

Hall-D L3 trigger status

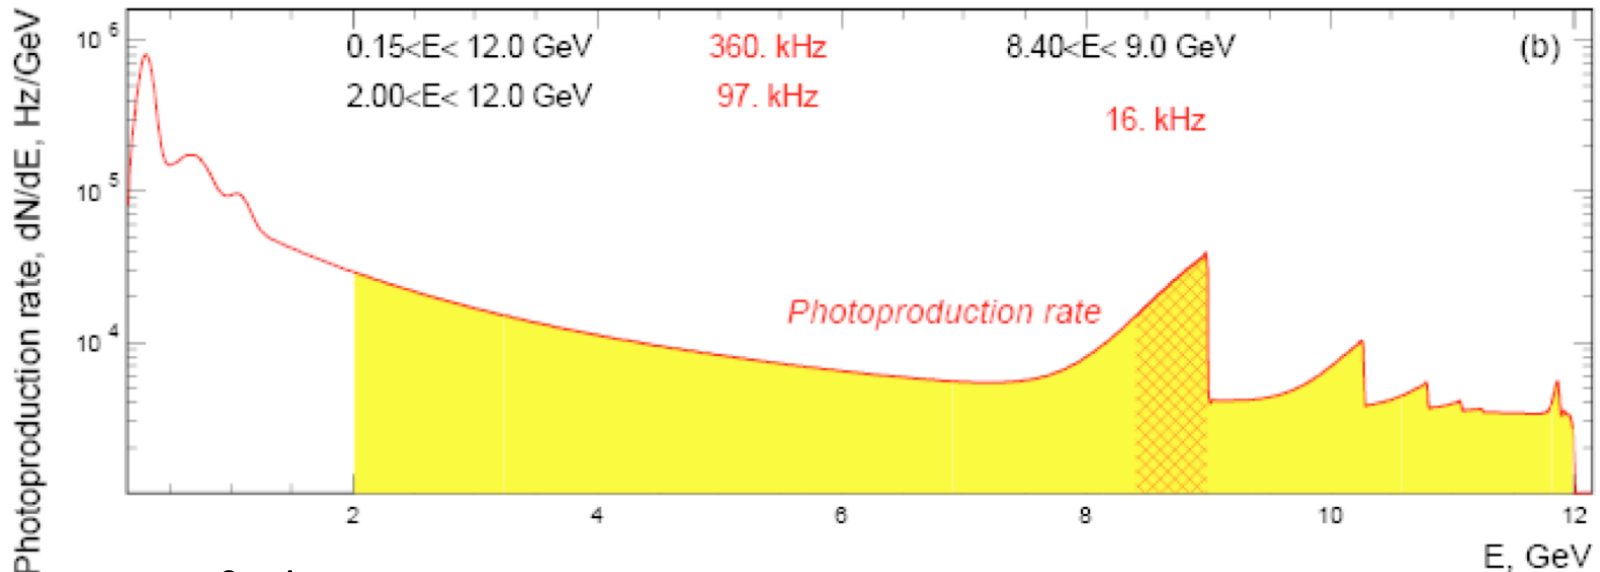


David Lawrence JLab

July 22, 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: $14 \text{ kB} + 0.05 \text{ kB}/nA^*$**
 - 15 kB/event @ 200 kHz = 3000 MB/s (front end)
 - L3 reduction by factor of 10 = 300 MB/s to RAID disk

L1 trigger rates

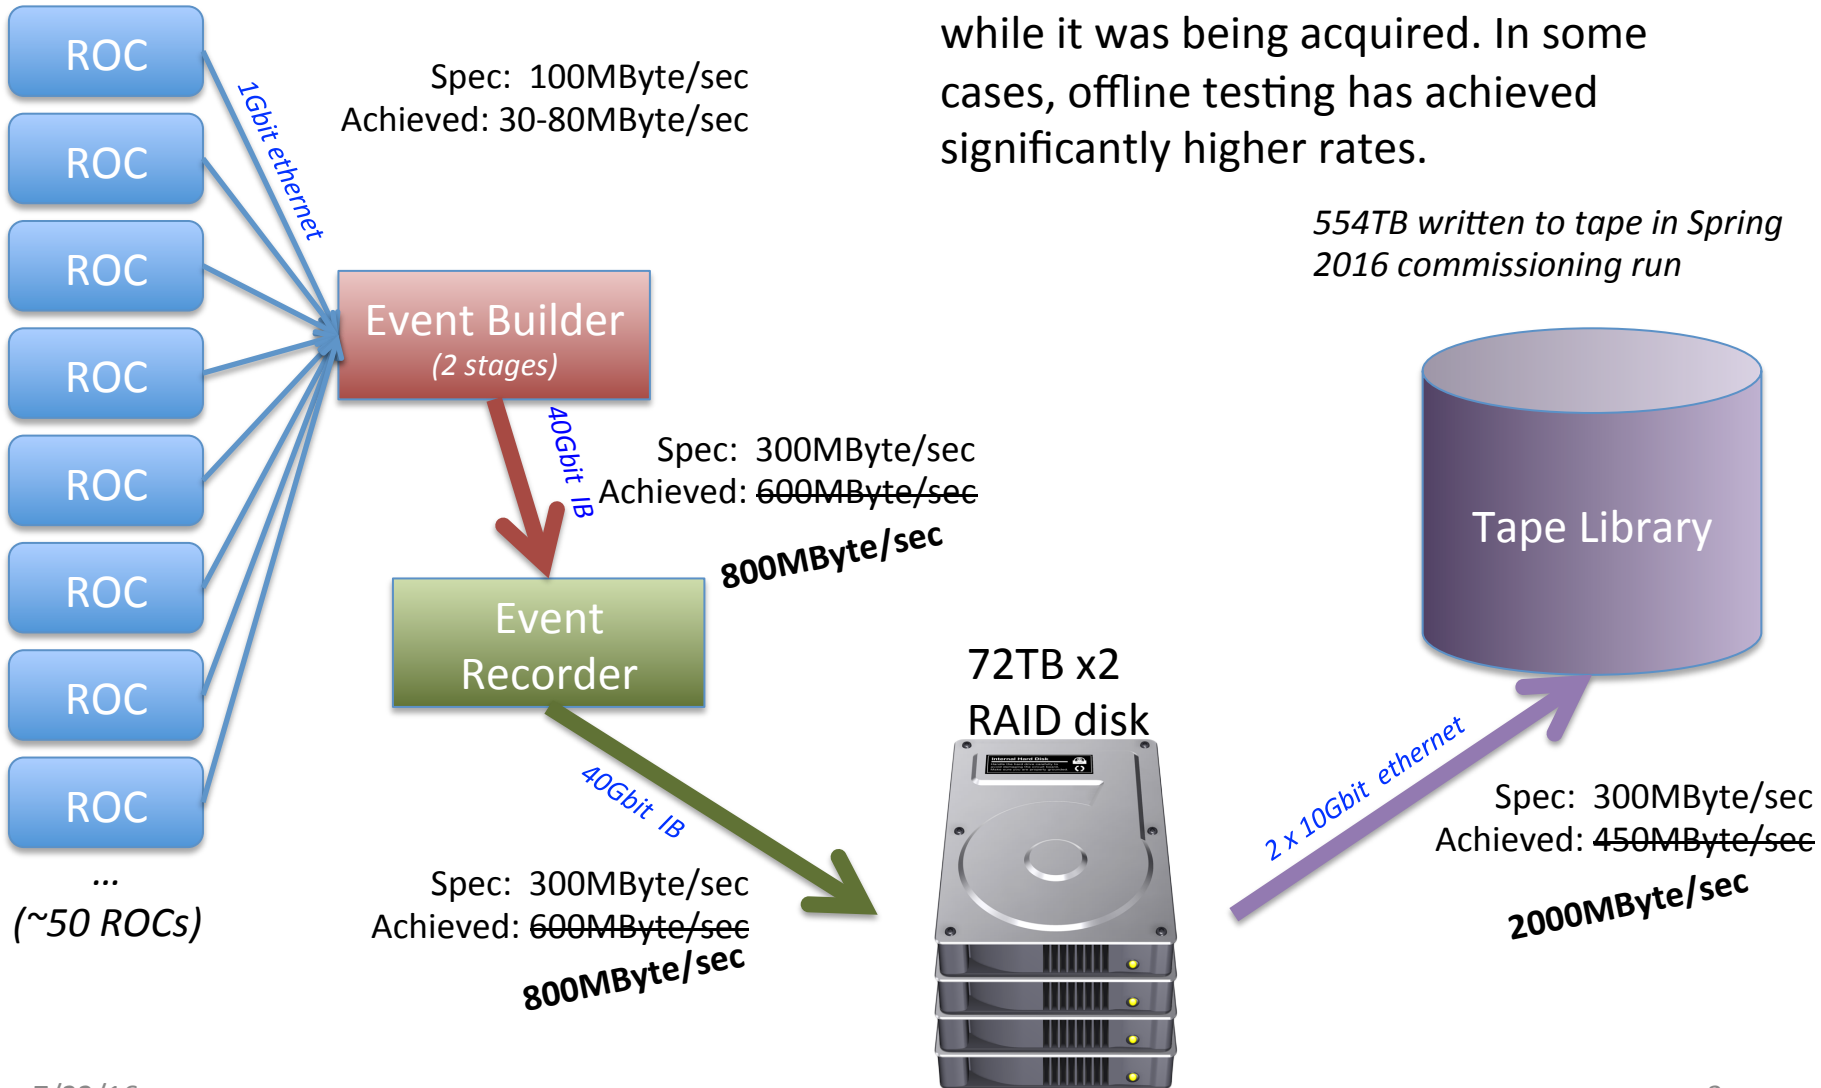
Fall 2014: 2kHz

Spring 2015: 3.5kHz

Spring 2016: 30kHz

Data Rates

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



554TB written to tape in Spring 2016 commissioning run

L1 trigger rates

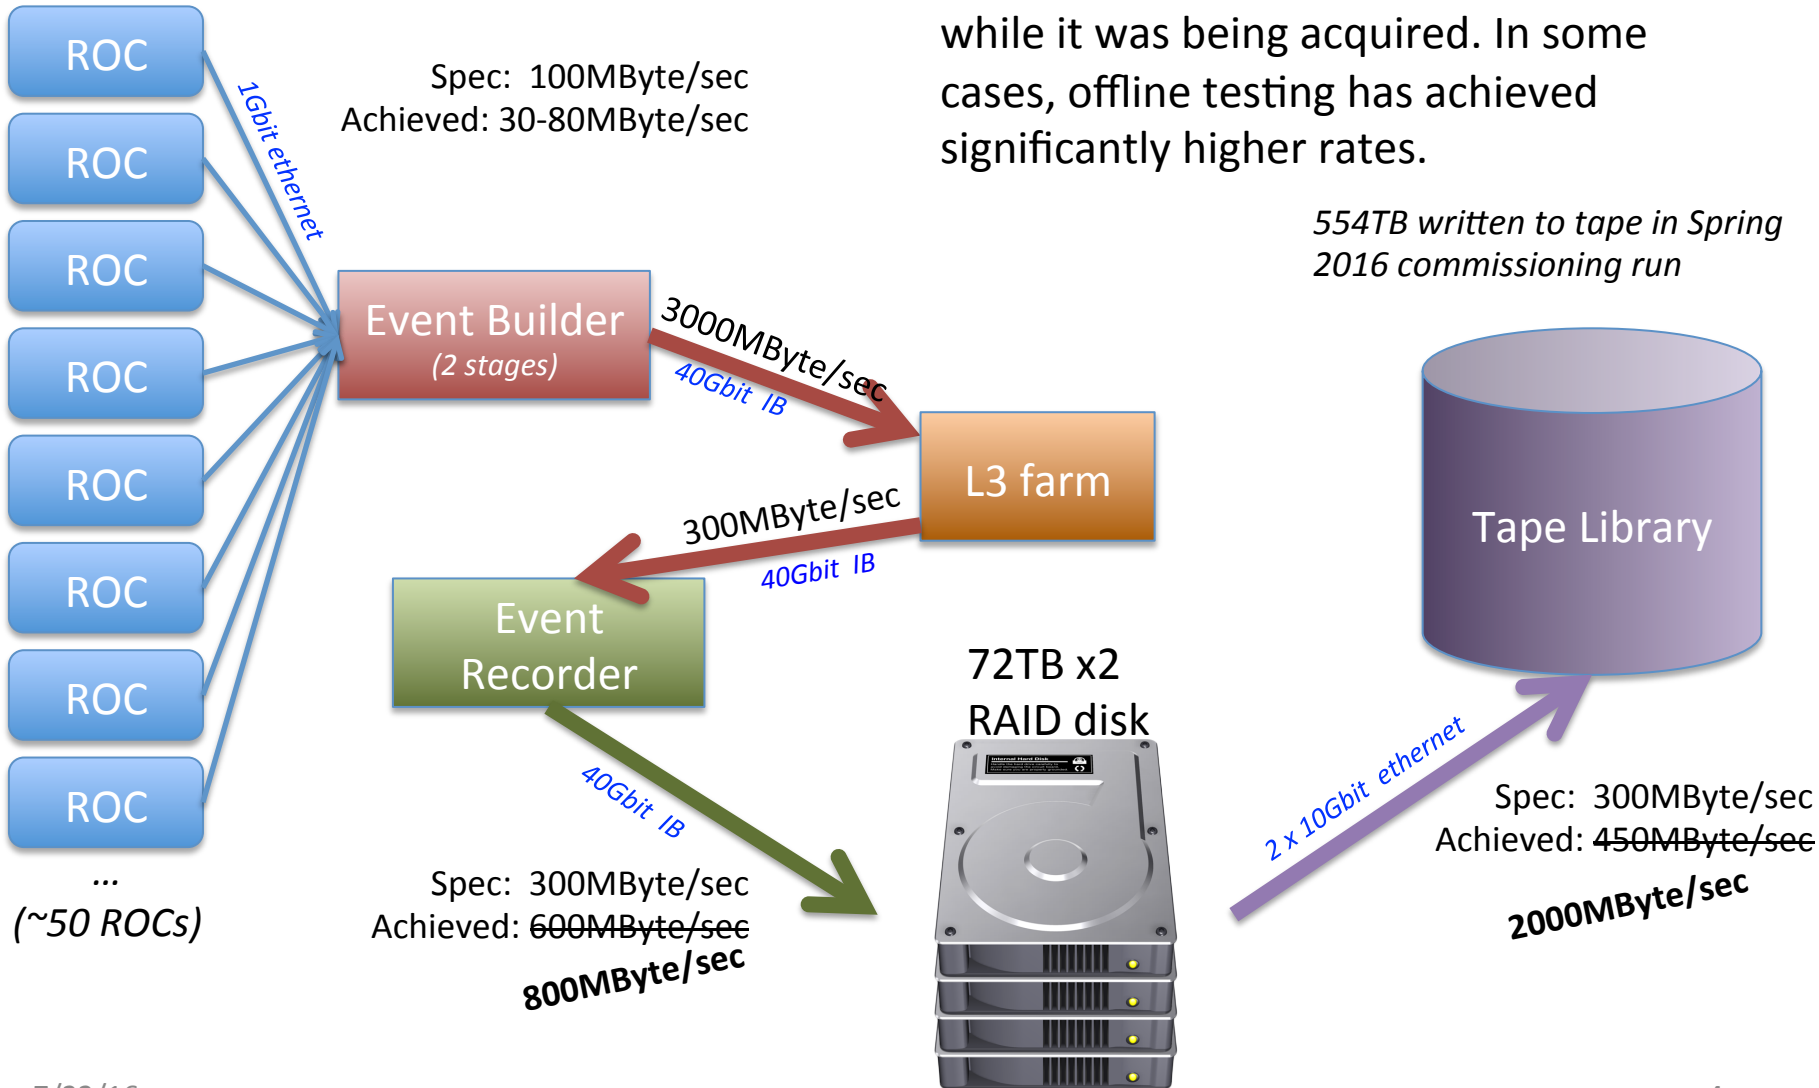
Fall 2014: 2kHz

Spring 2015: 3.5kHz

Spring 2016: 30kHz

Data Rates

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



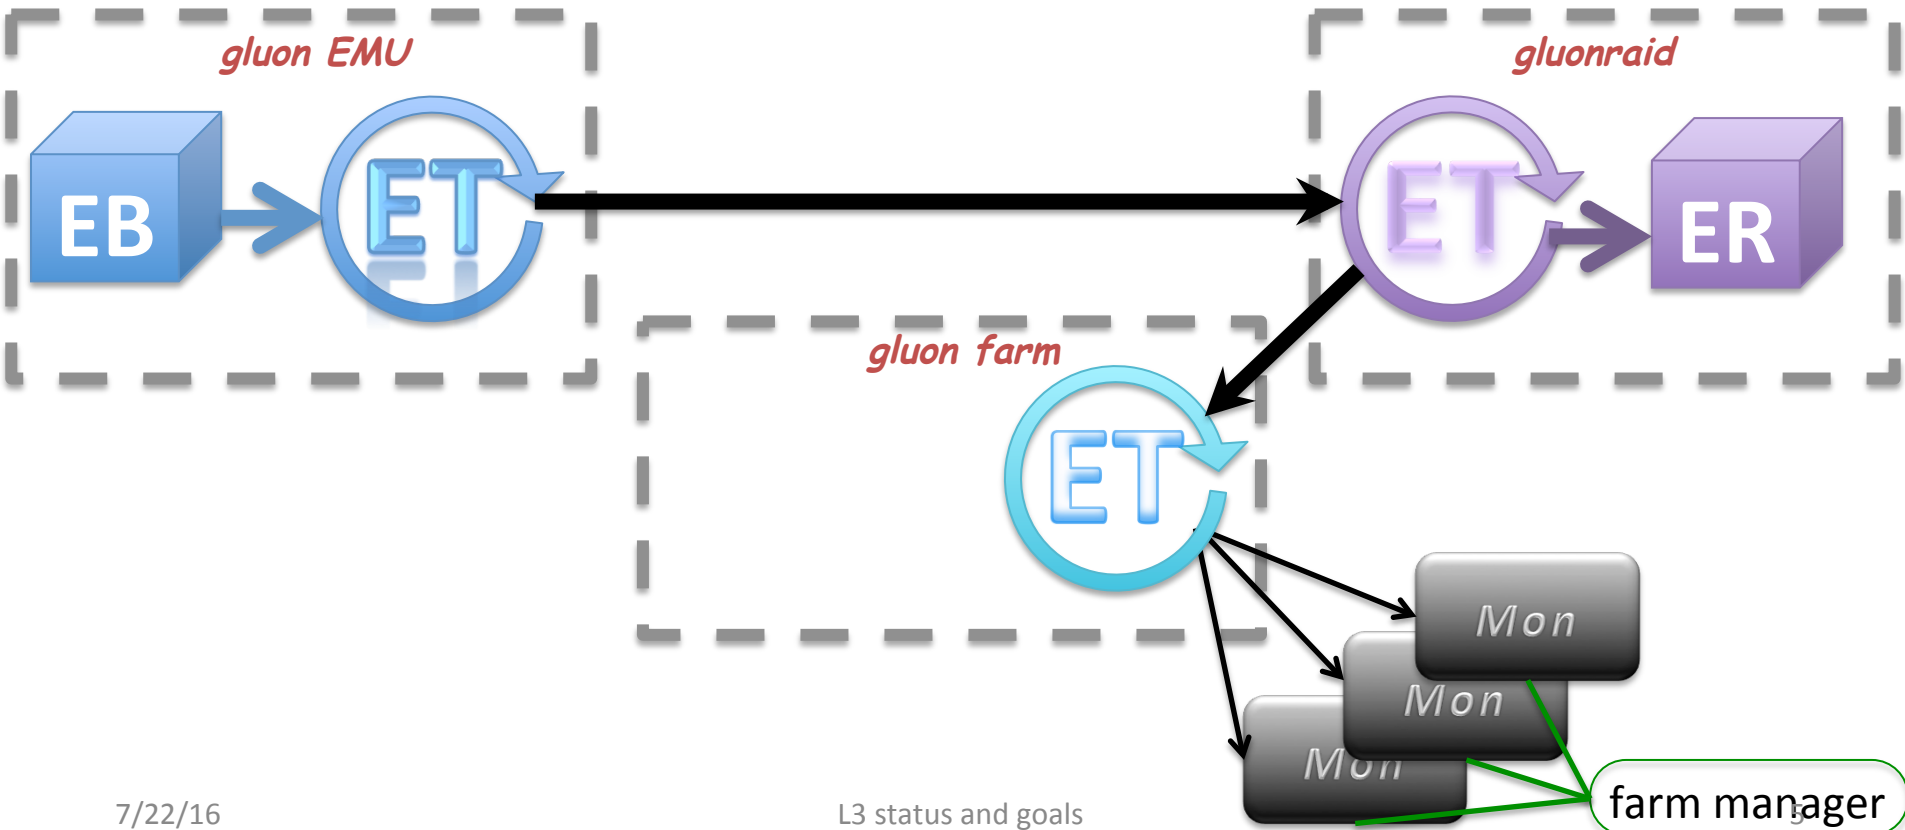
554TB written to tape in Spring 2016 commissioning run

L3 and monitoring architecture

(Data flows from left to right)



Spring 2016 Configuration

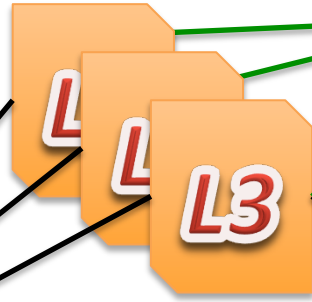


L3 and monitoring architecture


(Data flows from left to right)



farm manager
(part of CODA)

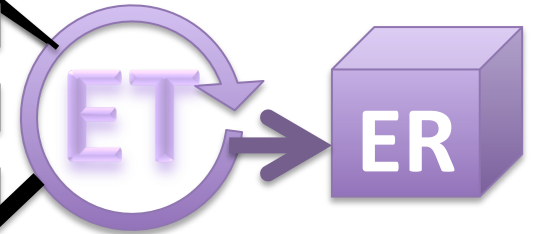


gluon EMU



Each L3 node runs JANA process with potential of doing full event reconstruction

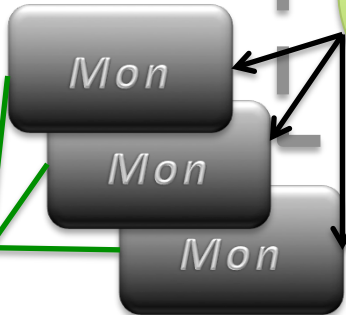
gluonraid



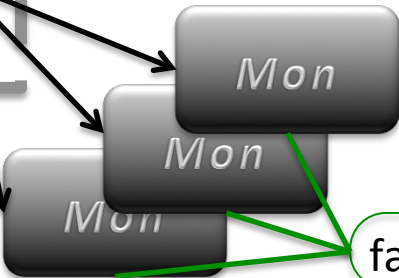
gluon farm



farm manager



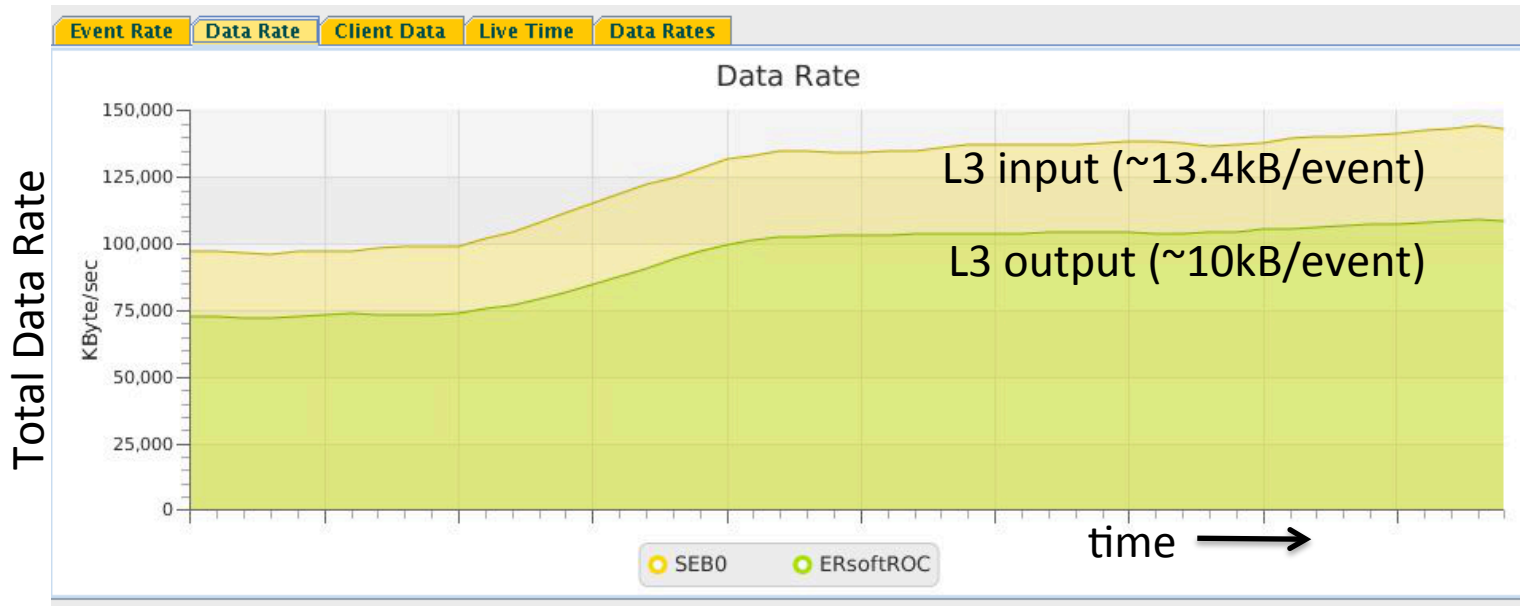
L3 status and goals



farm manager

L3 running in pass-through with beam

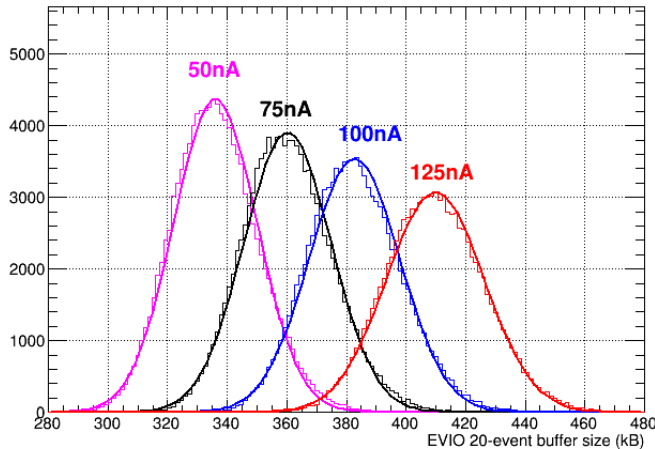
- Events are read in blocks of 20
- Events must be disentangled and reconstituted as single events before writing to disk
- Redundant headers may be dropped to reduce event size



Event Size vs. Beam Current

buffer size

March 14, 2016 DL
git revision #52de6c7
runs 10864-10867



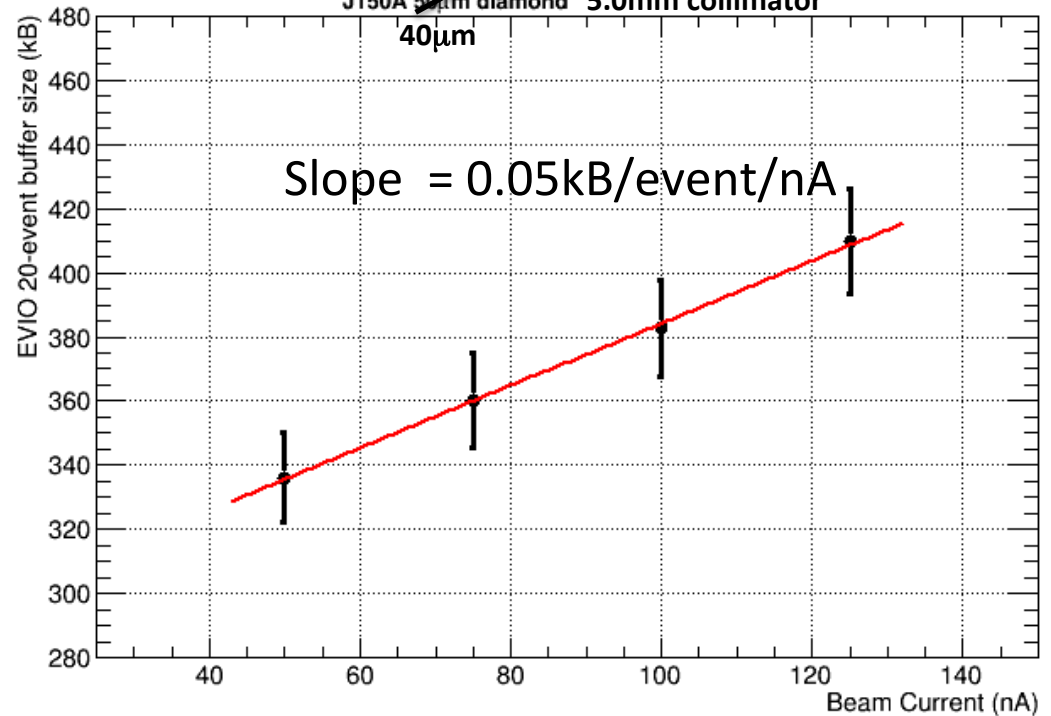
Event Size: ~19kB

40 μ m diamond
5mm collimator
120nA = 0.8x10⁷ γ /s
750nA = 5x10⁷ γ /s

buffer size vs. beam current

J150A 50 μ m diamond 5.0mm collimator

March 14, 2016 DL
git revision #52de6c7
runs 10864-10867



extrapolate to $I_{\text{beam}} = 0 \rightarrow 14.25\text{kB/event}$
(size of clean event with no accidentals)

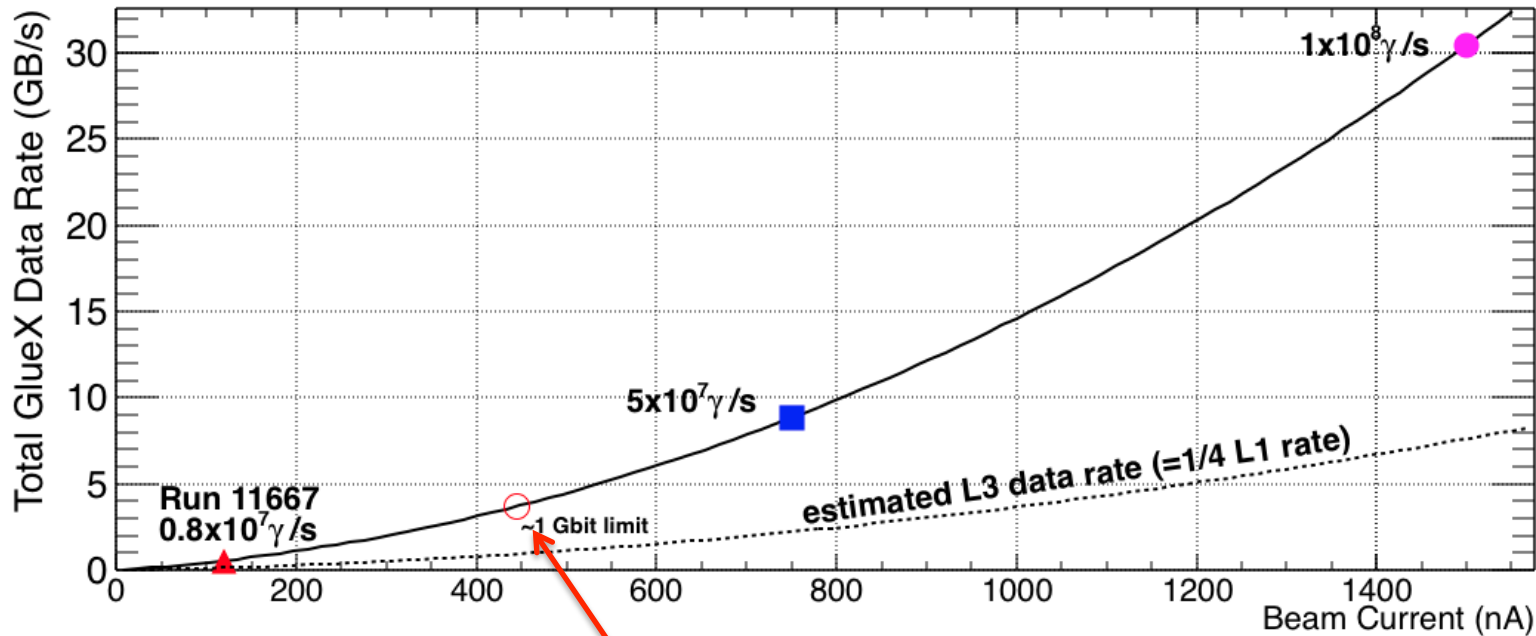
Accidental data fraction (by volume):

$$\frac{(0.05\text{kB/nA})(I_{\text{beam}} \text{ nA})}{(0.05\text{kB/nA})(I_{\text{beam}} \text{ nA}) + (14.25\text{kB})}$$

100nA: 26% of data is due to accidentals

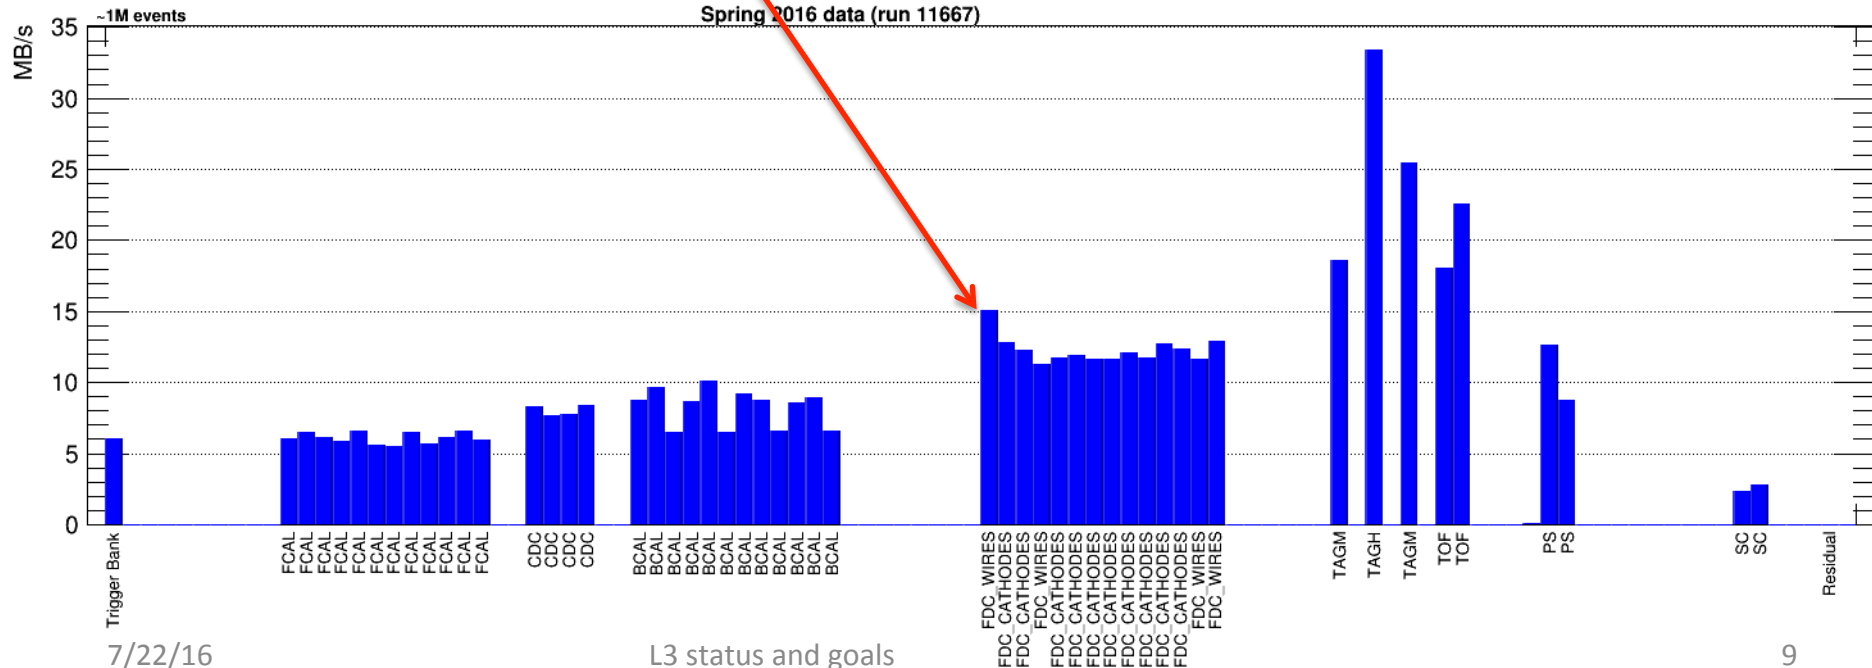
200nA: 41% of data is due to accidentals

750nA: 72% of data is due to accidentals



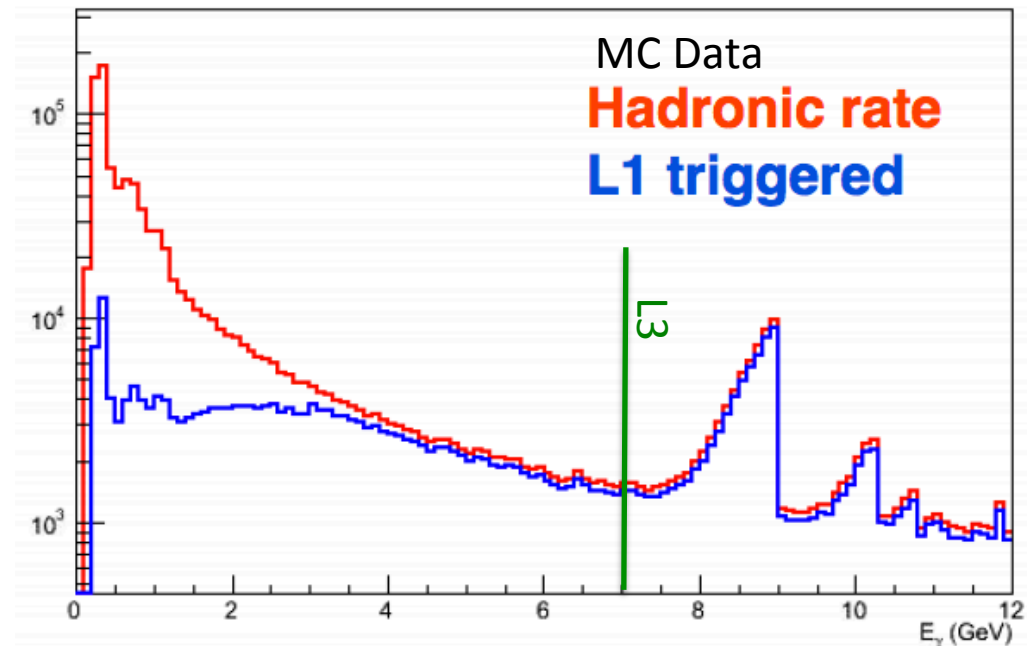
MB per sec by VME crate for run 11667

Spring 2016 data (run 11667)



L3 Algorithm Strategy

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



- Use Multivariate Analysis such as BDT or ANN to classify events
- 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

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

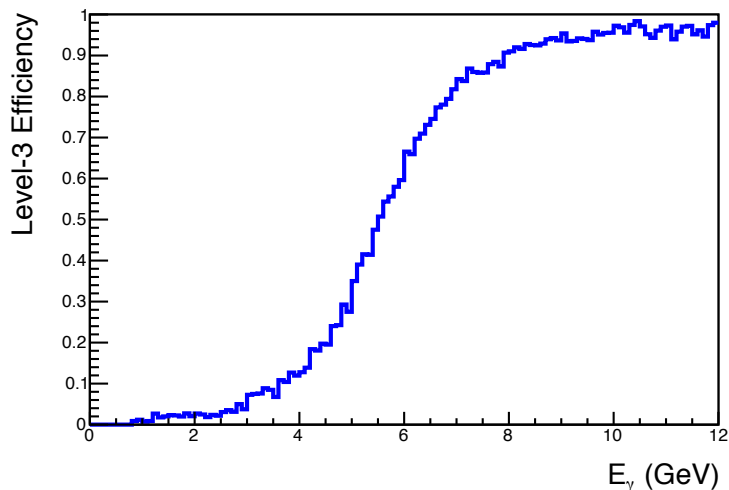
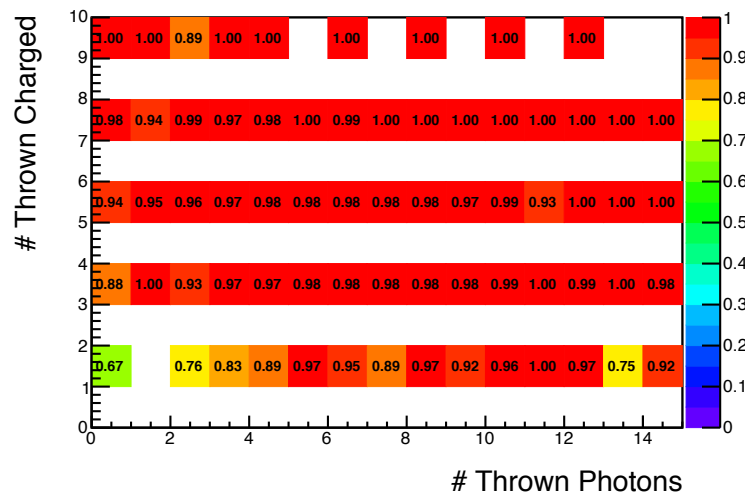
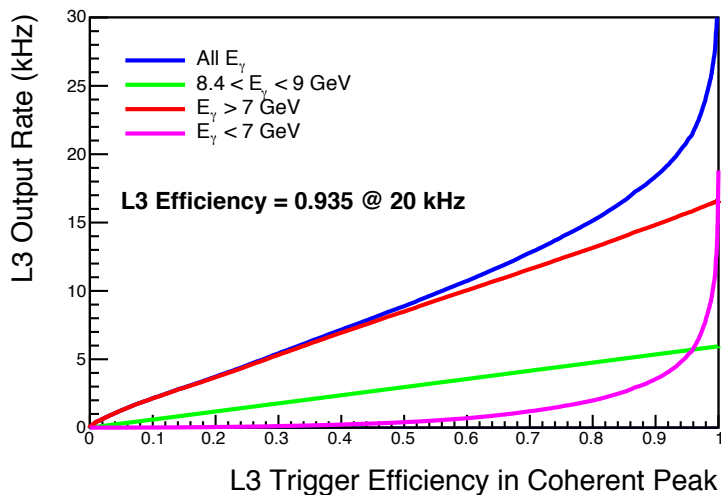
events	ALGORITHM	INPUT OBJECT	RECO TIME [s]	RECO TIME/event [ms]
10000	nominal reco	DNeutralShower	4.410	0.028
DATA	approx reco	DBCALShower	2.750	0.017
	approx reco	DFCALShower	8.010	0.050
	full tracking	DTrackTimeBased	18669.810	116.69
	approx tracking	DTrackWireBased	7397.300	46.23

hdgeant_smeared_14980

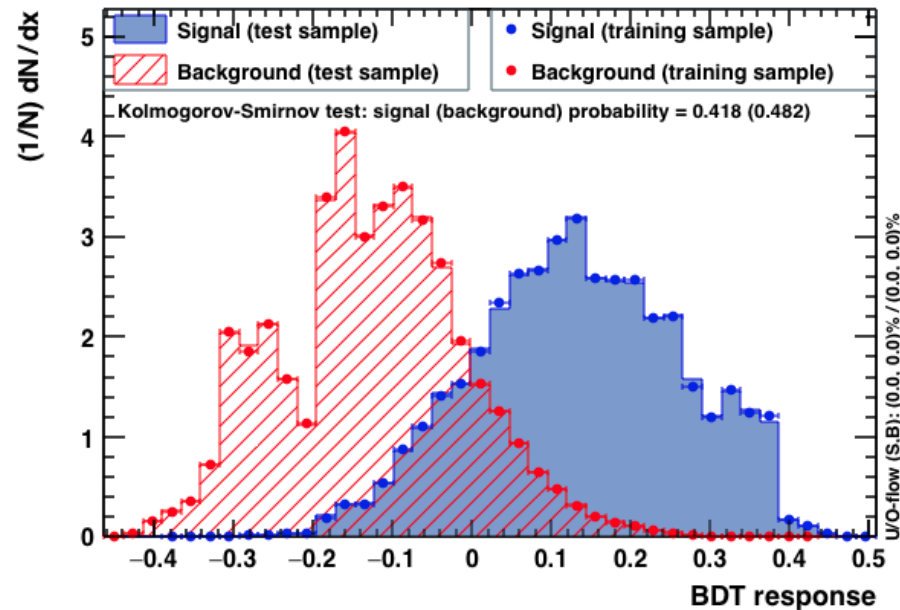
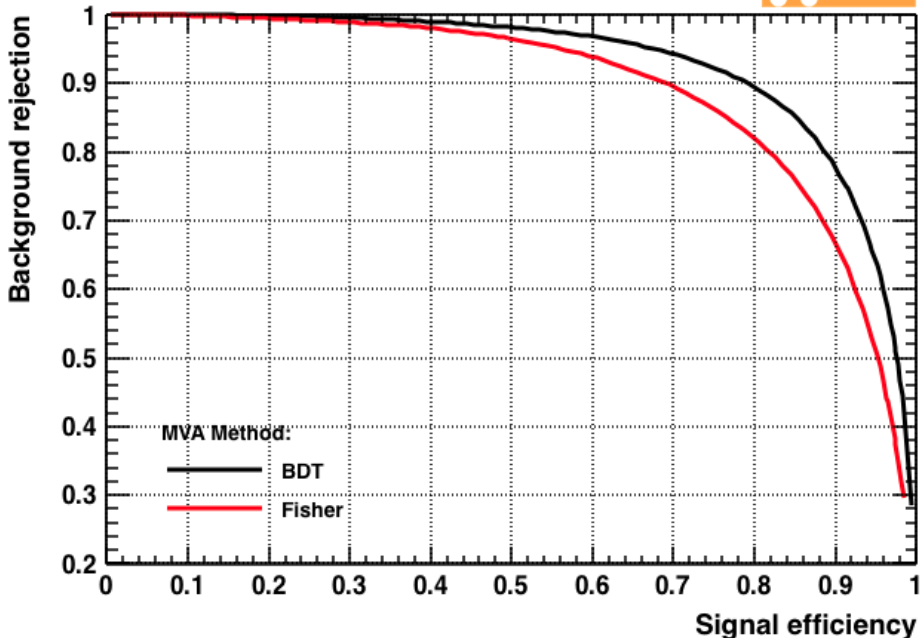
events	ALGORITHM	INPUT OBJECT	RECO TIME [s]	RECO TIME/event [ms]
10000	nominal reco	DNeutralShower	5.730	0.036
MC	approx reco	DBCALShower	3.430	0.021
	approx reco	DFCALShower	7.010	0.044
	full tracking	DTrackTimeBased	23878.840	149.243
	approx tracking	DTrackWireBased	12778.340	79.865

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

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



```

--- 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
  
```

\$SHALLD_HOME/src/plugins/Utilities/I3bdt

Estimated number of L3 nodes

- From Spring 2016 running: $\sim 30\text{kHz}/0.8 \times 10^7 \gamma/\text{s}$
- For $5 \times 10^7 \gamma/\text{s}$: $\sim 190\text{kHz}$
- 2013 Ivy Bridge nodes
 - 2.5-13kHz parsing only
 - 4kHz parsing+neutrals recon
 - 2kHz final algorithm (*rough estimate*)
- Newer nodes assume x2 faster
 - 4kHz/node
- Total number of nodes required:
 - $190\text{kHz}/4\text{kHz} = \mathbf{48}$



Schedule



- Continue L3 development, testing infrastructure and MVA efficiency through Fall
- Integrate testing schedule into Fall 2016 run
- If confident, take portion of Spring 2017 data in pass-through mode
 - Events on tape will be reconstituted single events
 - Event tagging may be used if fast algorithm available
- Full L3 deployment Fall 2018

Potential Issues



- fADC125 high rate performance
 - Busy signals (*currently being tested in Hall-D*)
- Individual VME crate saturates 1Gbit link
 - 10Gbit ethernet cards in ROCS
 - Split crates (*either backplanes or add additional crates*)
 - Parallel module readout (*VXS, or new fADC125 daughters*)
- 50 ROCs pushing 9GB/s through full CODA system
 - Minimize single crate data size to allow high luminosity testing
 - May require tuning IB network or splitting streams

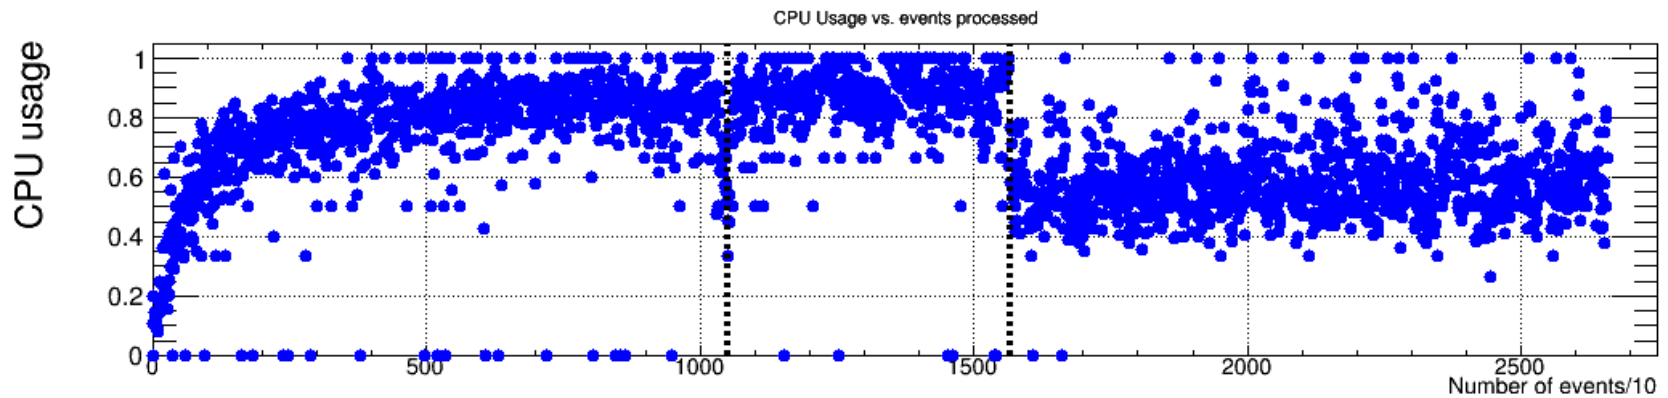
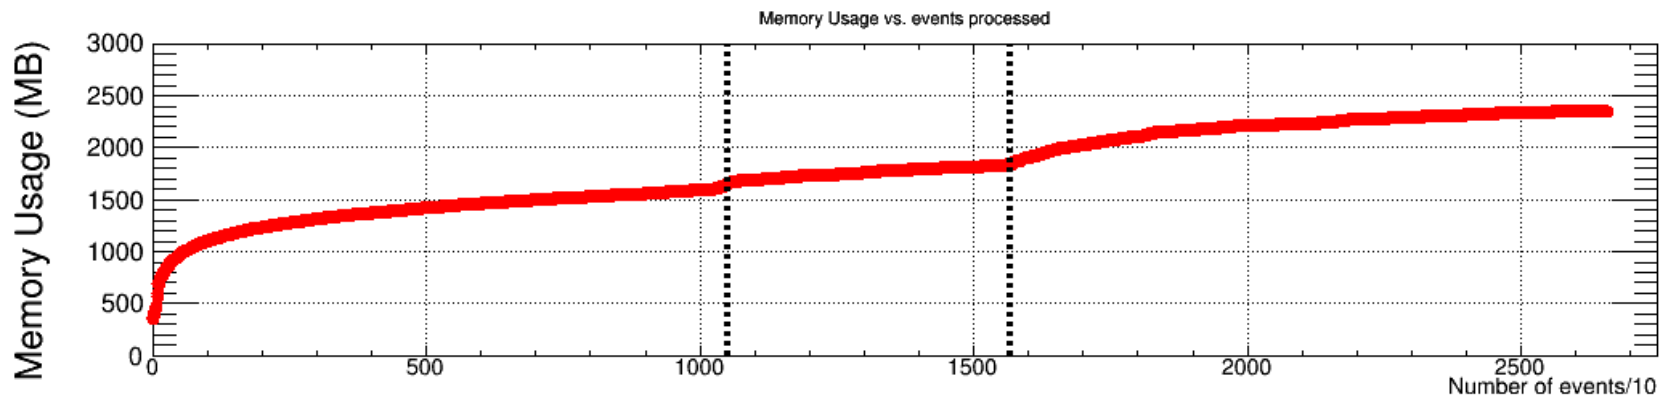
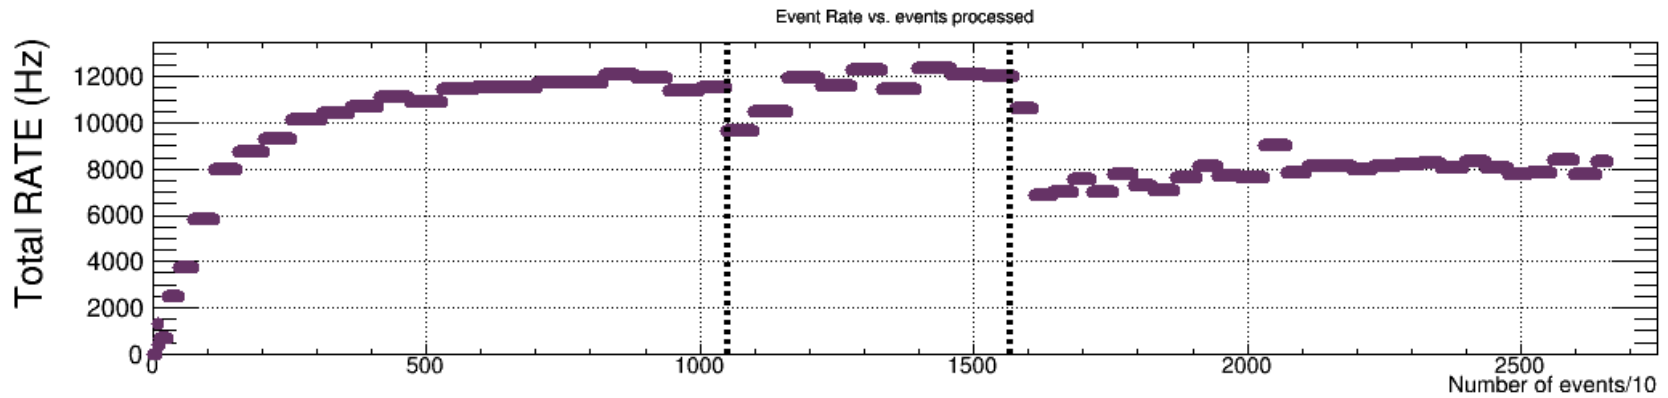
Backup Slides



Detectors At High Rate

- FDC aging effects for wires closest to beamline
- FCAL innermost blocks radiation damage
 - Curing
 - Replace with rad-hard insert

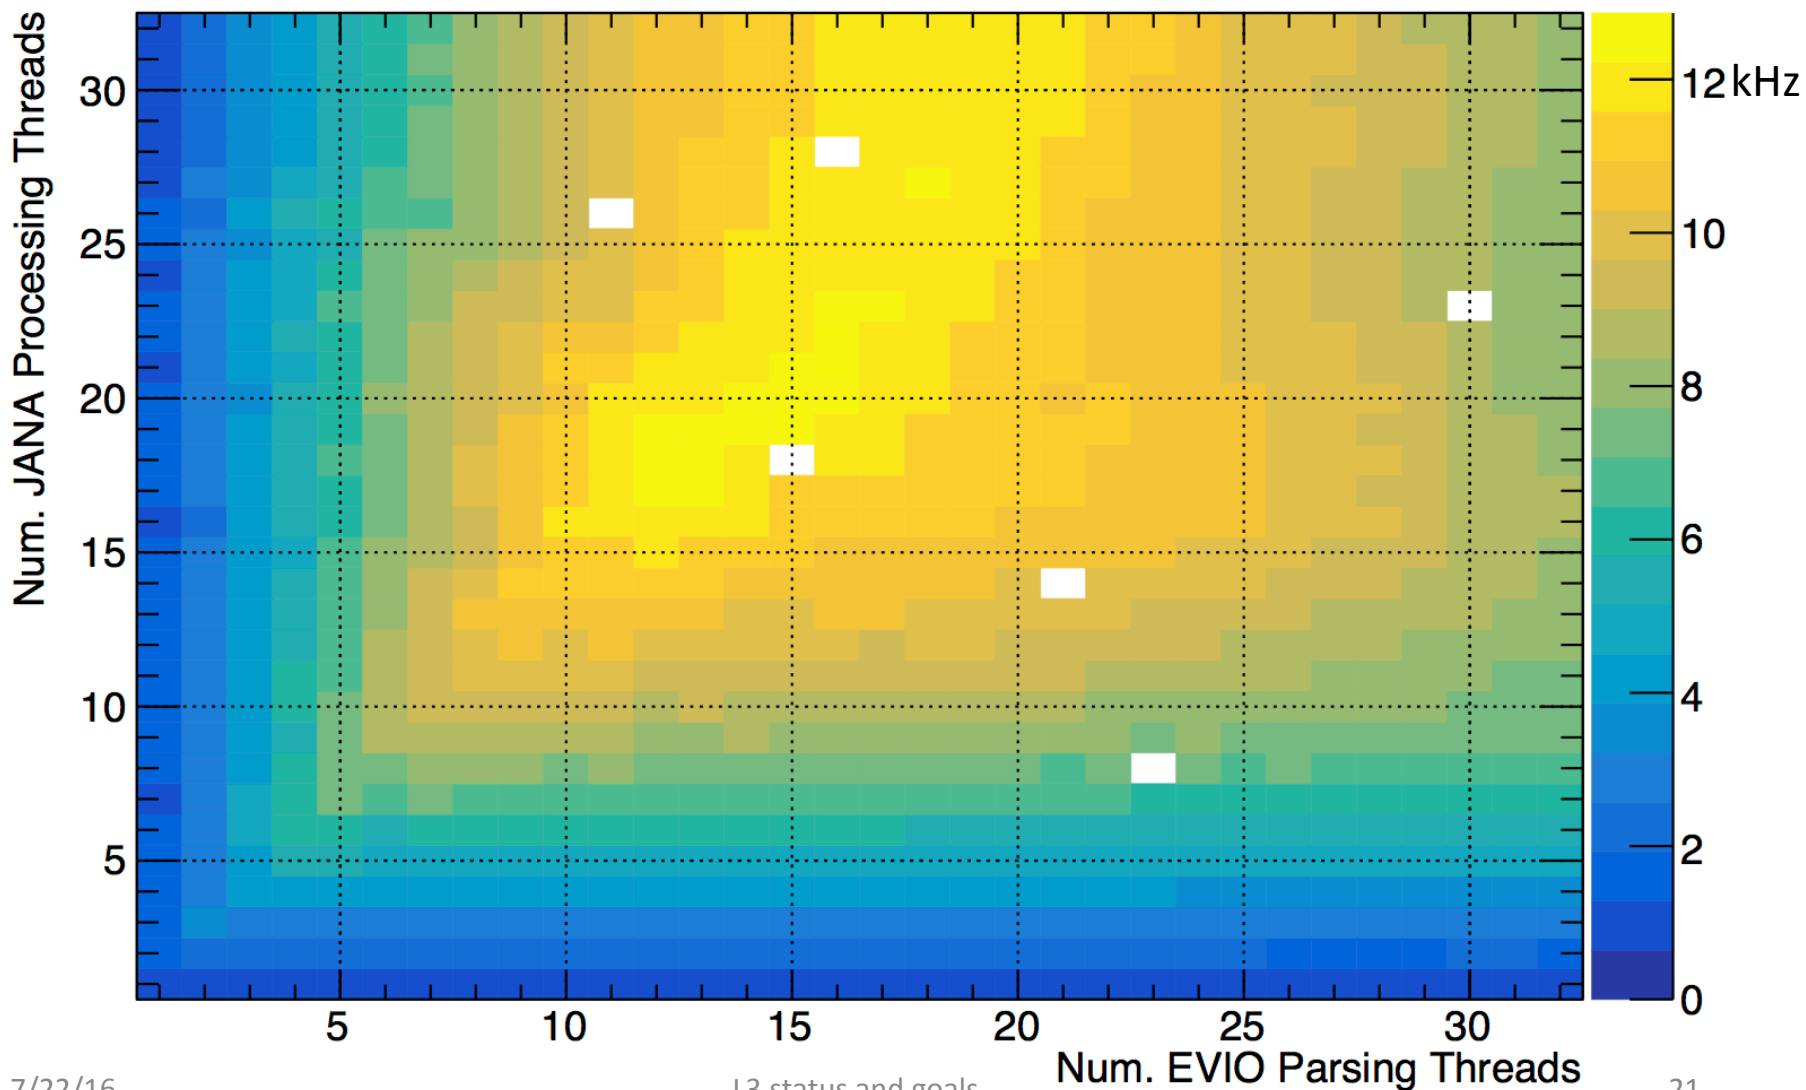
- Potential concern for the CDC and FDC is the aging of the wires: high currents in the chambers which, especially in presence of organic contamination, can result in formation of deposits on the wire surface (like polymers) and manifest as a decrease of the gas gain.
- The relevant quantity is the charge per unit wire length. For the nominal low intensity during the commissioning/ engineering running we had in the inner FDC HV sector (20 wires) about 5 uA currents, which however is concentrated to the area closest to the beam line, where we expect to have up to 1 uA/cm. In the CDC the currents on 24 inner straws was about 10 uA, but distributed evenly along each wire and between the 24 wires, corresponding to current densities of about 5 nA/cm.
- No wire aging effects at all have been observed in Ar/CO₂ gas mixtures in studies with up to 1C/cm, which however corresponds to only ~12 beam-on-target days at low intensity. Up to now (??total charge/cm??) we have not seen any wire aging. Even wire aging is possible at higher intensities, it will be concentrated in a small area around the beam line.
- Conclusions: don't expect problems with the CDC; in FDC aging is possible but only in a small area around the beam.
- In order to mitigate potential aging effects due to radiation we do add alcohol to the gas. This is proven to stop the progress of aging if already present and inhibit aging from happening in the first place. As an example HallC operated their wire chambers for more than a decade now using alcohol in the gas mix and to my knowledge no aging has been observed. (This statement should be confirmed by Howard Fenker)
- One potential issue could be radiation damage to the FCAL lead glasses over the prolonged period of high-luminosity running.



Parsing only

(no linking, no reconstruction)

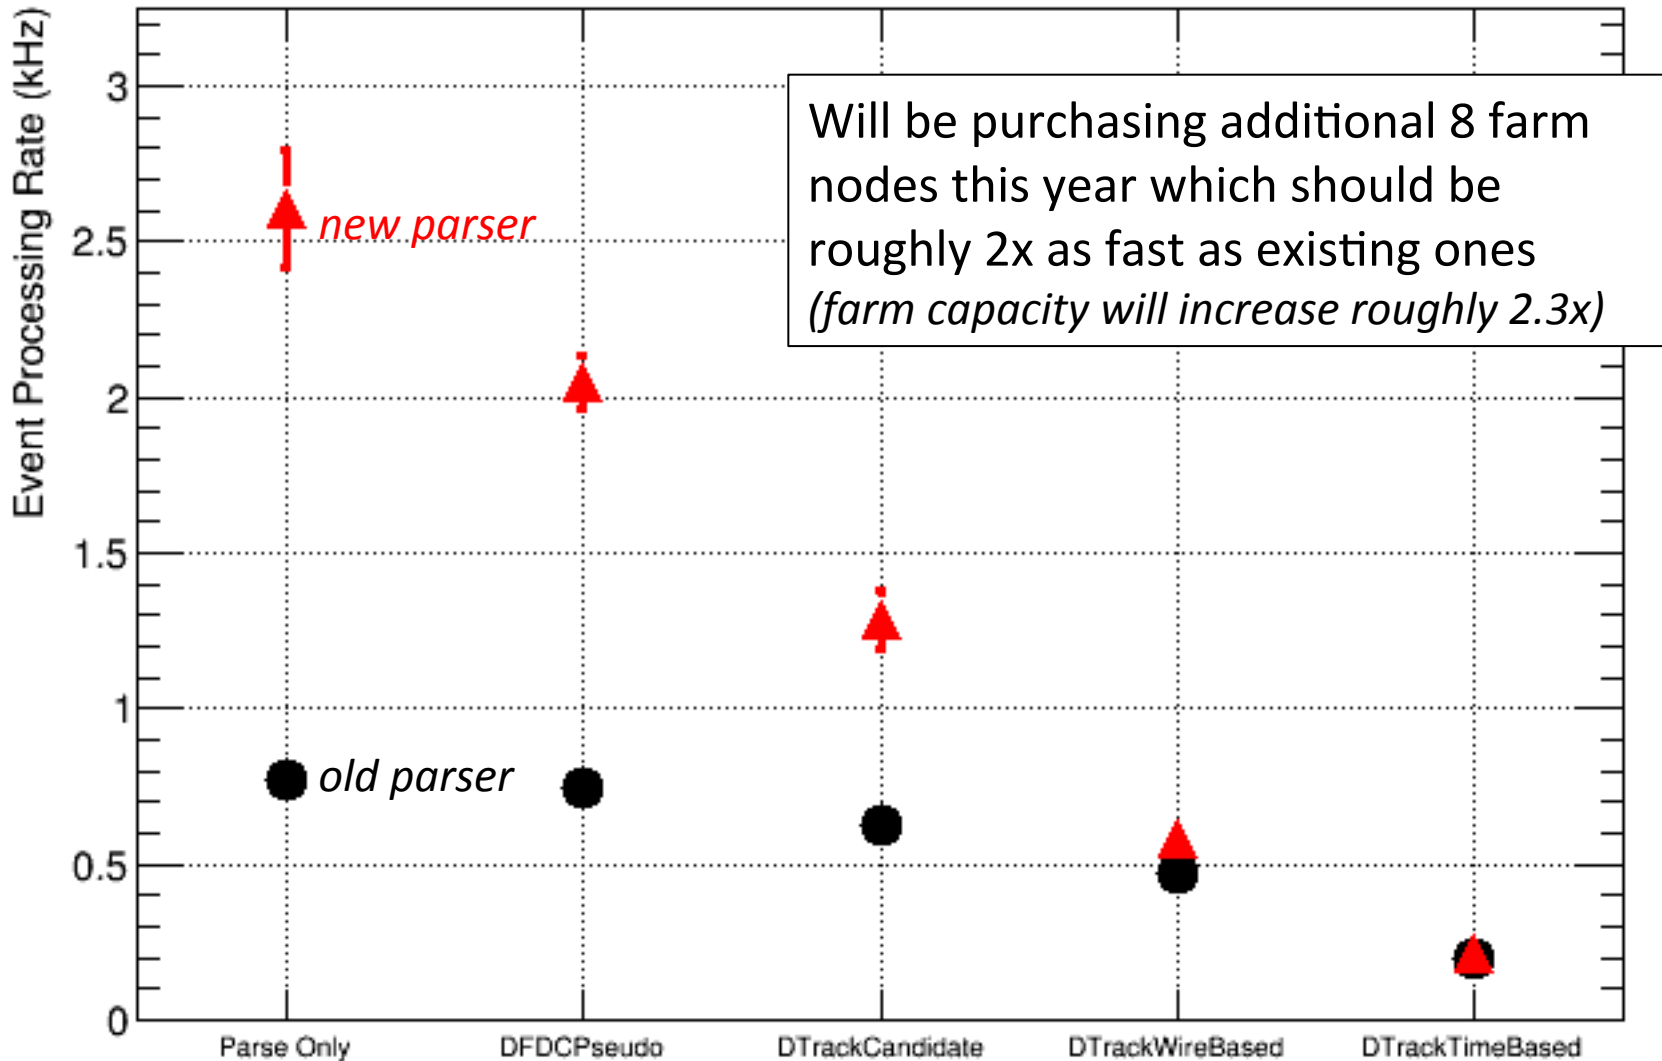
Steady state rate



New Parser Processing Rates

May 7, 2016 DL
git revision a347b7

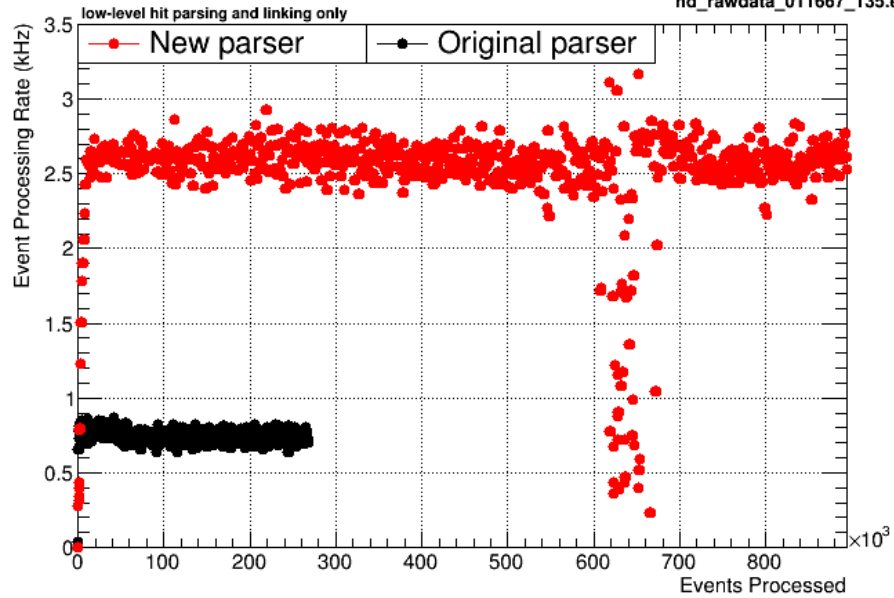
hd_rawdata_011667_135.evio



Event Parsing Rate

May 3, 2016 DL
git revision #8e65dc3

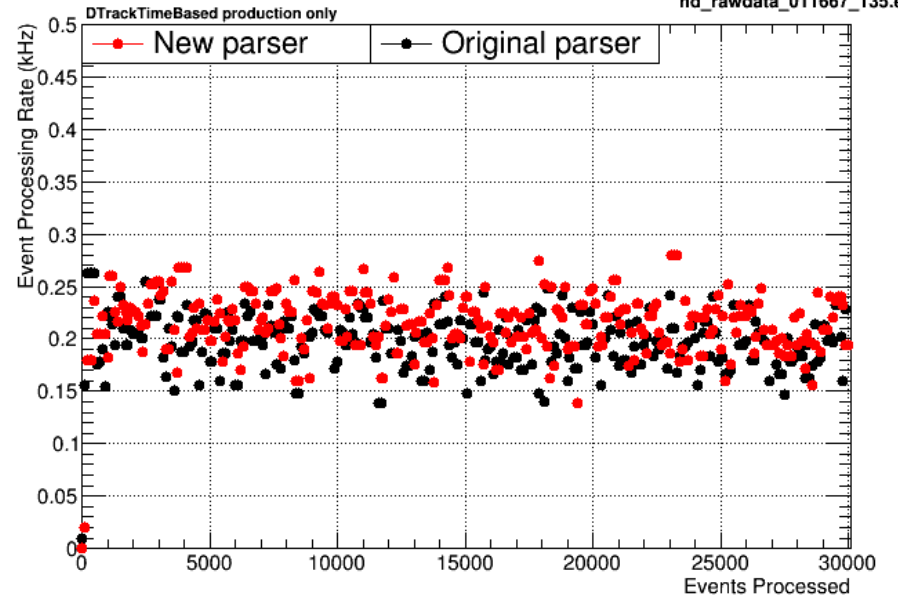
hd_rawdata_011667_135.evio



DTrackTimeBased Rate

May 7, 2016 DL
git revision a347b7

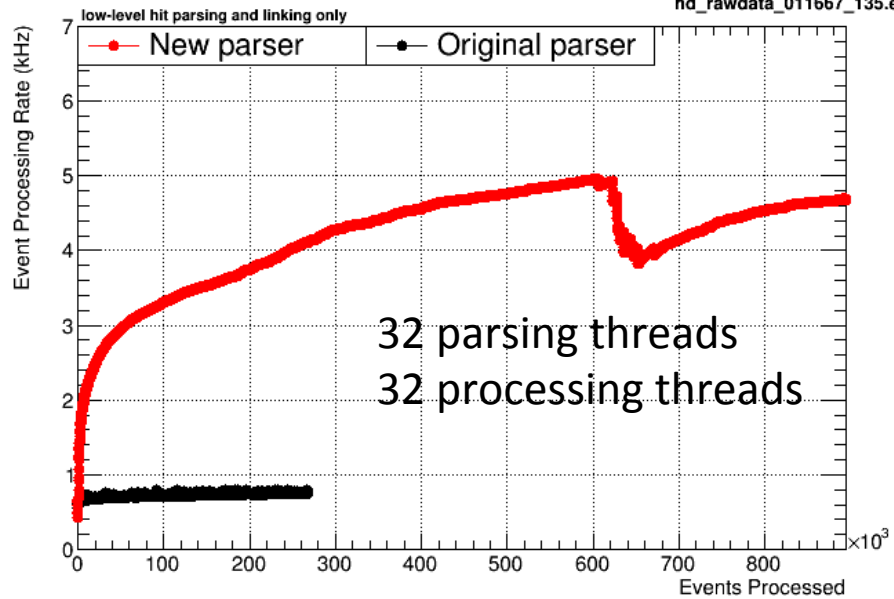
hd_rawdata_011667_135.evio



Event Parsing Memory Usage

May 3, 2016 DL
git revision #8e65dc3

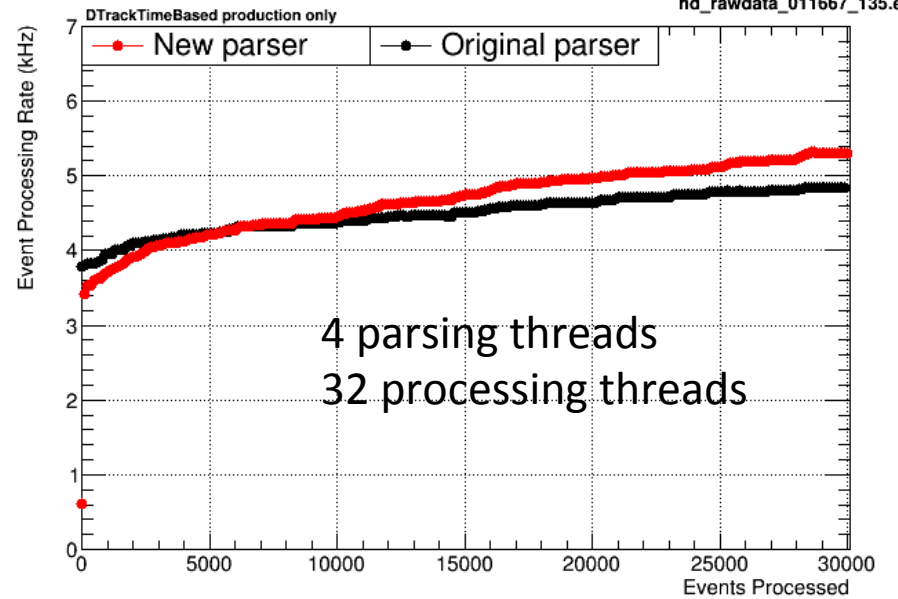
hd_rawdata_011667_135.evio



DTrackTimeBased Memory Usage

May 7, 2016 DL
git revision a347b7

hd_rawdata_011667_135.evio



Storage Group

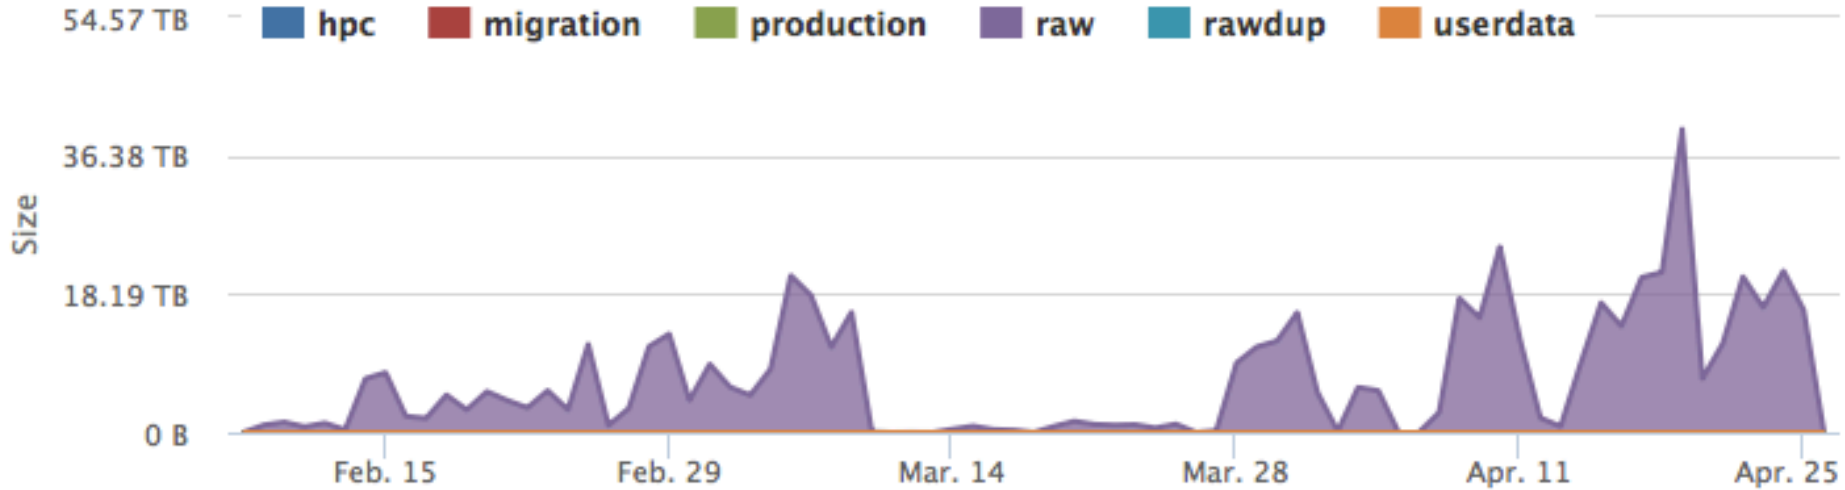
halld

Volume Set

RunPeriod-2016-02-raw

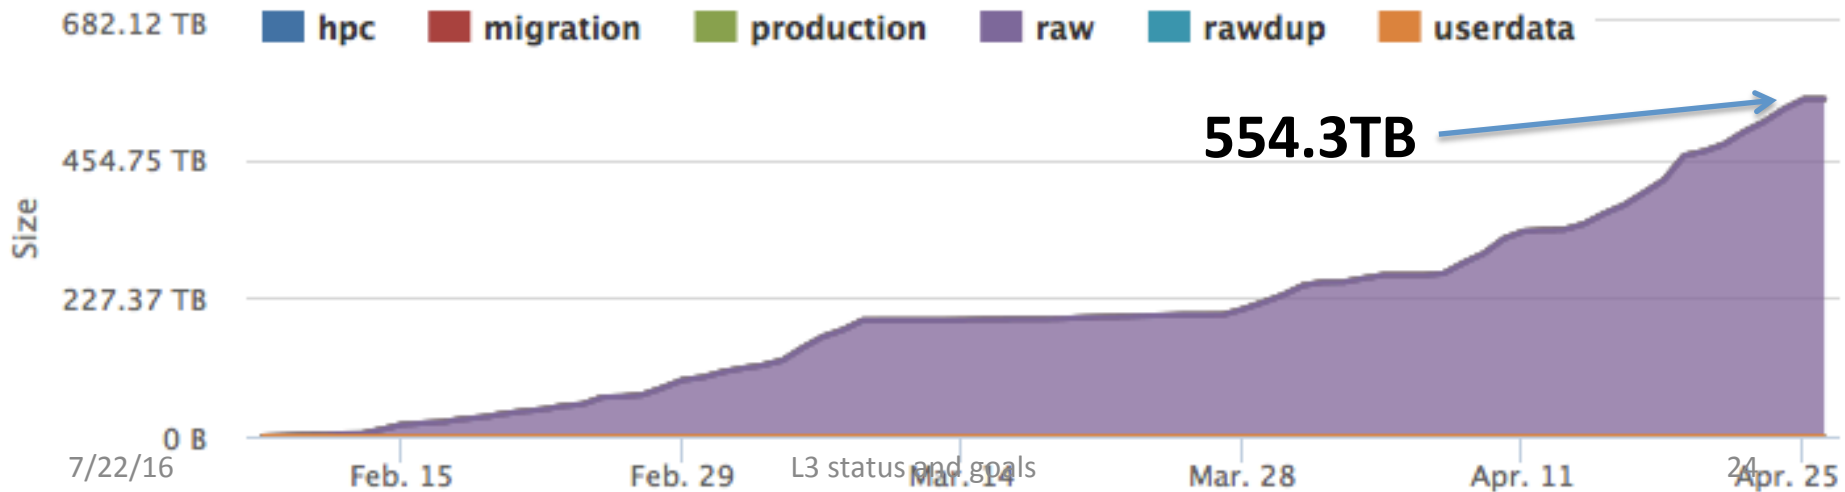
Data Written

Ancestor: halld, Storage Group: halld, Volume Set: RunPeriod-2016-02-raw



Data Written (Cummulative)

Ancestor: halld, Storage Group: halld, Volume Set: RunPeriod-2016-02-raw

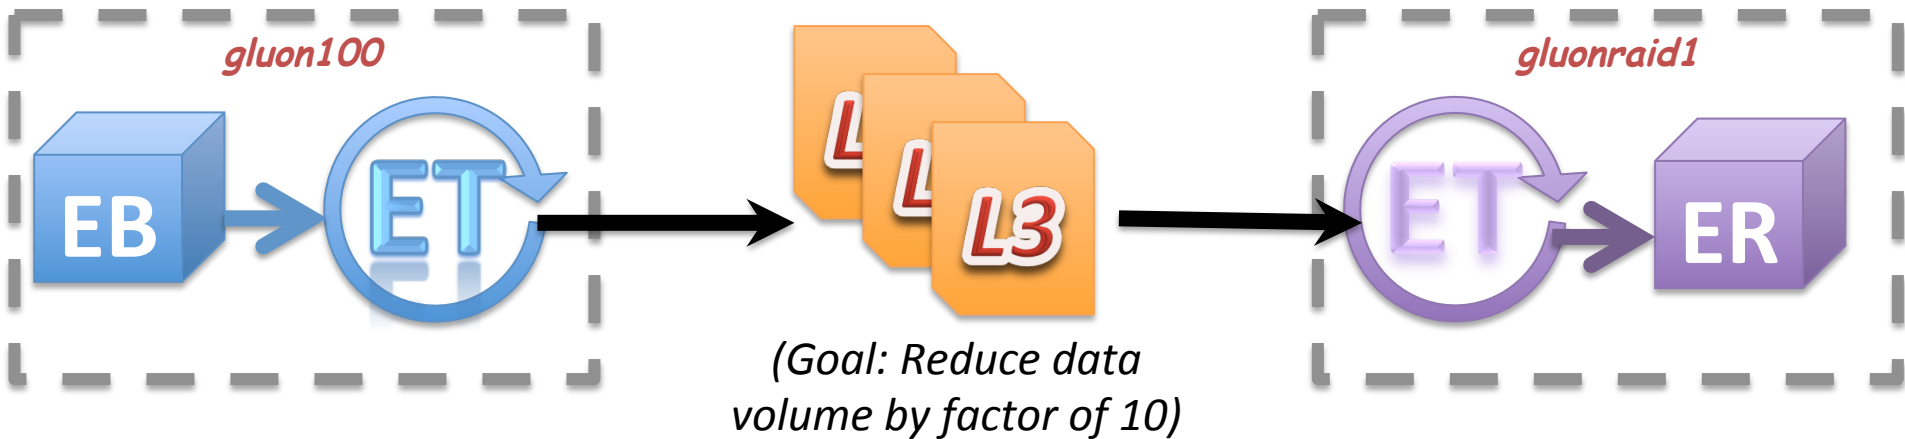


Online Storage Capacity

- 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



L3 Testing



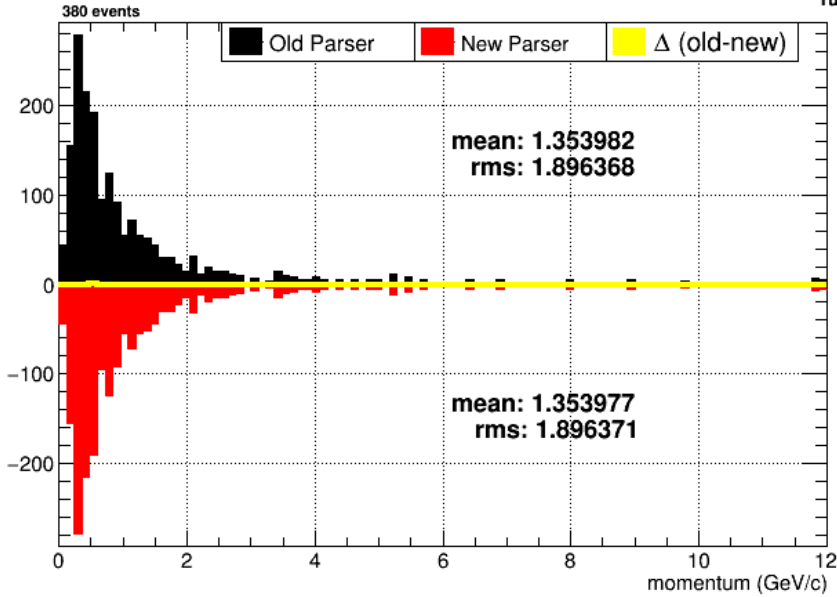
Spring 2016 L3 Testing

- Primarily infrastructure testing done
- Limited by 12 node farm to ~200MB/s parsing rate
- (~10kHz input rate)

Reconstructed parameters for old and new parsers

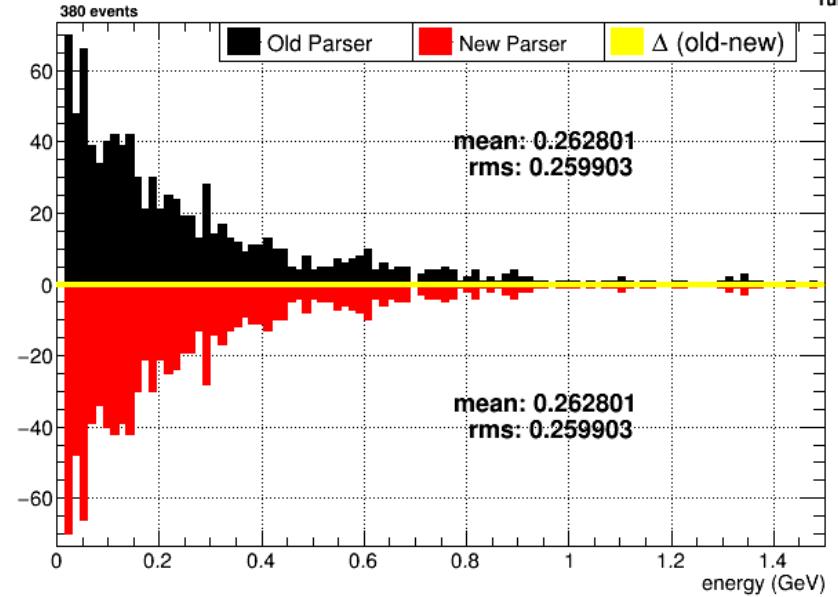
Momentum of Time-based tracks

May 8, 2016 DL
git revision 8347b7
run 11667



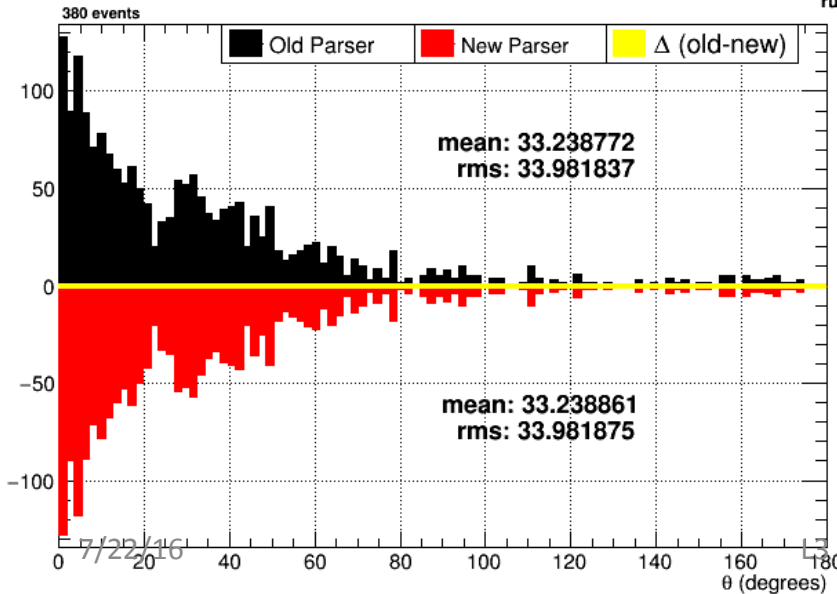
Reconstructed BCAL Showers

May 8, 2016 DL
git revision 8347b7
run 11667



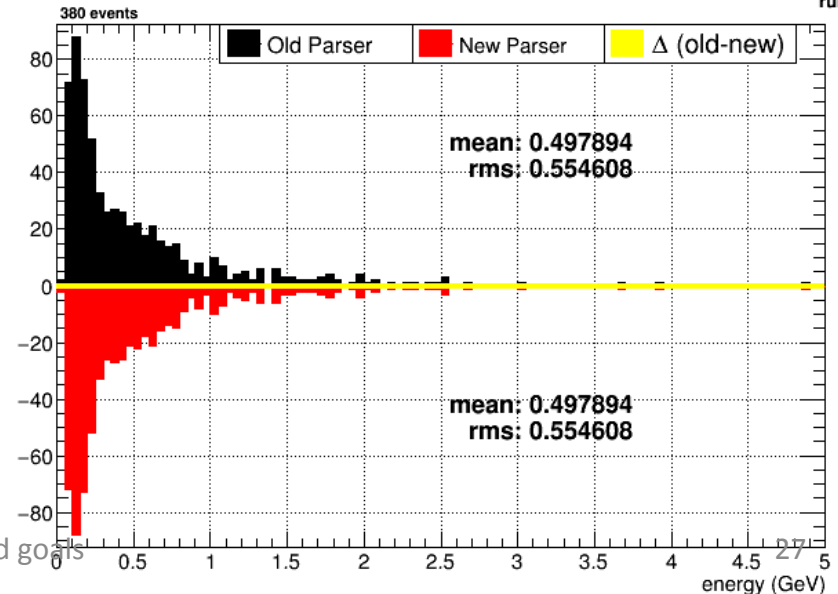
θ angle of Time-based tracks

May 8, 2016 DL
git revision 8347b7
run 11667



Reconstructed FCAL Showers

May 8, 2016 DL
git revision 8347b7
run 11667



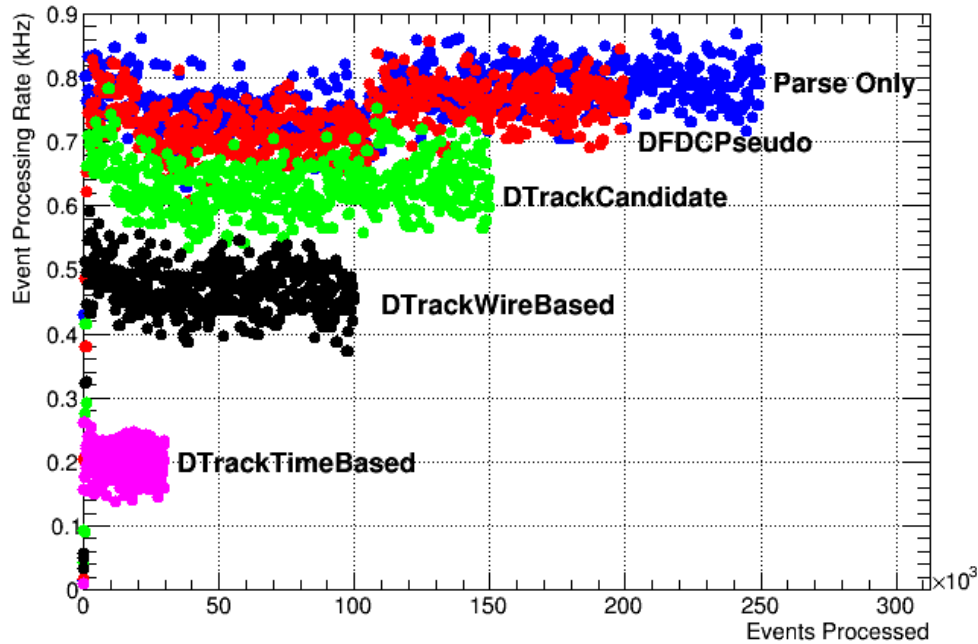
7/22/16

status and goals

27

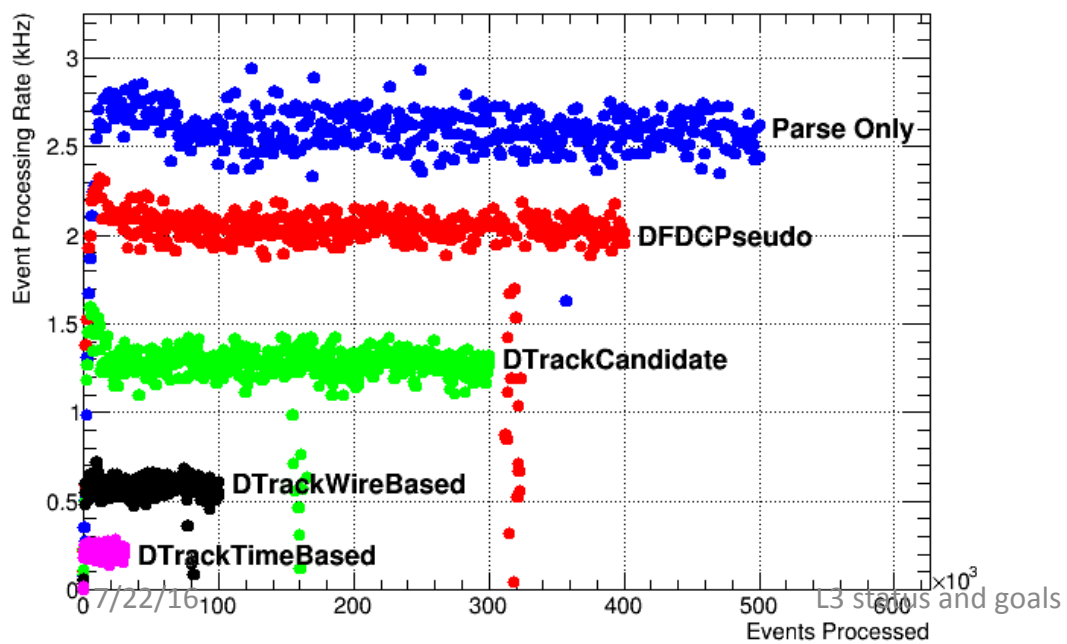
Old Parser Processing Rates

May 7, 2016 DL
git revision a347b7
hd_rawdata_011667_135.evio

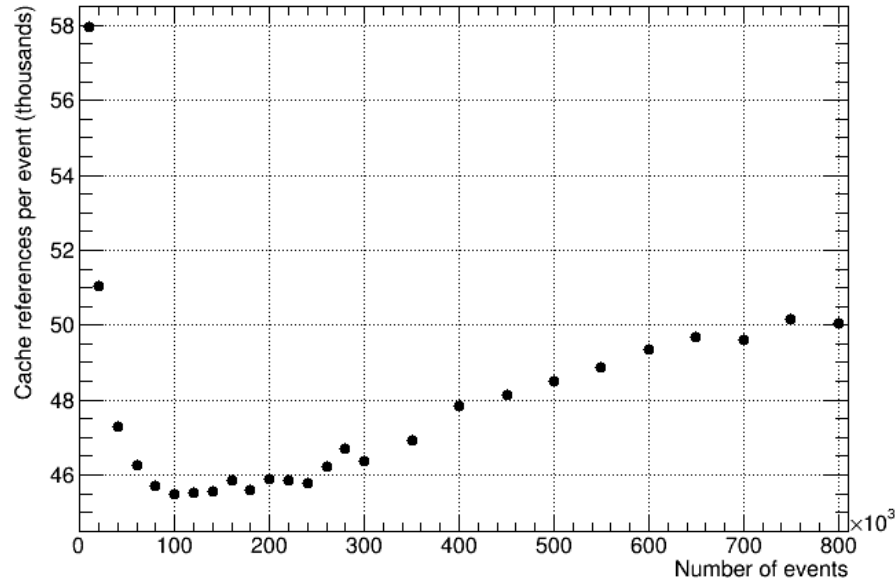


New Parser Processing Rates

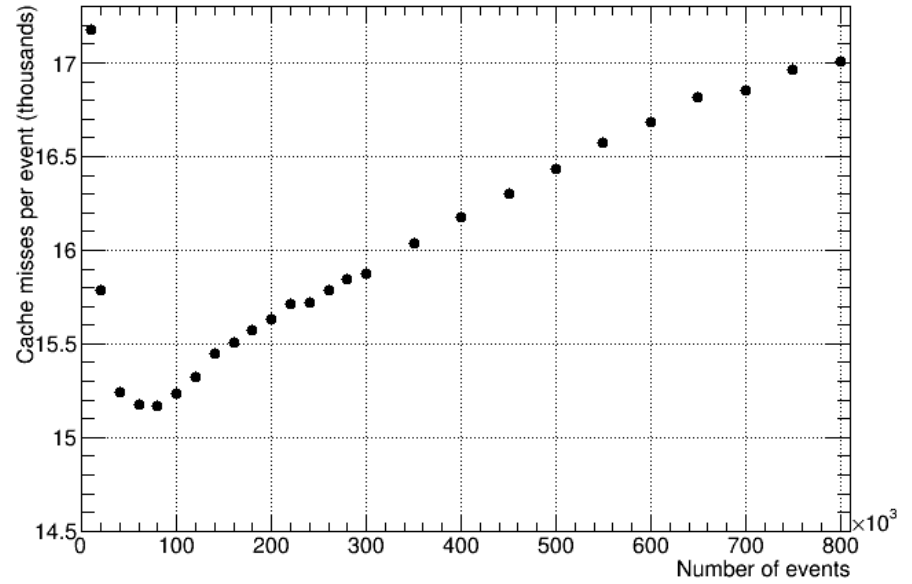
May 7, 2016 DL
git revision a347b7
hd_rawdata_011667_135.evio



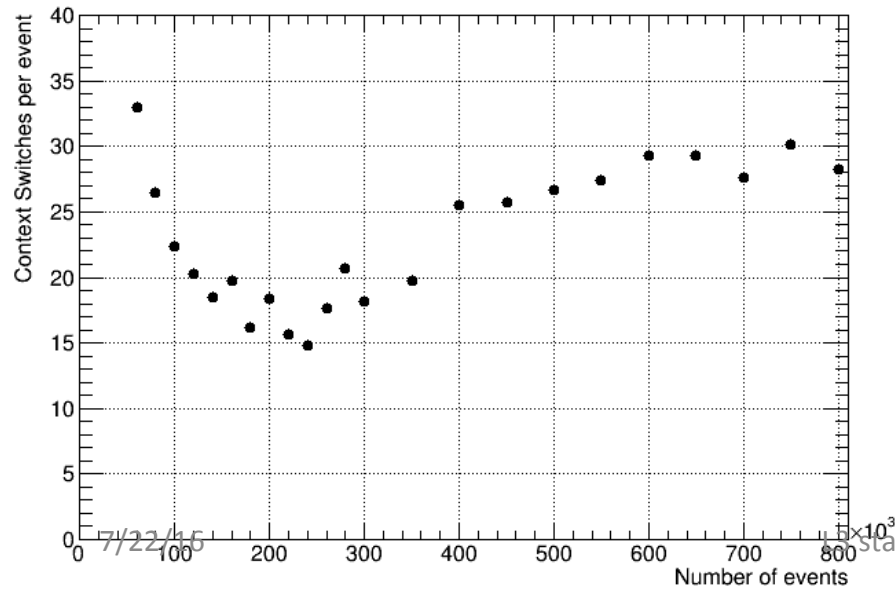
Cache references/event



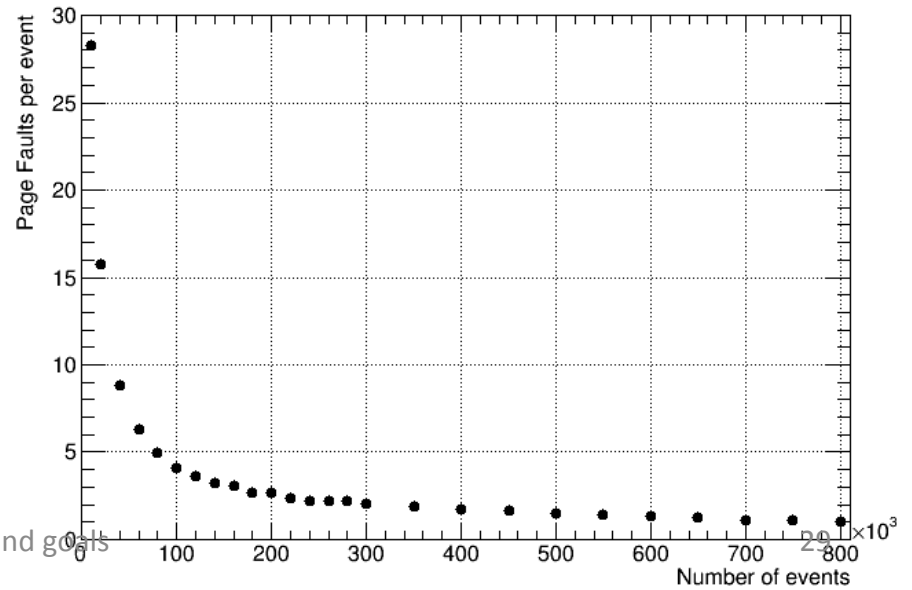
Cache misses/event



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

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

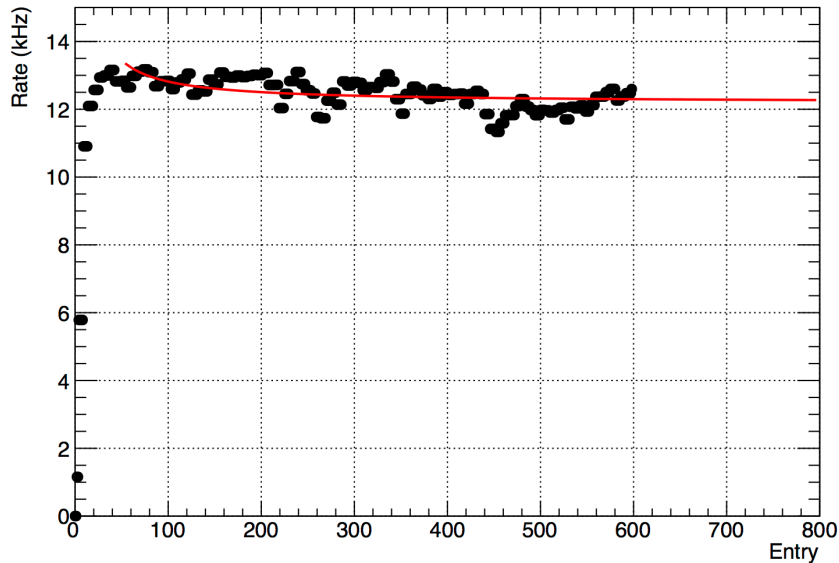
Fitting event rate vs. time

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

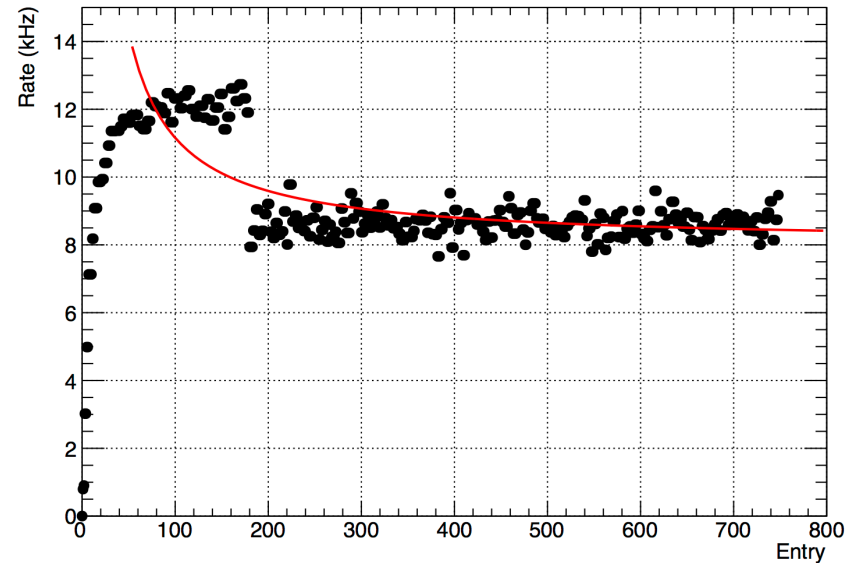
R_o = asymptotic rate

Q = relaxation term

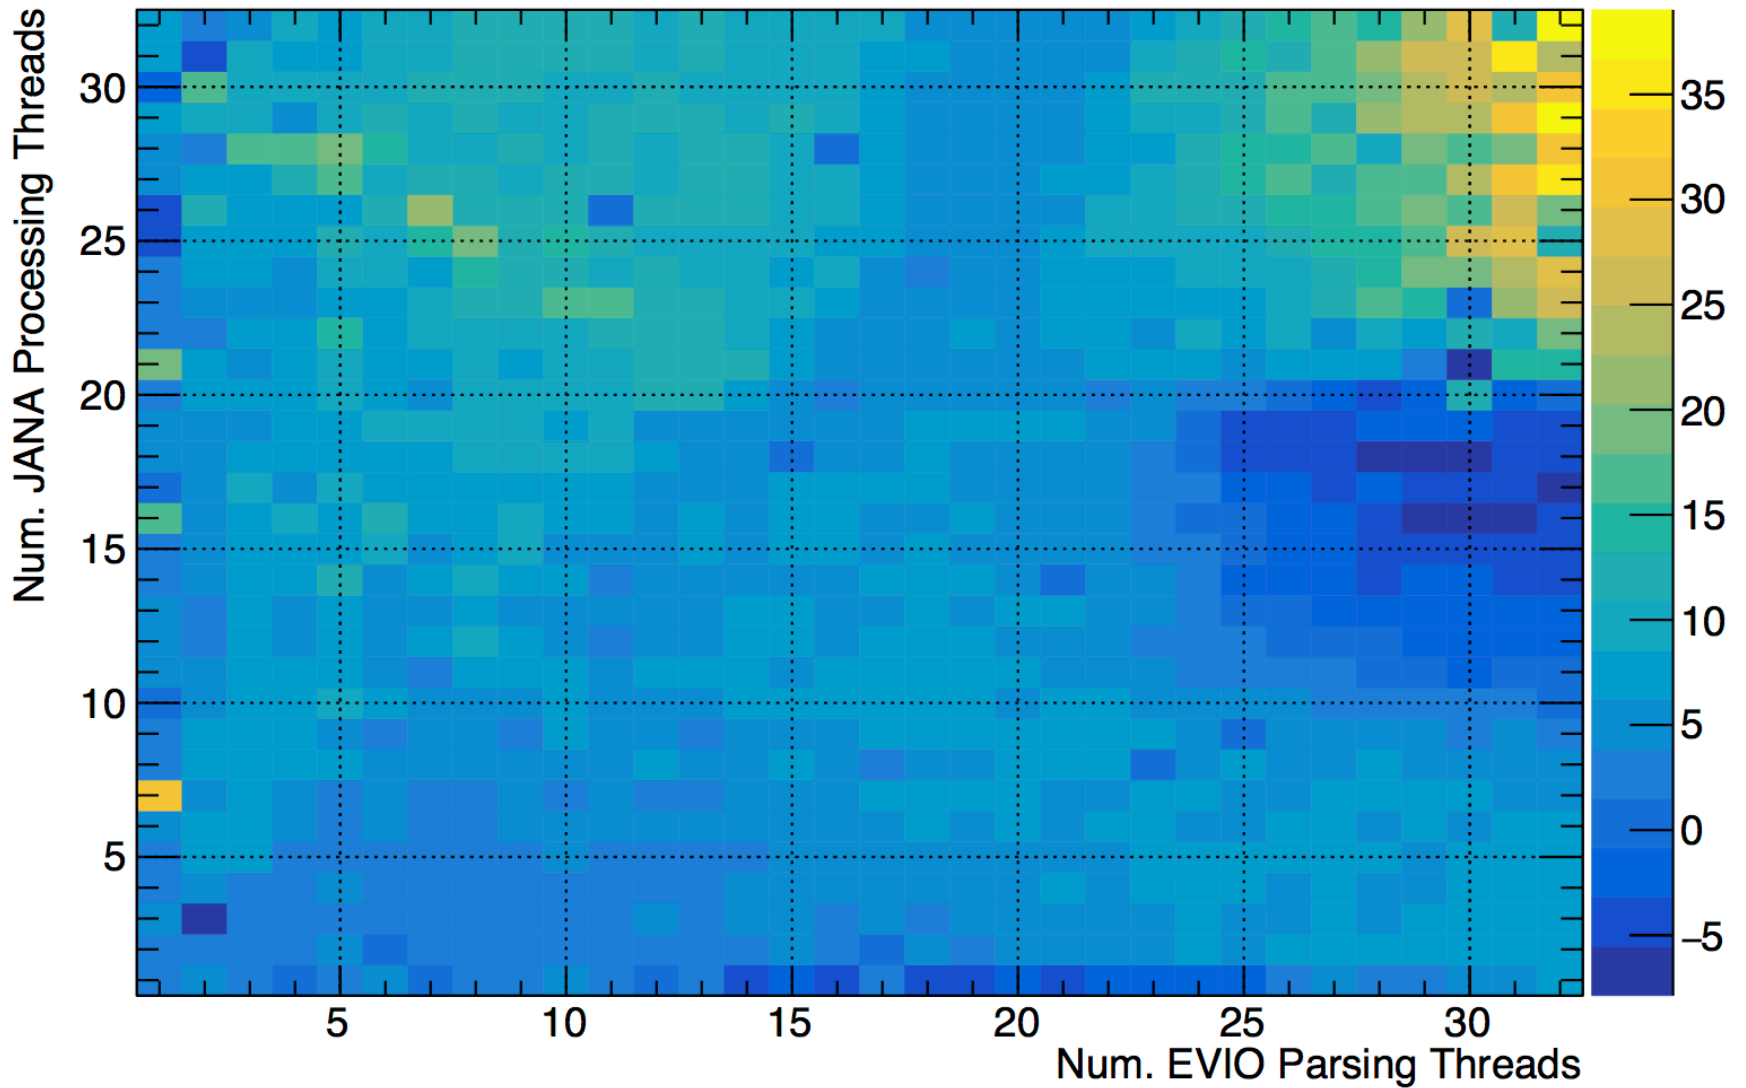
14 Workers 18 Processors



32 Workers 32 Processors

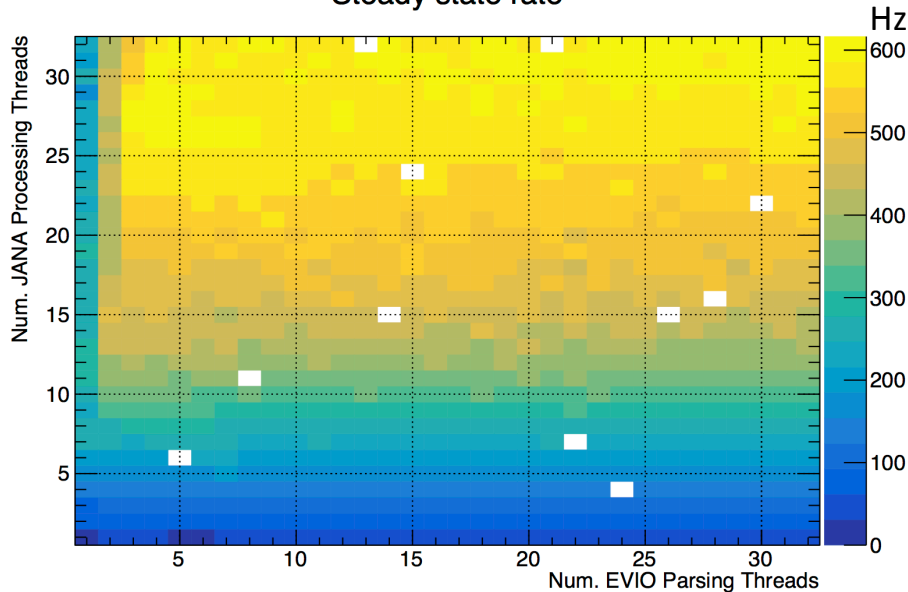


Relaxation term

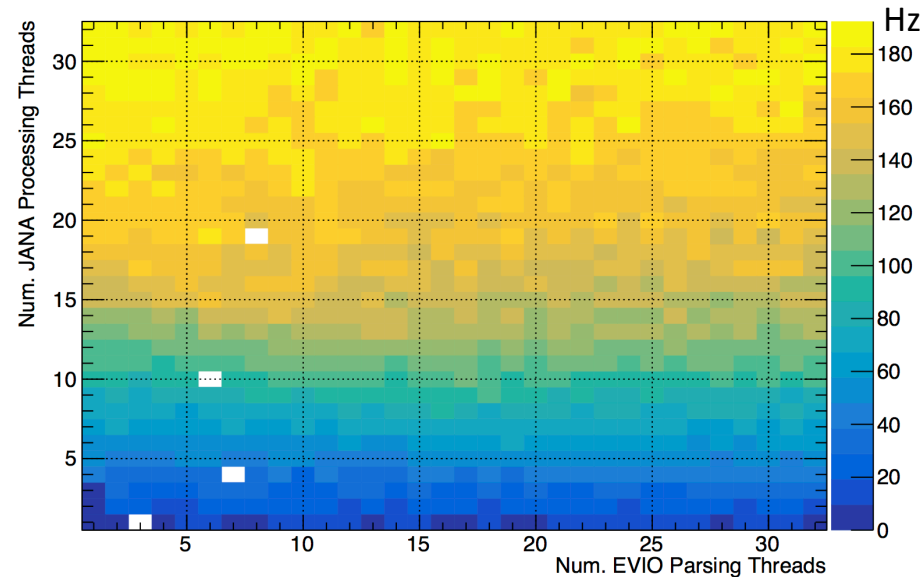


Event rates with tracking

Wire-based Tracking
Steady state rate



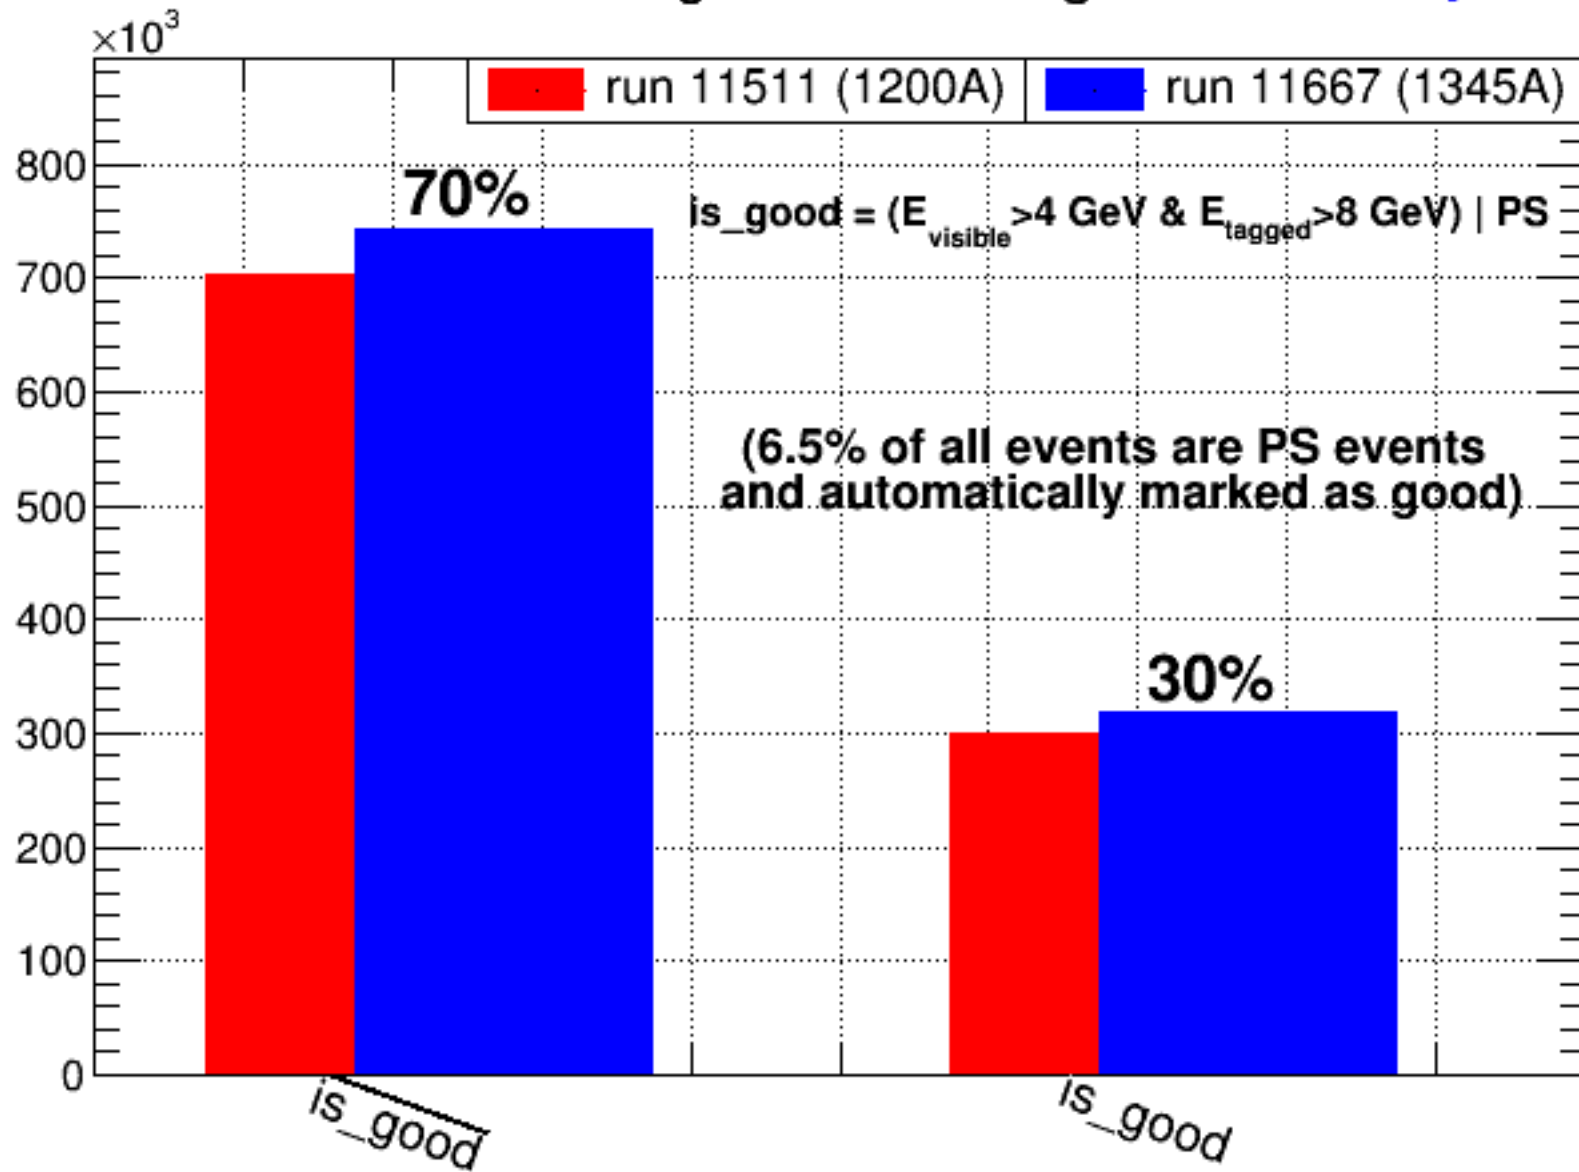
Time-based Tracking
Steady state rate



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

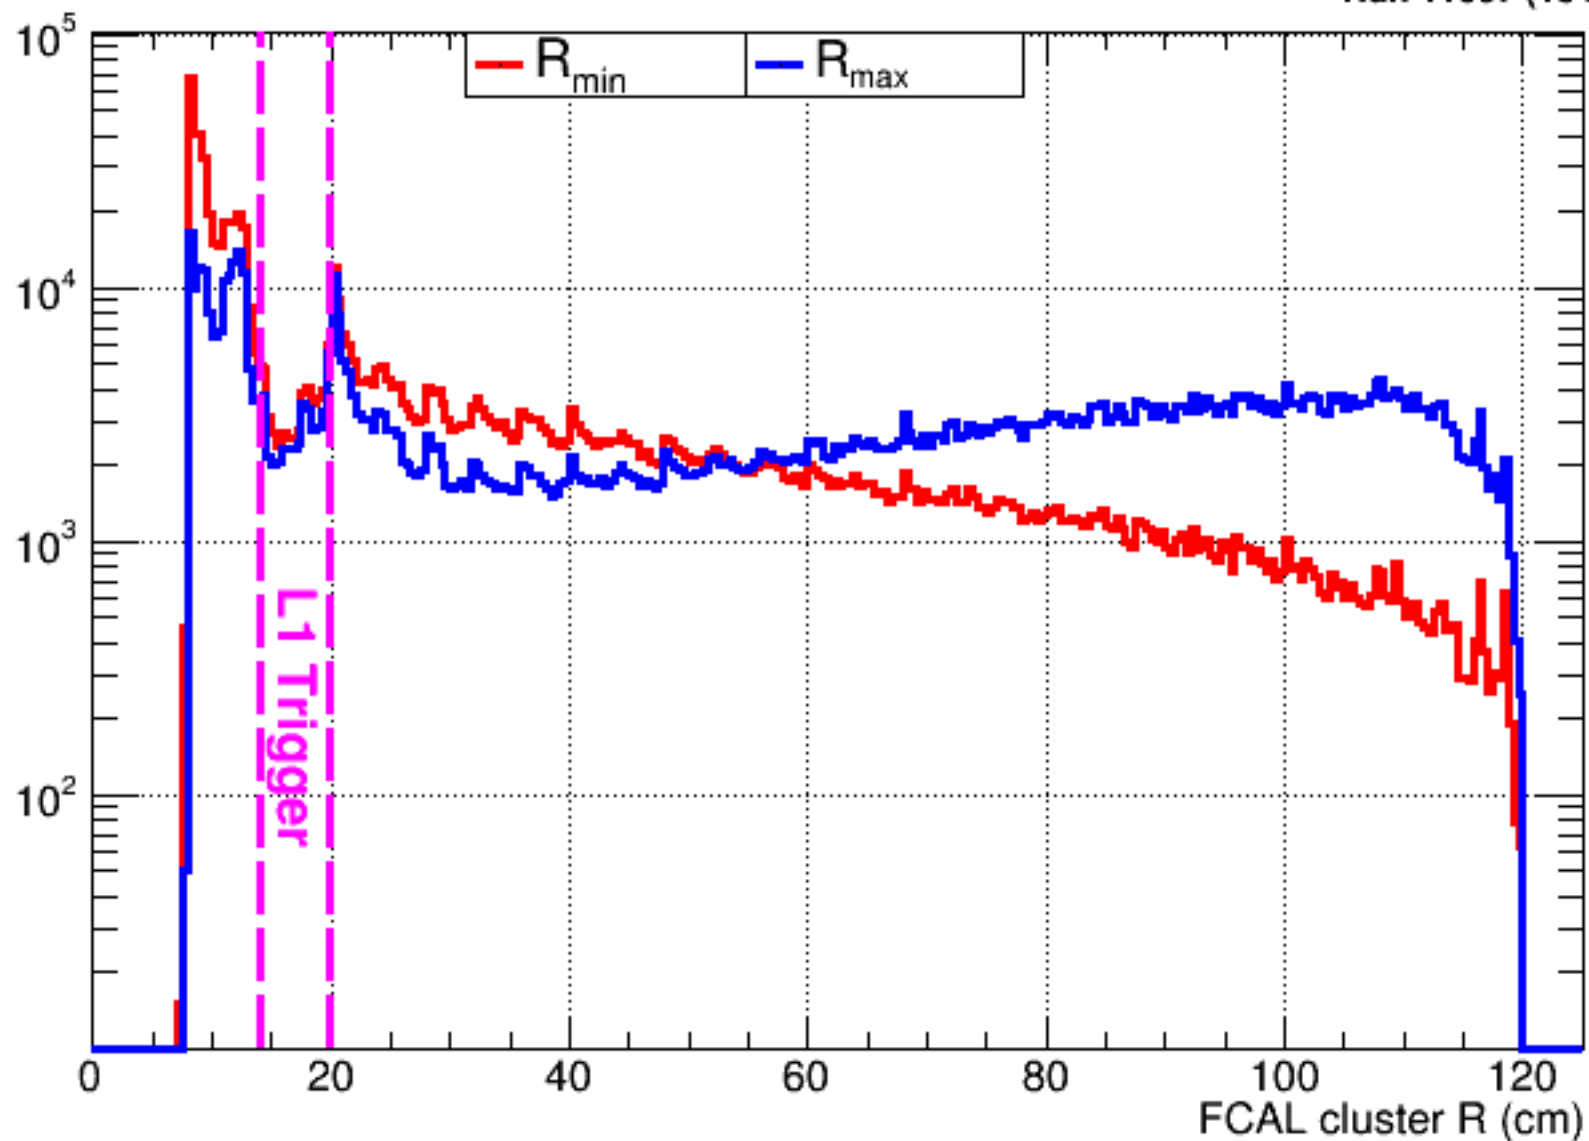
L3 good event flag

July 7, 2016 DL
git revision #7118b85

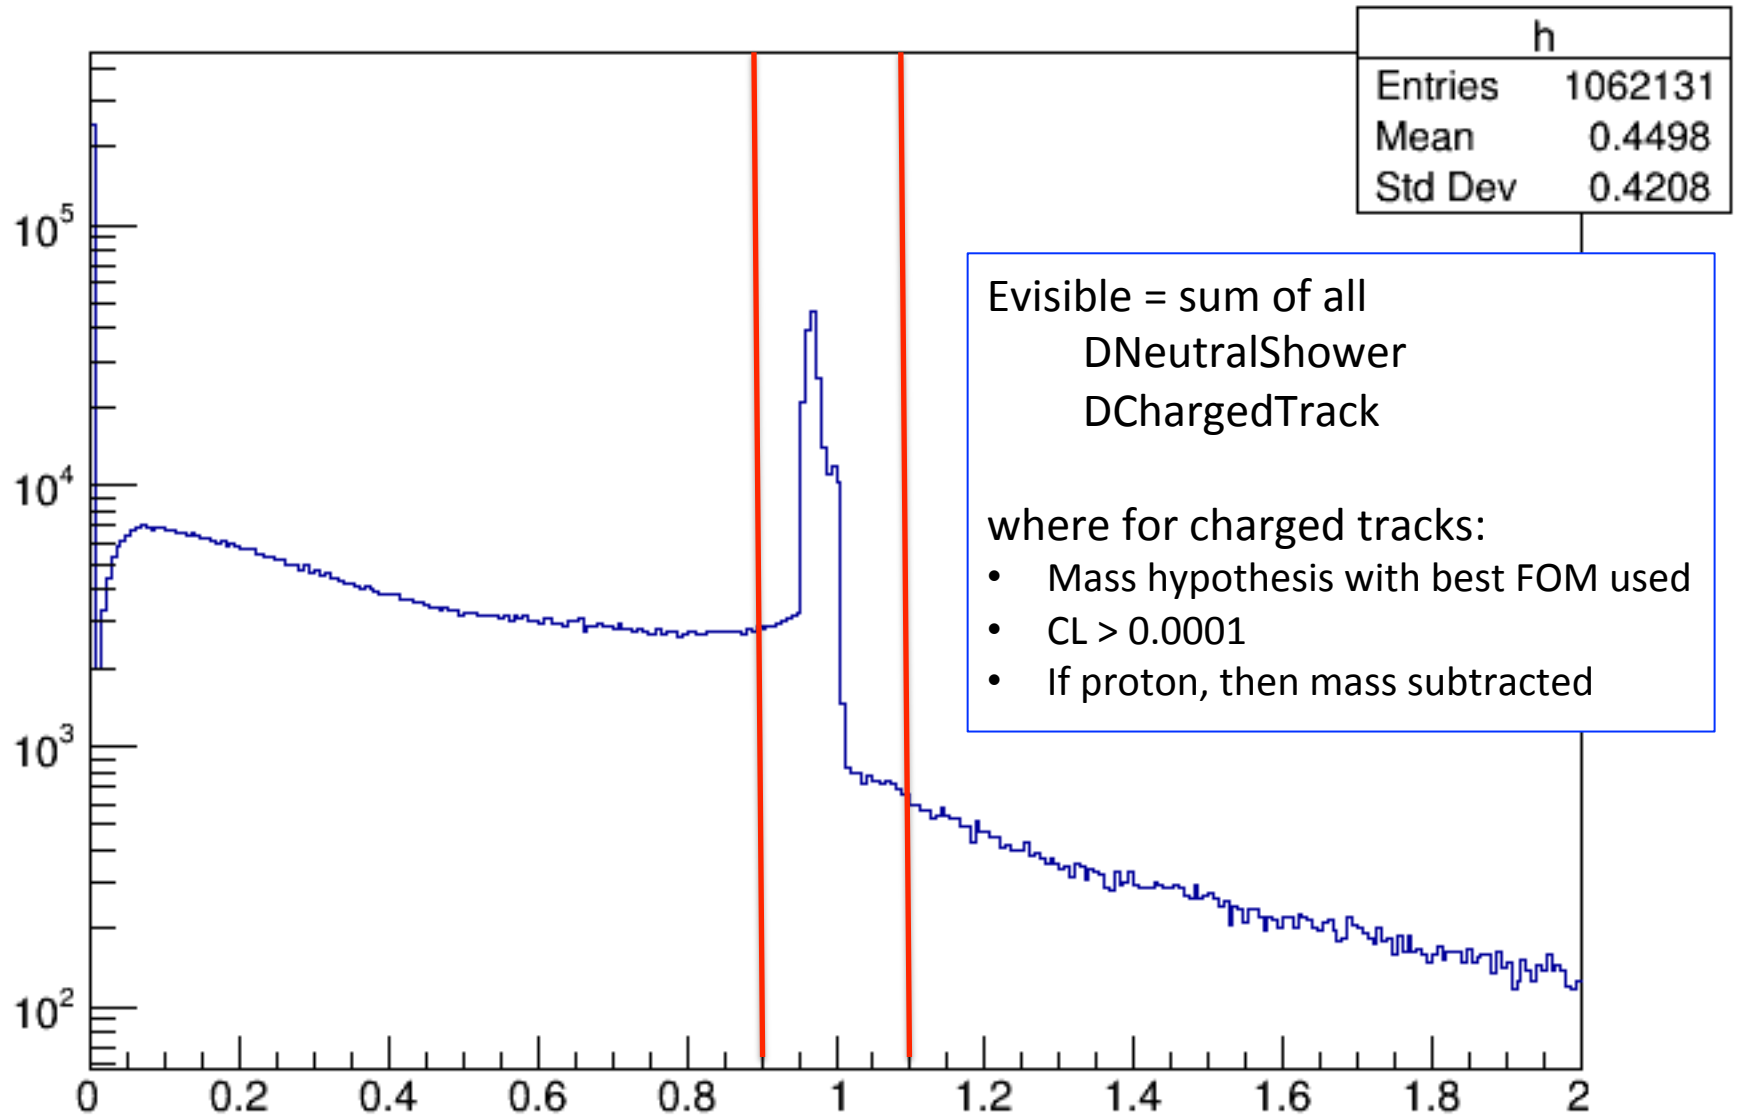


FCAL cluster R

July 7, 2016 DL
git revision #7118b85
Run 11667 (1345A)

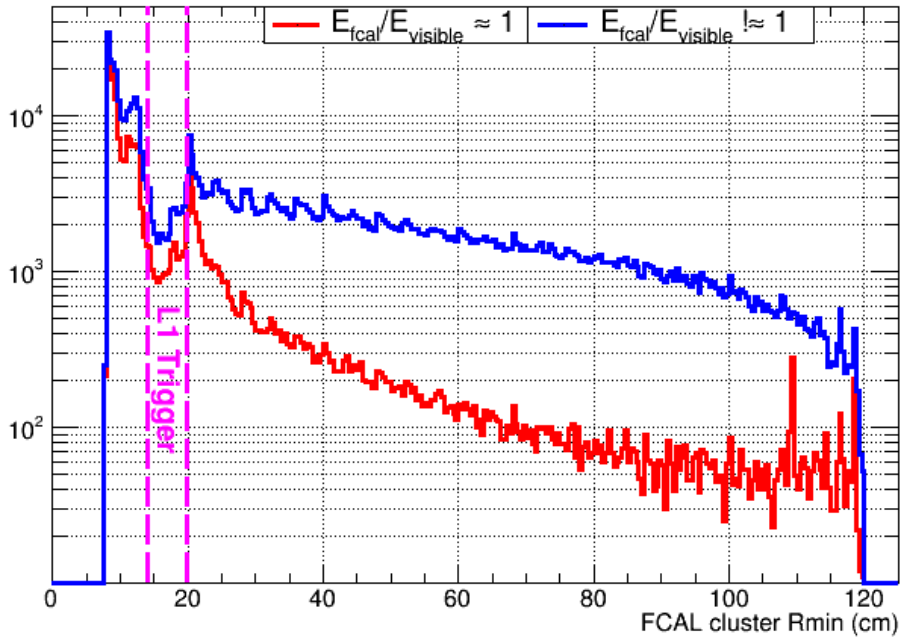


Efcals_clusters/Evisible



FCAL cluster Rmin

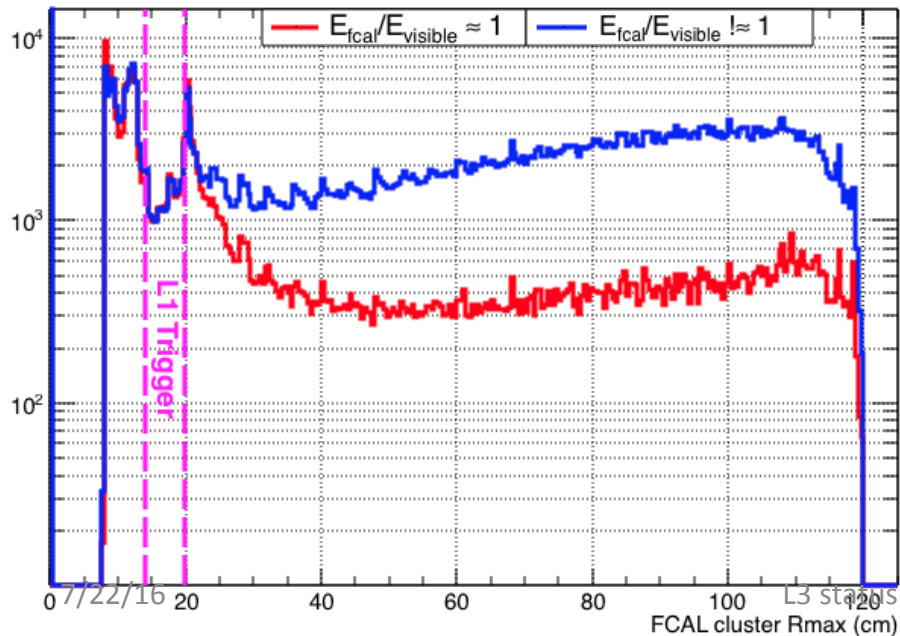
July 7, 2016 DL
git revision #7118b85
Run 11667 (1345A)



FCAL Rmin and Rmax
for when most of Evisible is
inside (outside) of FCAL

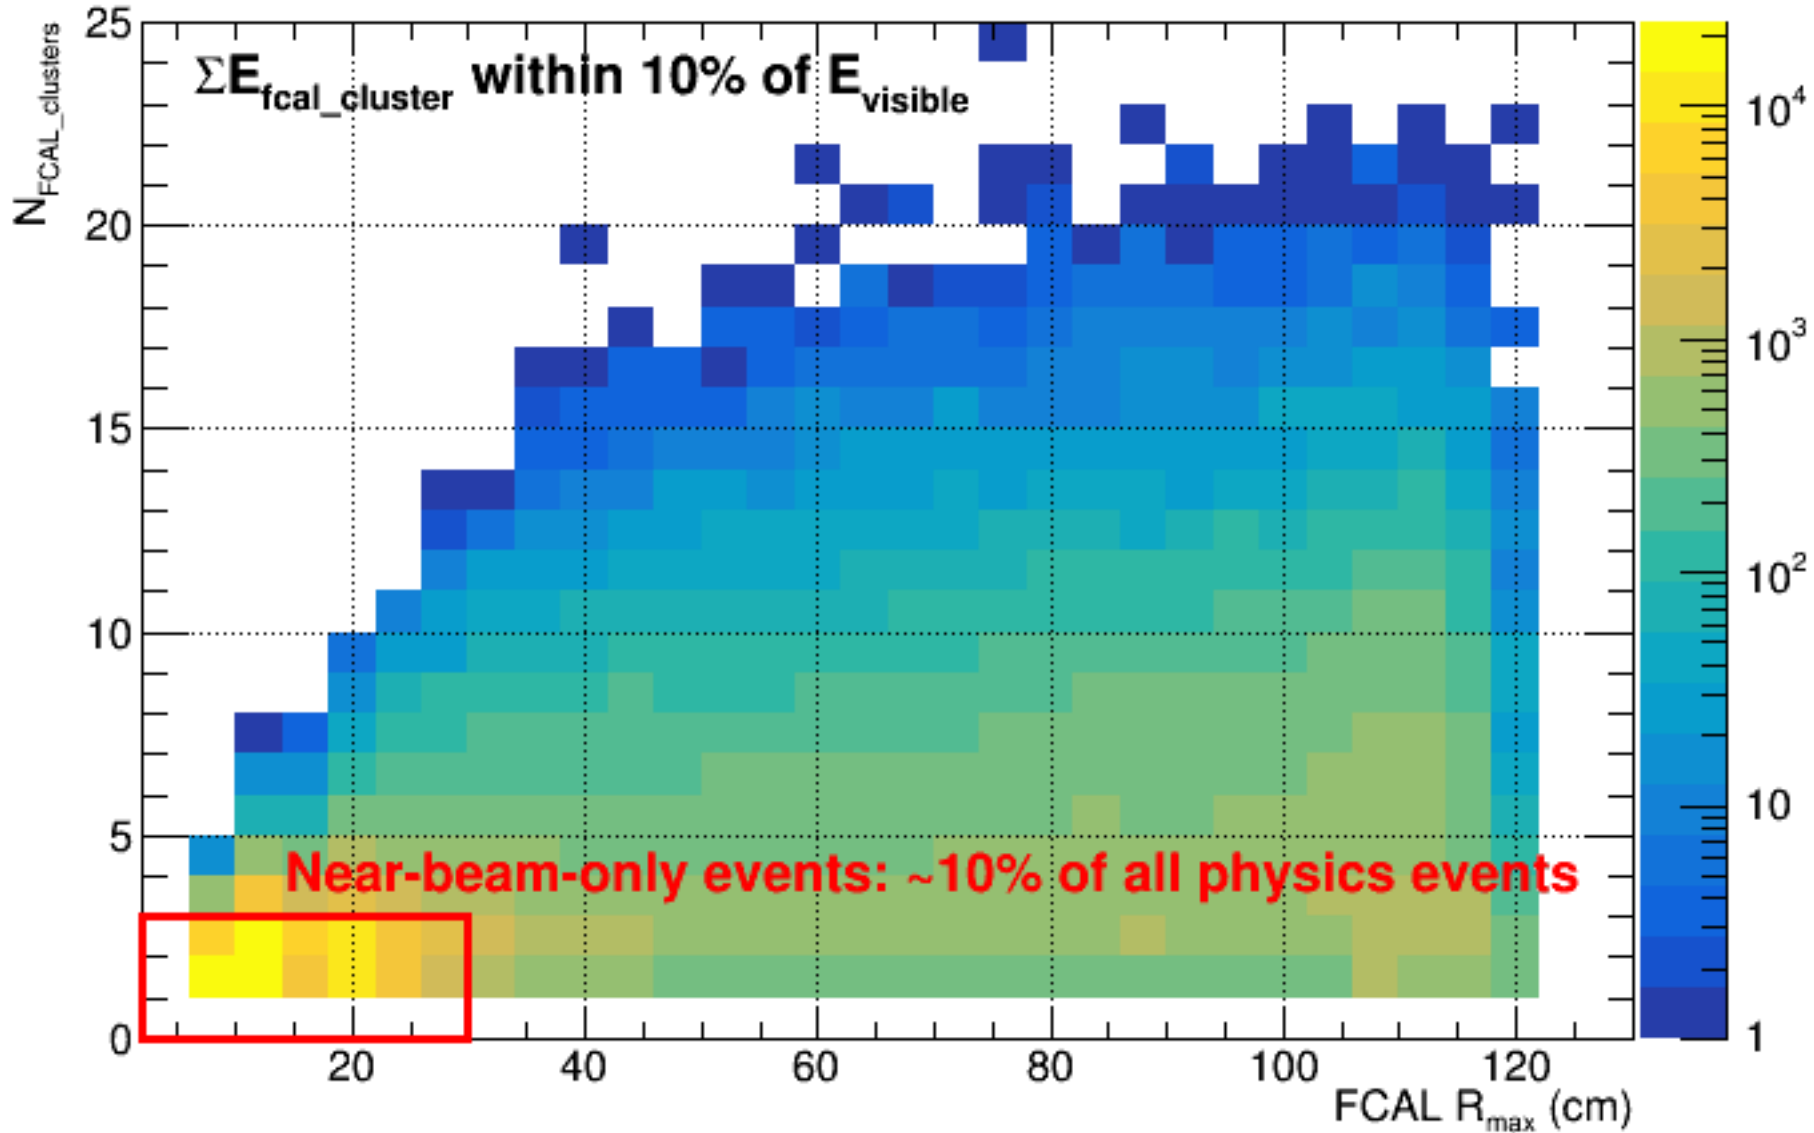
FCAL cluster Rmax

July 8, 2016 DL
git revision abc1904
Run 11667 (1345A)

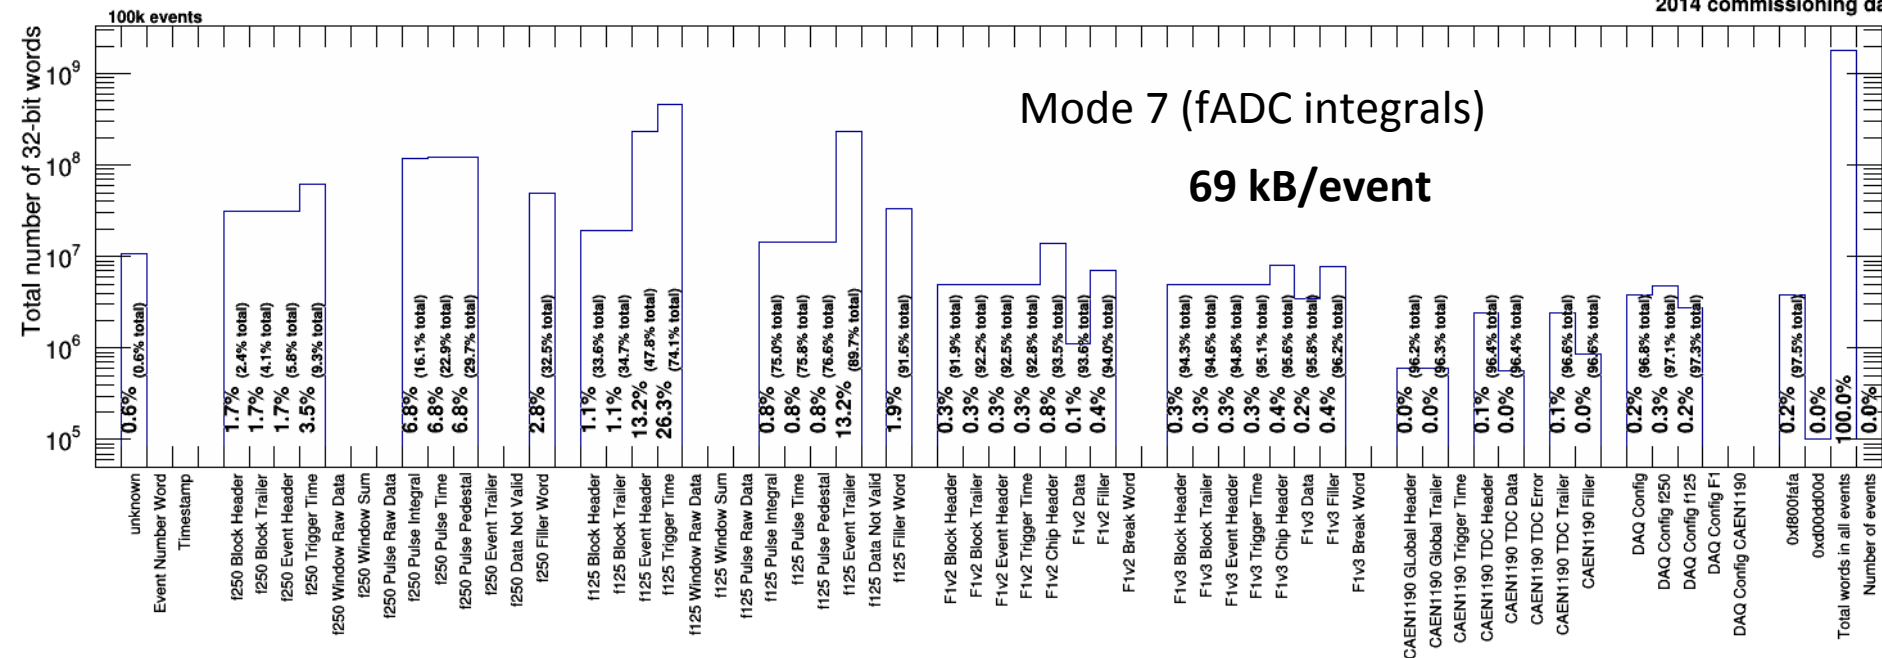


L3 status and goals

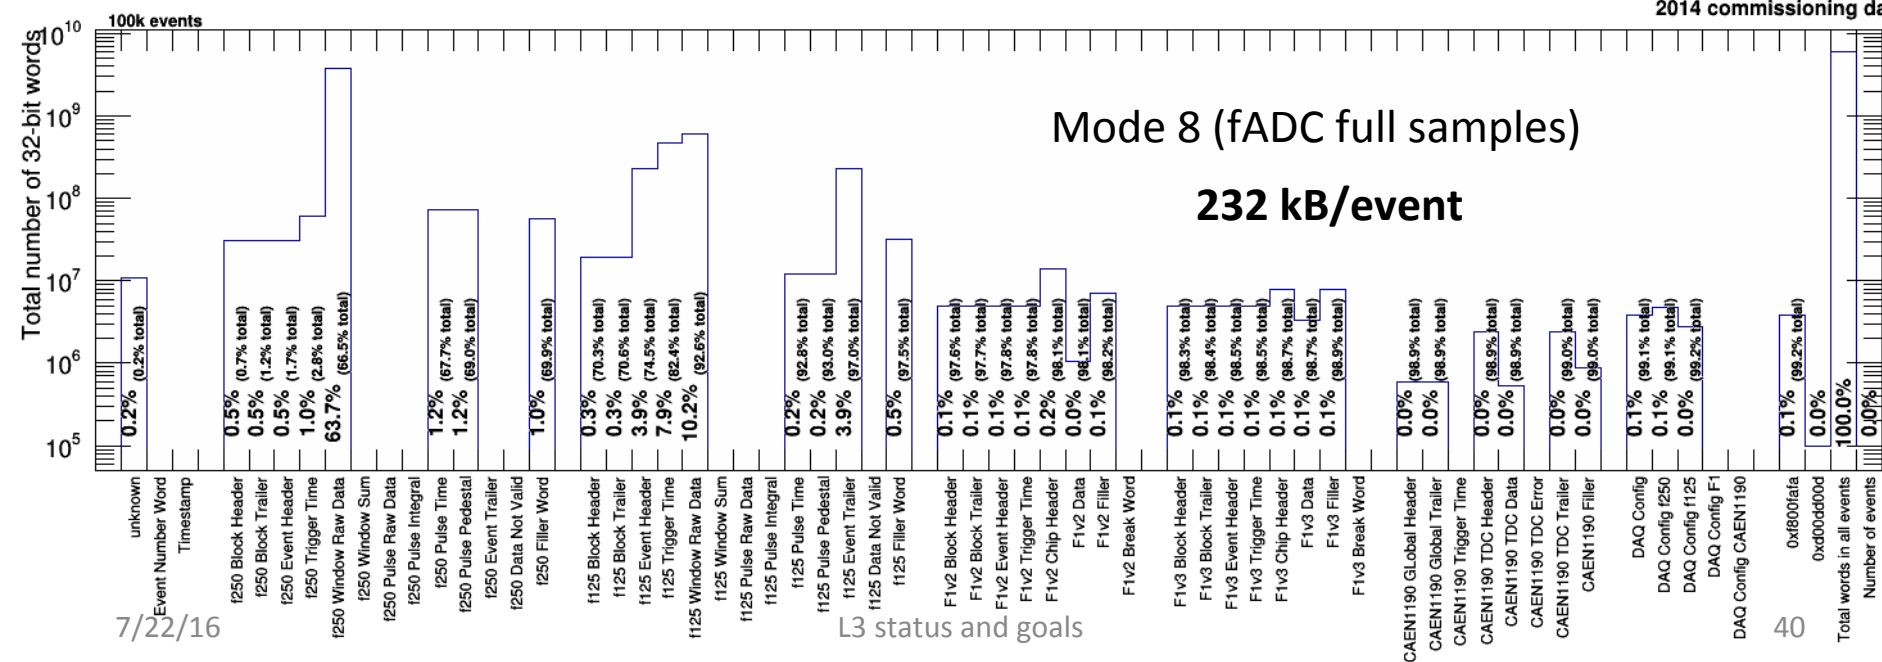
FCAL N_{clusters} vs. R_{max}



Number of words in 100k events for run 2179



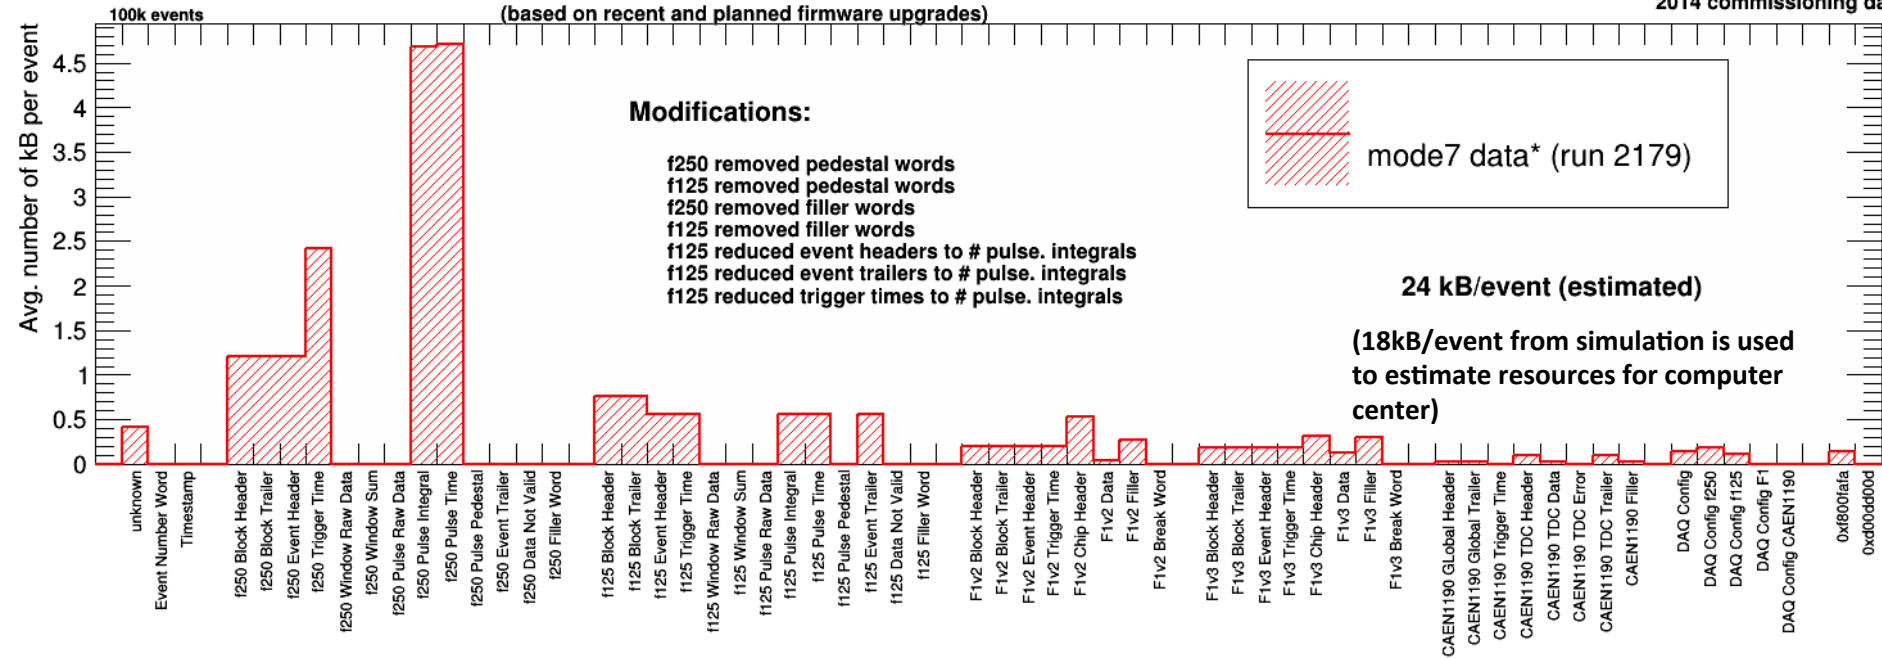
Number of words in 100k events for run 2180



Adjusting profile of 2014 commissioning data based on recent or planned firmware upgrades is used to estimate event size for production data in the future.

February 3, 2015 DL
sim-recon: svn 17000, JANA svn 2115
2014 commissioning data

Estimate of future production running EVIO words/event (based on recent and planned firmware upgrades)



(Additional compression is expected when disentangled data is rebuilt after L3 into an as yet undetermined format.)


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)

- Each 32bit word in the EVIO file tallied to identify what file space is being used for
- Comparison between mode 7 and mode 8 data made

Example: some of the fADC250 word types

Mode 7 (fADC Integrals)

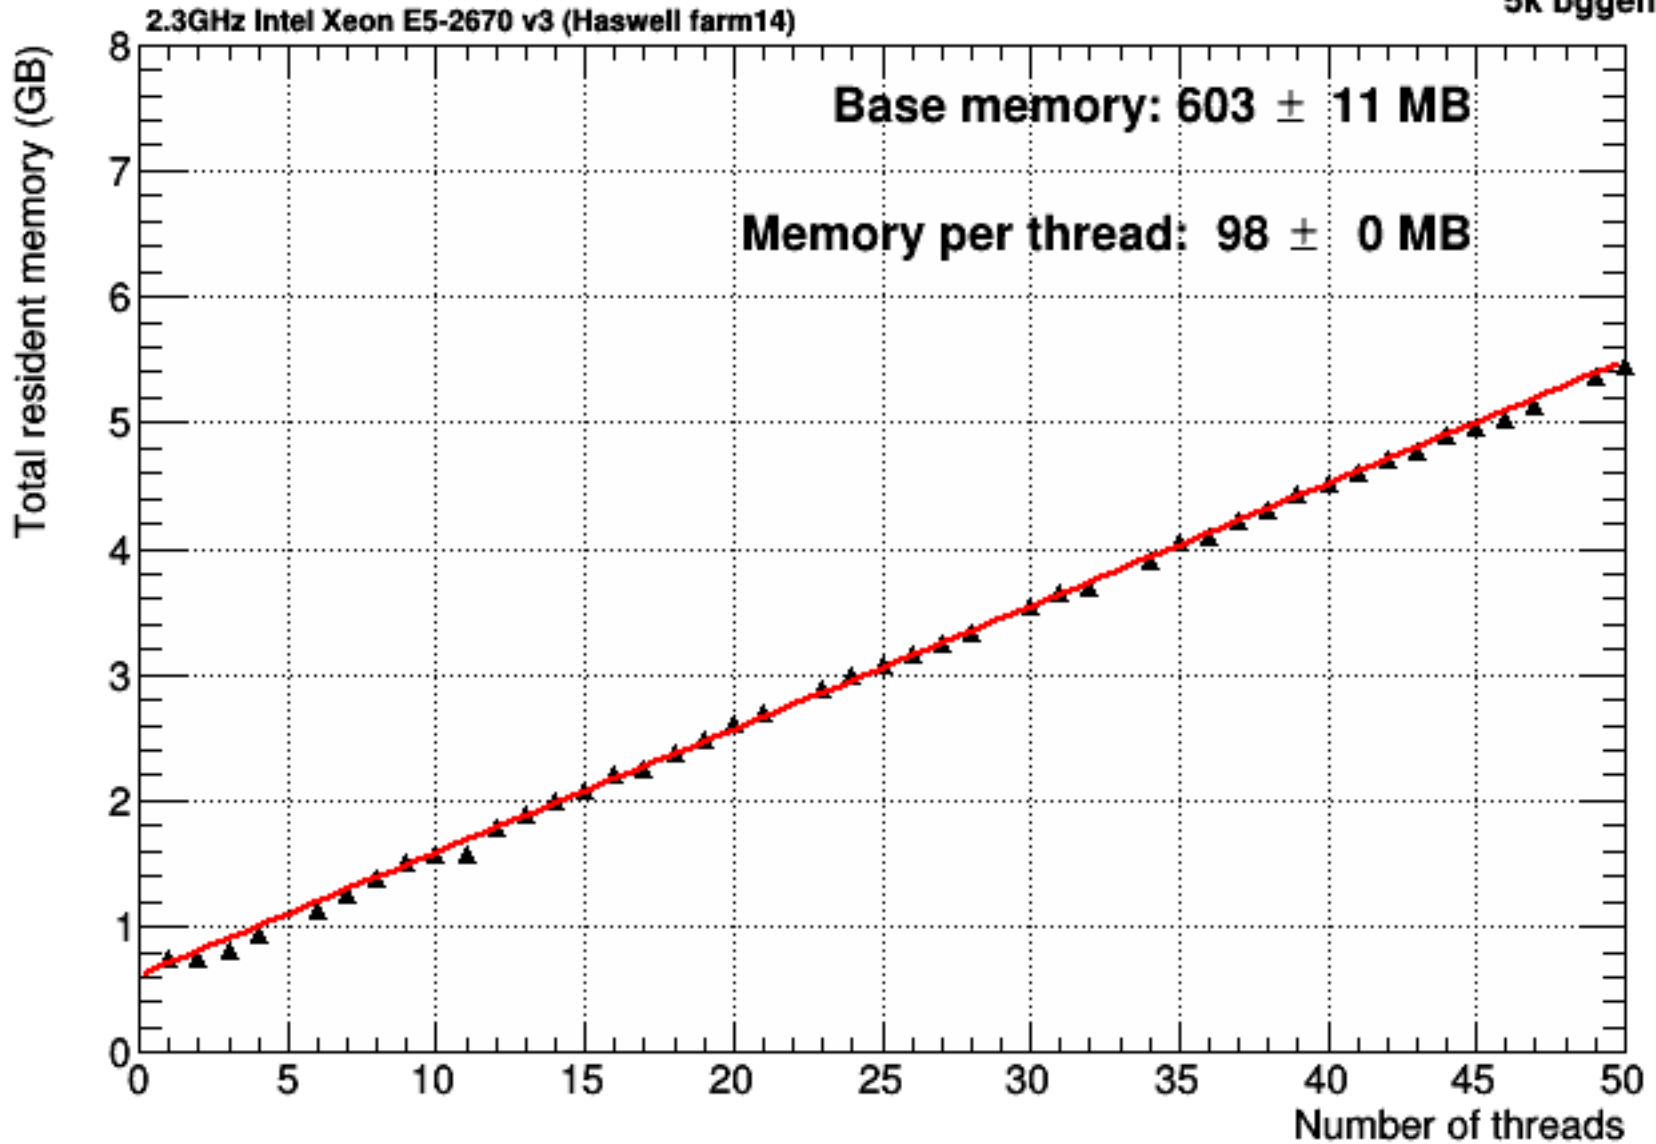



Mode 8 (fADC full samples)



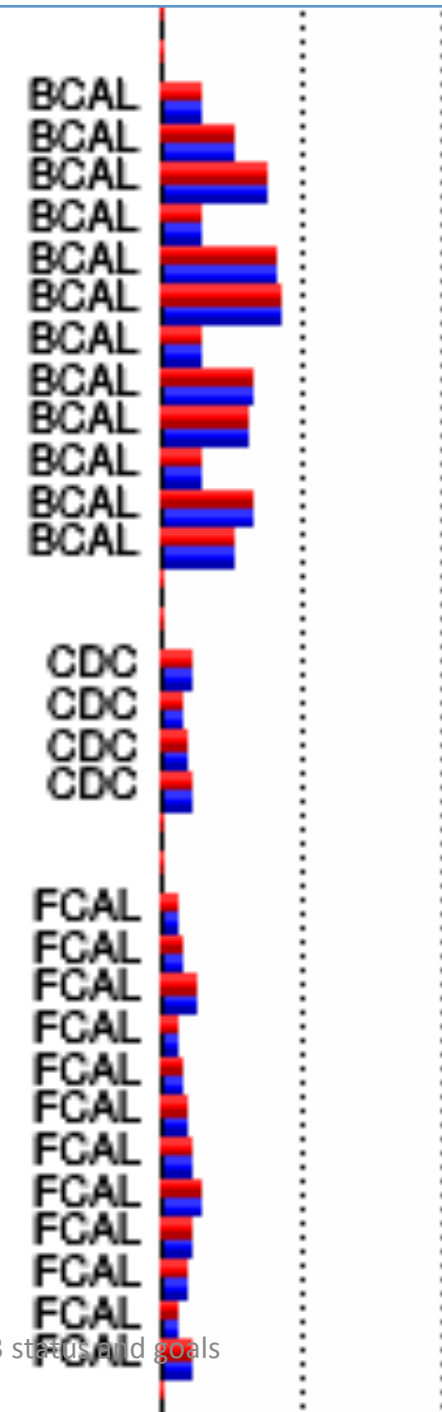
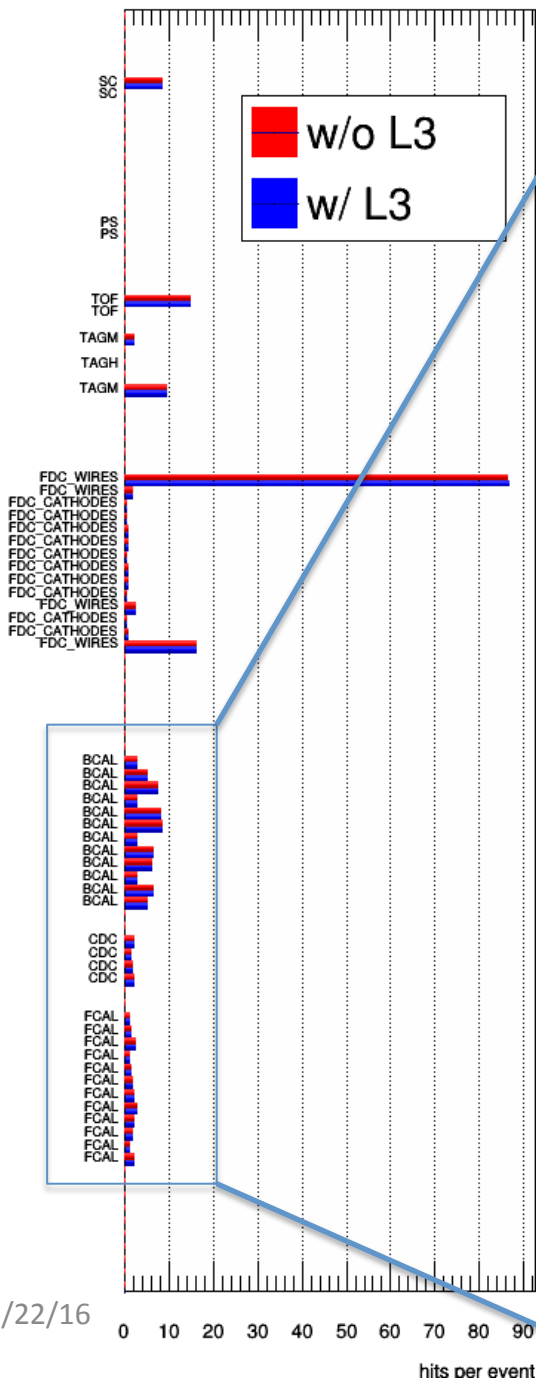
Memory Usage vs. Num. Threads

January 23, 2015 DL
sim-recon: svn 17000, JANA svn 2115
5k bggen events



name	node	level 	Nthr	Nevents	rate (Hz)	idle
gluon101.jlab.org_11057	gluon101.jlab.org	--	32	212874	751.1	75.2%
gluon49.jlab.org_12748	gluon49.jlab.org	--	32	206140	646.7	22.8%
gluon111.jlab.org_5726	gluon111.jlab.org	--	32	206700	643.1	58.1%
gluon110.jlab.org_30100	gluon110.jlab.org	--	32	208726	700.7	4.1%
gluon109.jlab.org_4409	gluon109.jlab.org	--	32	213494	369.9	75.2%
gluon108.jlab.org_10935	gluon108.jlab.org	--	32	247875	390.6	77.0%
gluon107.jlab.org_23963	gluon107.jlab.org	--	32	211054	406.0	72.1%
gluon106.jlab.org_20172	gluon106.jlab.org	--	32	219621	895.5	49.8%
gluon105.jlab.org_10192	gluon105.jlab.org	--	32	212474	507.2	71.1%
gluon104.jlab.org_23134	gluon104.jlab.org	--	32	205814	549.7	66.7%
gluon102.jlab.org_22561	gluon102.jlab.org	--	32	206803	451.5	71.7%
gluon48.jlab.org_18860	gluon48.jlab.org	--	32	214251	519.2	49.7%

Hits/event vs. rocid



Comparing
number of hits
before and
after rewriting
with L3