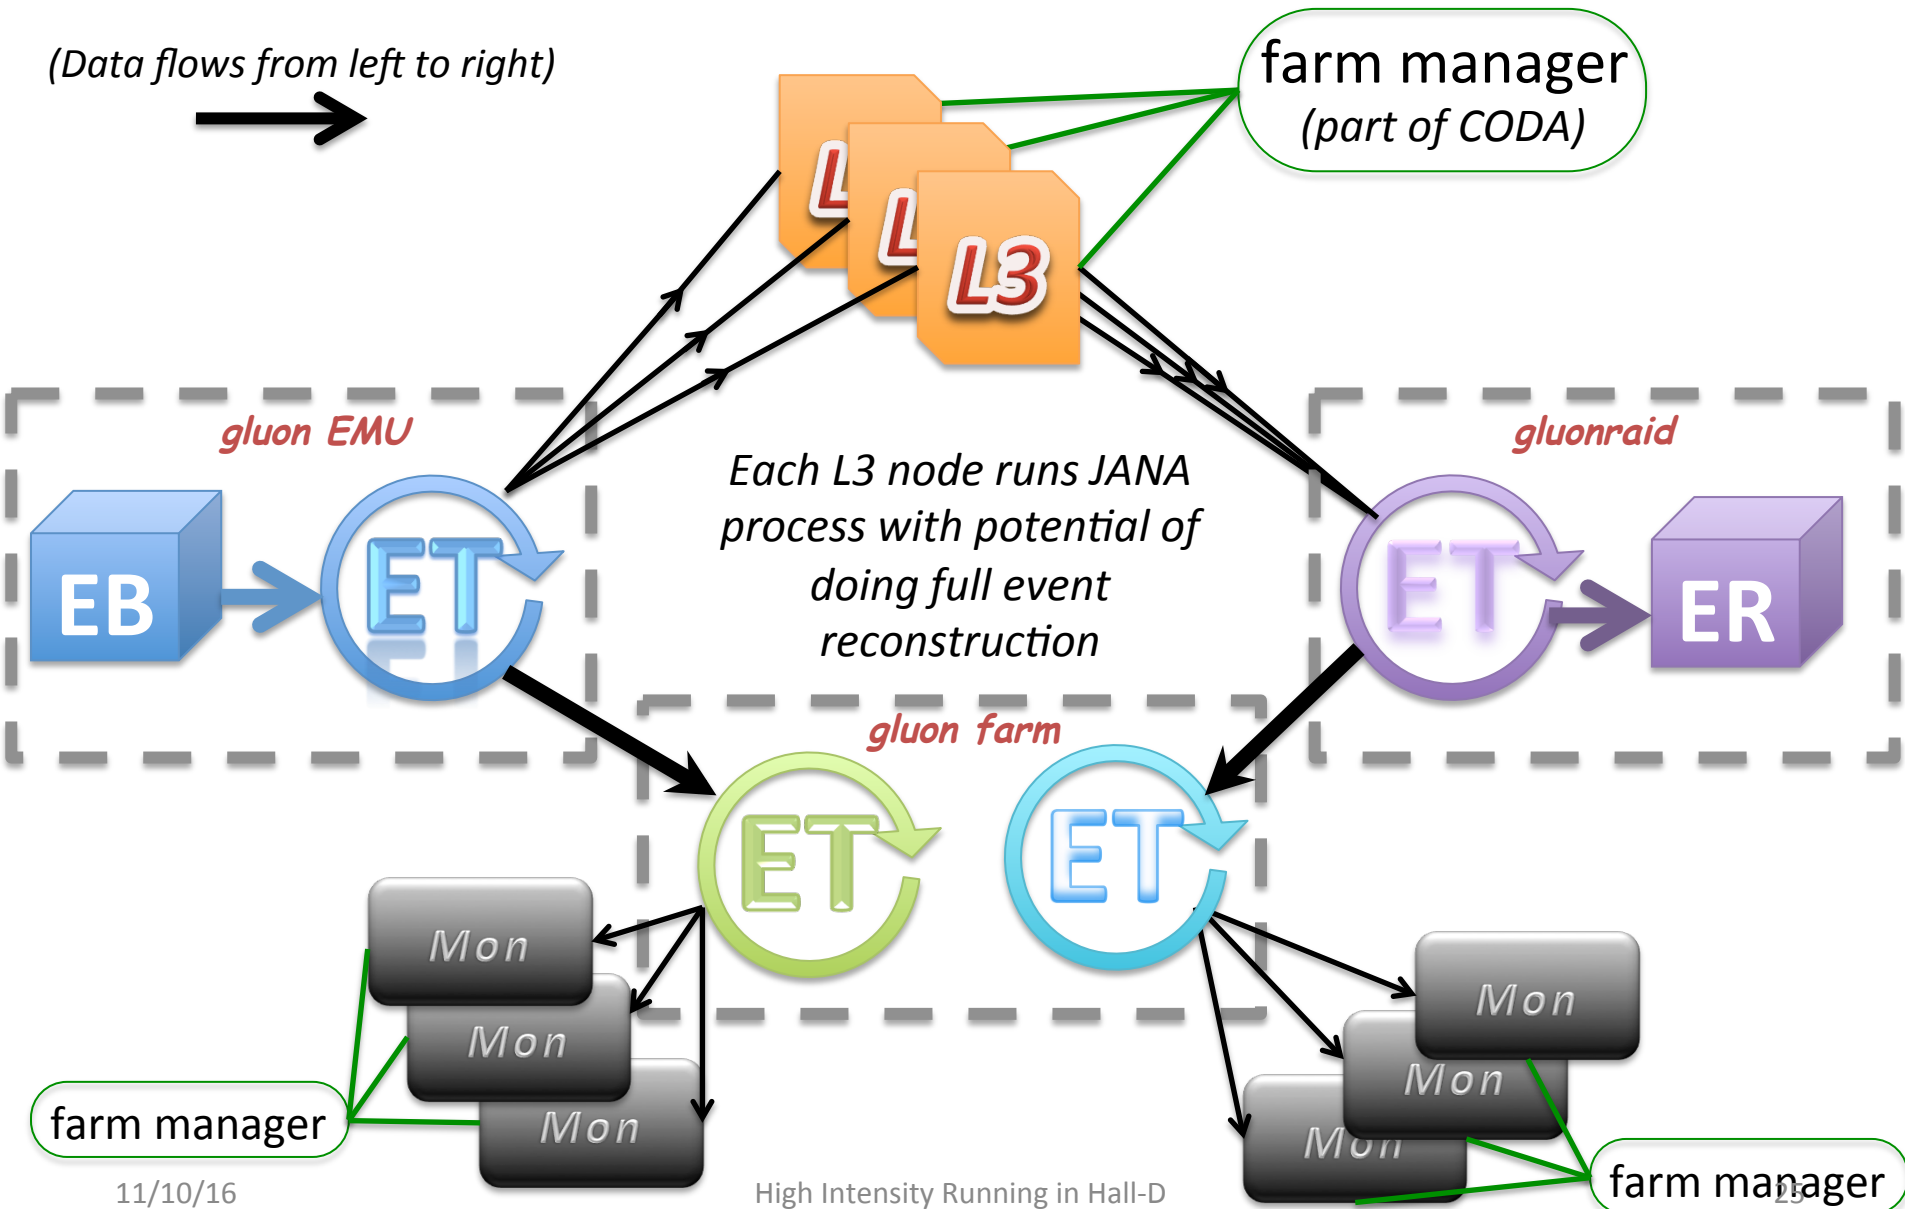
RAW data






hdl3 data

L3 and monitoring architecture

(Data flows from left to right)






Fitting event rate vs. time

$$R(t) = R_o (1 + Q/t)$$


R_o = asymptotic rate

Q = relaxation term

14 Workers 18 Processors




32 Workers 32 Processors




Parsing only

(no linking, no reconstruction)

Steady state rate

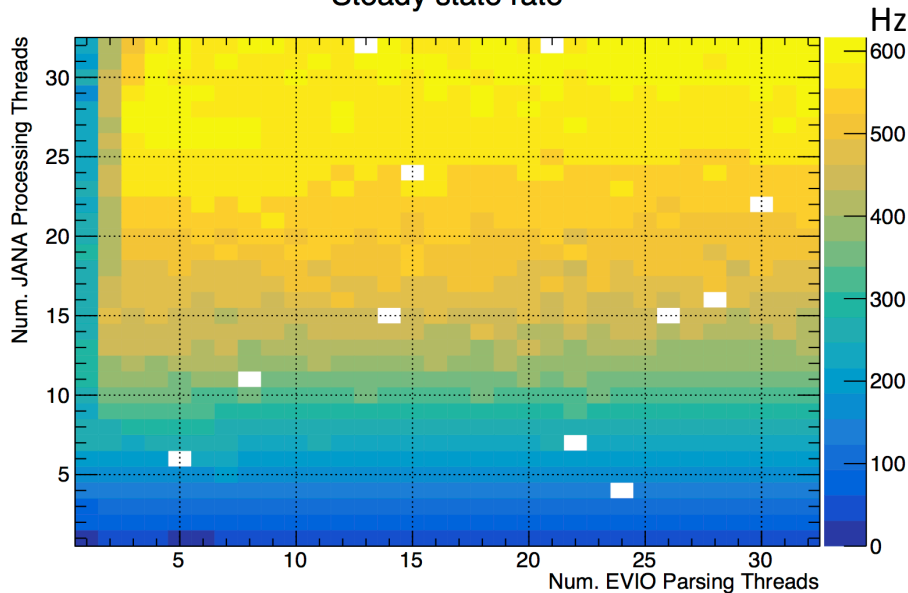


Relaxation term




Event rates with tracking

Wire-based Tracking
Steady state rate



Time-based Tracking
Steady state rate




- With associated object linking, parsing threads run about 4 times slower
- Single parsing thread with full linking: $\sim 250\text{Hz}$

EVIO Parsing Time

Rate (kHz)	Time/core/event (ms)	Condition
2.5	8.0	All linking enabled
2.9	6.9	All linking except TriggerTime
3.8	5.3	All linking except BORConfig
3.0	6.7	All linking except Config
4.8	4.2	All linking except TriggerTime and BORConfig
5.9	3.4	Hit linking only
8.0	2.5	No Linking


Old Parser Processing Rates

May 7, 2016 DL
git revision a347b7
hd_rawdata_011667_135.evio



New Parser Processing Rates


May 7, 2016 DL
git revision a347b7
hd_rawdata_011667_135.evio



New Parser Processing Rates

May 7, 2016 DL
git revision a347b7


hd_rawdata_011667_135.evio



Event Parsing Rate

May 3, 2016 DL
git revision #8e65dc3


hd_rawdata_011667_135.evio



DTrackTimeBased Rate

May 7, 2016 DL
git revision a347b7


hd_rawdata_011667_135.evio



Event Parsing Memory Usage

May 3, 2016 DL
git revision #8e65dc3


hd_rawdata_011667_135.evio




DTrackTimeBased Memory Usage

May 7, 2016 DL
git revision a347b7


hd_rawdata_011667_135.evio




Cache references/event




Cache misses/event



Context Switches/event



Page Faults/event




Input Test File

- `hd_rawdata_011667_135.evio`
 - 120nA, 50 μ m diamond (PERP), 5.0mm collimator
 - $I_{\text{solenoid}} = 1345\text{A}$
 - 18kB/event
 - Measured I/O rate: $\sim 900\text{MB/s}$ (=50kHz)
 - `fspeed_reader`
 - `gluonraid2 -> gluon48`
 - Maximum sim-recon read speed: $\sim 33\text{kHz}$
 - Parsing and linking disabled

L3 good event flag

July 7, 2016 DL
git revision #7118b85




Counting house computer systems


Computer(s)	processor	General Purpose Network	DAQ Network	I.B. Network	comments
gluonfs1	N/A	X			~1.6TB with snapshot backup
gluonraid1-2	Intel E5-2630 v2 @2.6GHz	X	X	X	RAID disk host ER process
gluon01-05	i5-3570 @3.4GHz	X			Shift taker consoles
gluon20-23	AMD 2347	X			Controls 8core
gluon24-30	E5-2420 @1.9GHz	X			Controls (gluon24 is web/DB/cMsg server) 12core + 12ht
gluon40-43	AMD 6380	X	X	X	16core + 16"ht"
gluon46-49	E5-2650 v2 @2.6GHz	X	X (gluon47 &49)	X	16core + 16ht
gluon100-111	E5-2650 v2 @2.6GHz	X		X	16core + 16ht
rocdev1	Pentium 4 @2.8GHz	X			RHEL5 system for compiling ROLs for DAQ
hdguest0-3	i5-3470 @3.2GHz	X (outside network)			Guest consoles in cubicles (outside network)

New farm nodes

OLD GLUON FARM NODE




NEW GLUON FARM NODE



- 8 new Broadwell based farm nodes installed
- Previous farm consisted of 12 Ivy Bridge nodes
- Farm now has 2.5 times as much CPU as Spring

Raw Data Flow



4 **Each partition has same directory structure**

- `/.../active` Data written to “active”
- `/.../volatile` Data moved to “volatile” once partition is inactive
- `/.../staging` Data hard linked in “staging” for copy to tape

5 **Disk mapping + auto-deletion**


Old scheme:

- Entire RAID disk mapped via cron job
- Files deleted “by hand” just before switching disks


New scheme:

- Map only volatile directory
- Map only when data moved to volatile (via cron)
- Auto-delete files from volatile to ensure 16.4TB is free

Hall-D Online Monitoring Architecture

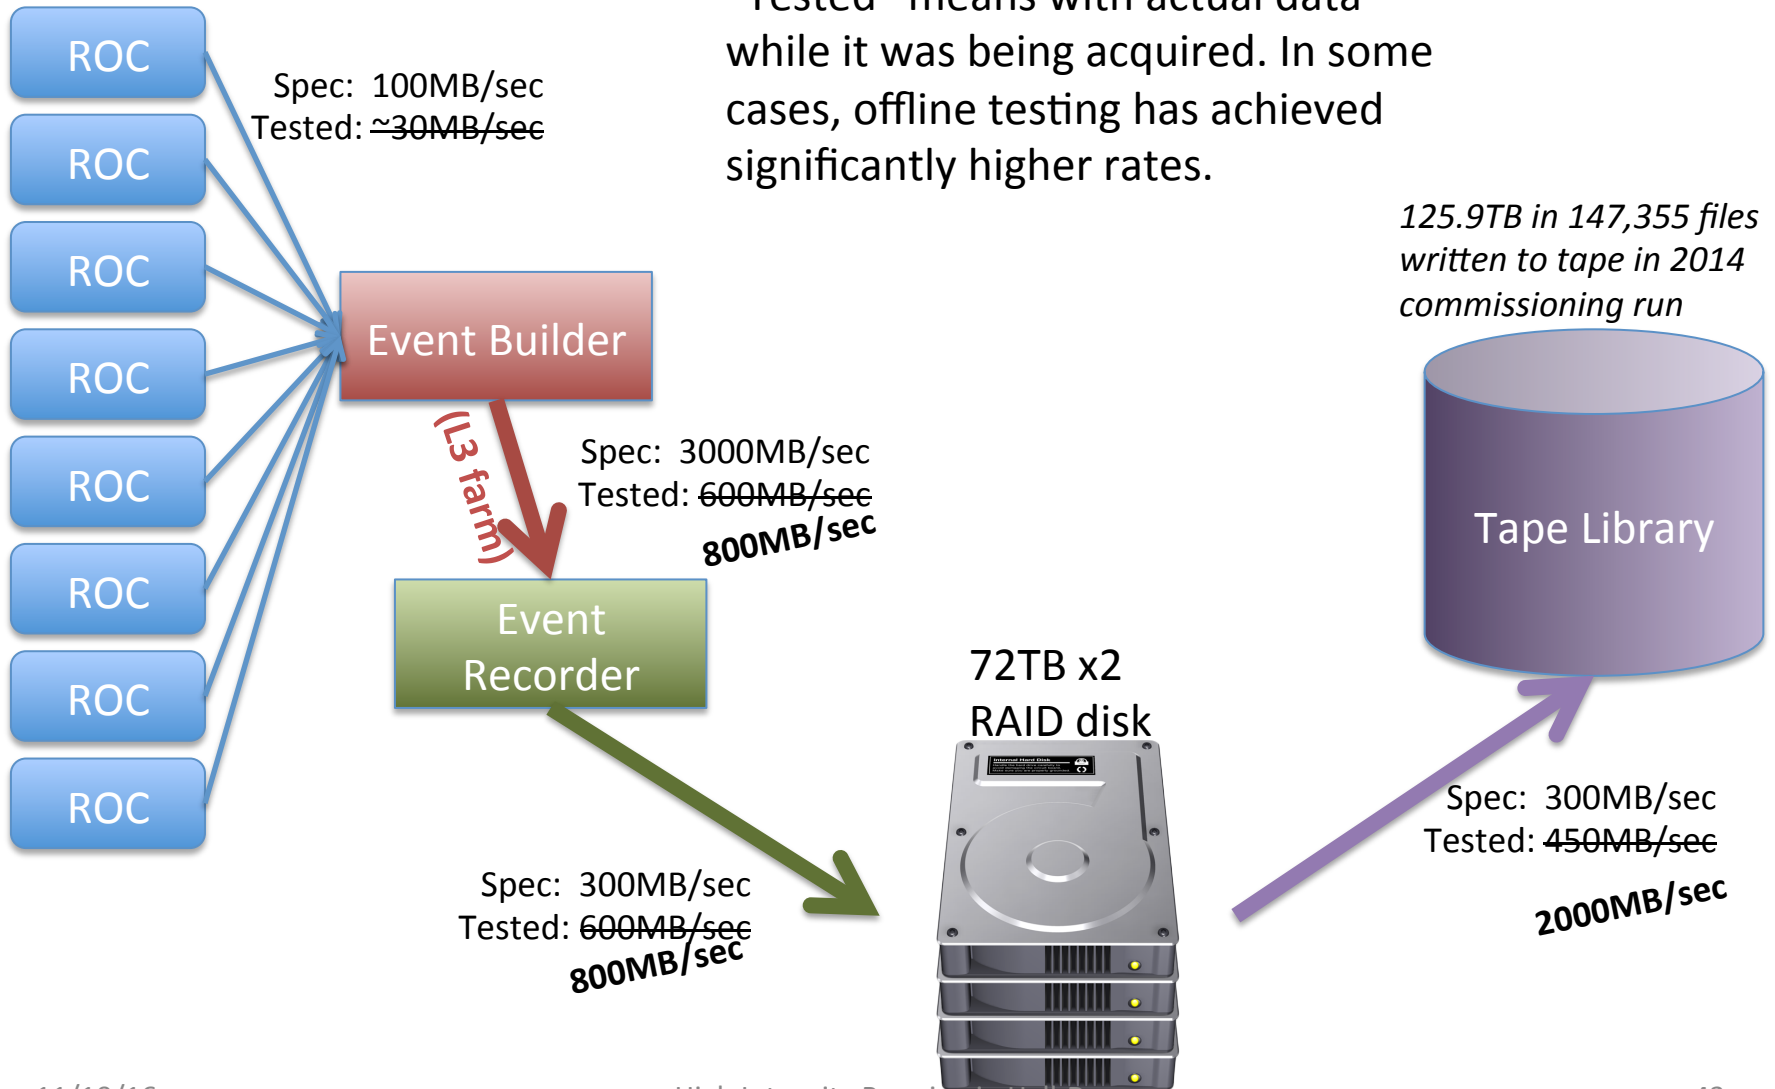


- Working on extracting L1 trigger rate from time between events
- Dropping events based on stricter L1 trigger thresholds would allow estimate of L1 trigger rate with these existing data files



Data Rates

“Tested” means with actual data while it was being acquired. In some cases, offline testing has achieved significantly higher rates.



Online Storage Capacity

- Two RAID disks with 72TB each of usable space
 - Maintain some portion of recent data
 - ~100TB effective space total for new data
- Need 72hr buffer in case of issue with link to tape library
- $100\text{TB} \div 800\text{MB/s} = 35\text{hr}$
- Need additional 100TB of RAID
 - Will purchase this summer

