

# Hall-D Online/Raw Data

David Lawrence, JLab

Feb. 10, 2015

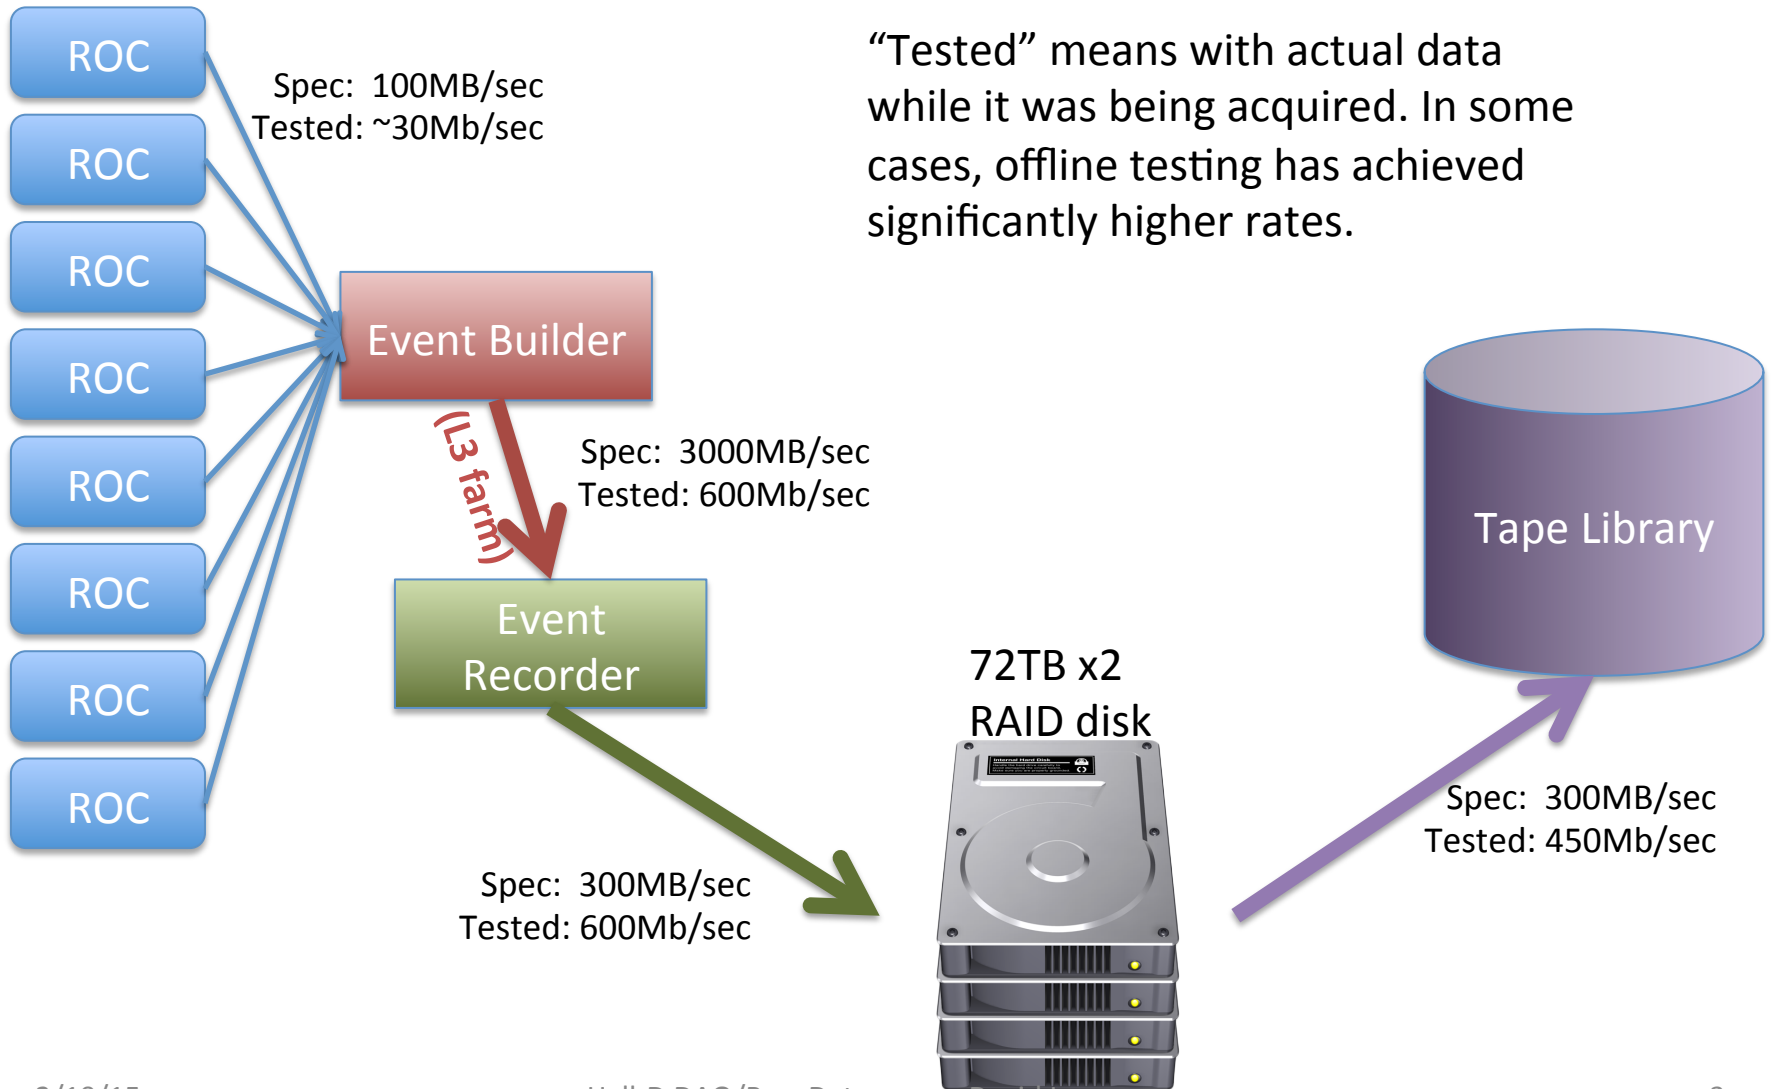
# Rough Specs. Review

- $10^8 \gamma/s$  on  $LH_2$  target  $\rightarrow$   $\sim 400\text{kHz}$  hadronic rate
- L1 trigger goal is to cut away  $\sim 50\%$  leaving  $200\text{kHz}$
- L3 trigger goal is to reduce by  $\sim 90\%$  leaving  $20\text{kHz}$
- Early simulation suggested  $\sim 15\text{kB/event}$
  
- Design specs\*:
  - $15\text{kB/event @ } 200 \text{ kHz} = 3000 \text{ MB/s}$  (front end)
  - L3 reduction by factor of 10 =  $300\text{MB/s}$  to RAID disk
  - 3 days storage on RAID =  $300\text{MB/s} * 3\text{days} = 78\text{TB}$
  - Maintain  $300\text{MB/s}$  transfer from RAID to tape

*\*L3 not officially part of 12GeV upgrade project*

# Data Rates


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




Example: some of the fADC250 word types

- 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

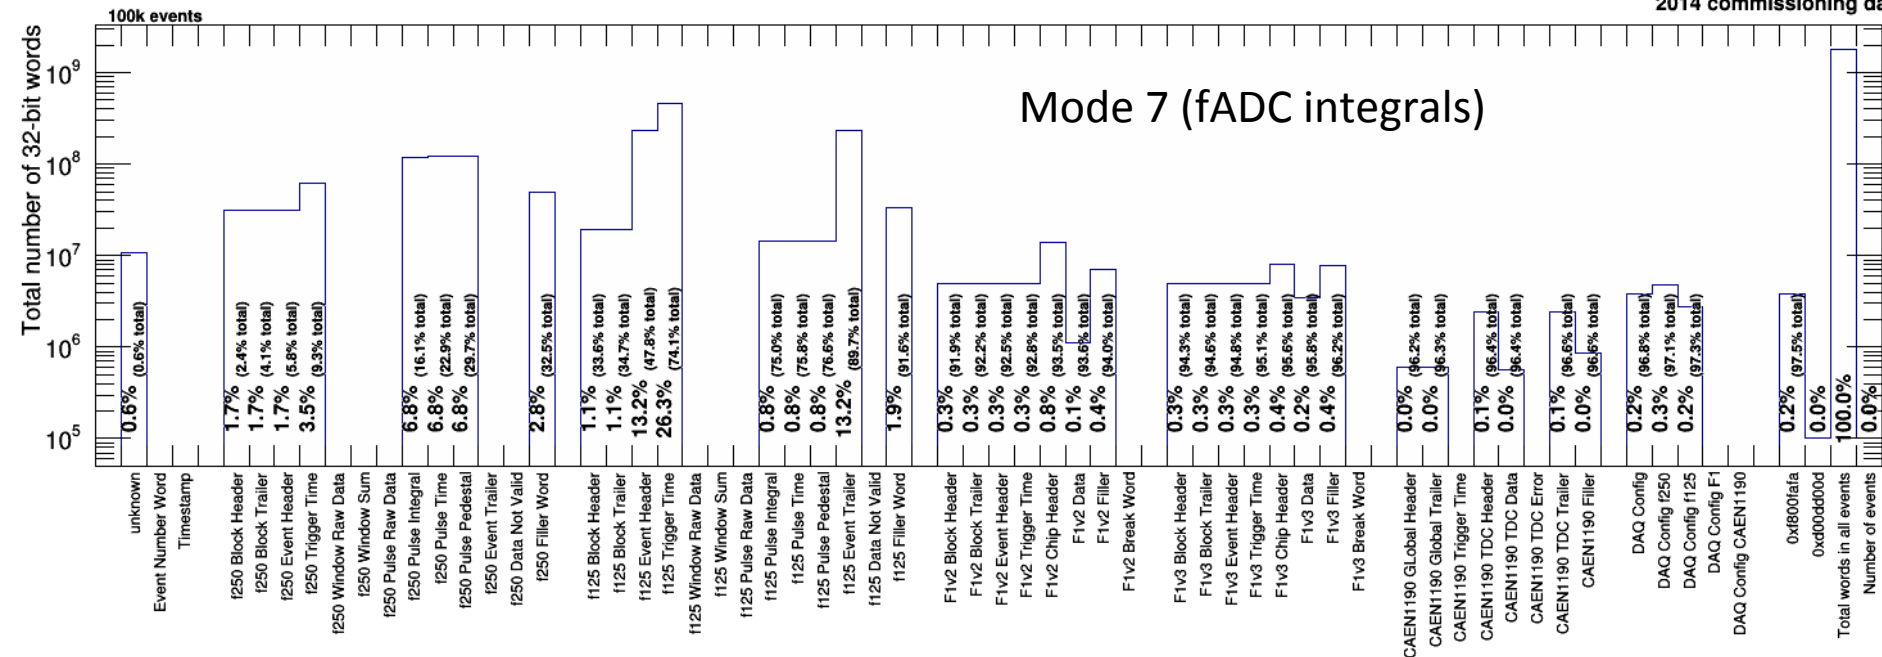
### Mode 7 (fADC Integrals)



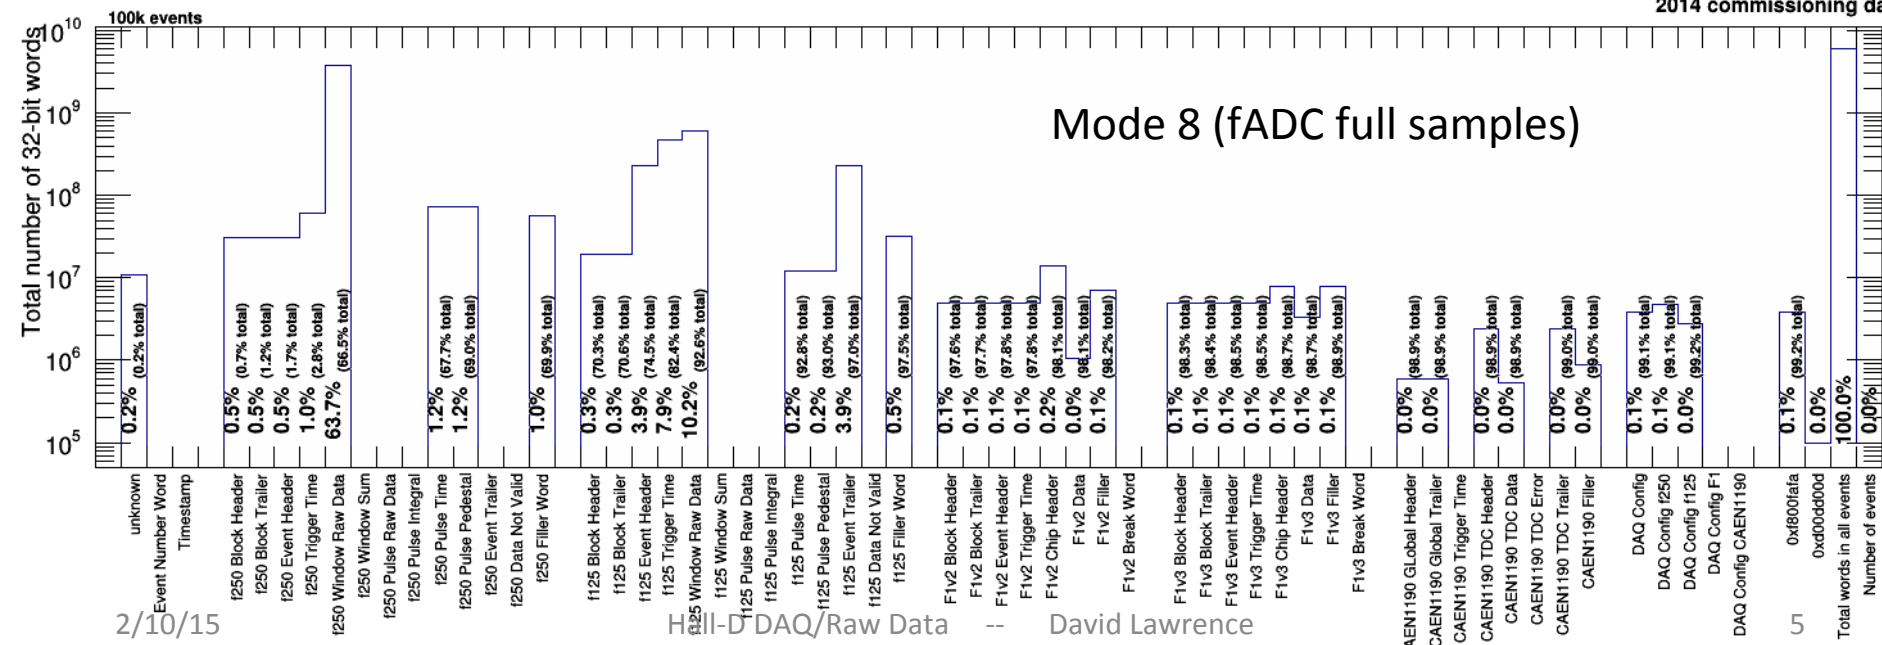
### Mode 8 (fADC full samples)




# Number of words in 100k events for run 2179



# Number of words in 100k events for run 2180



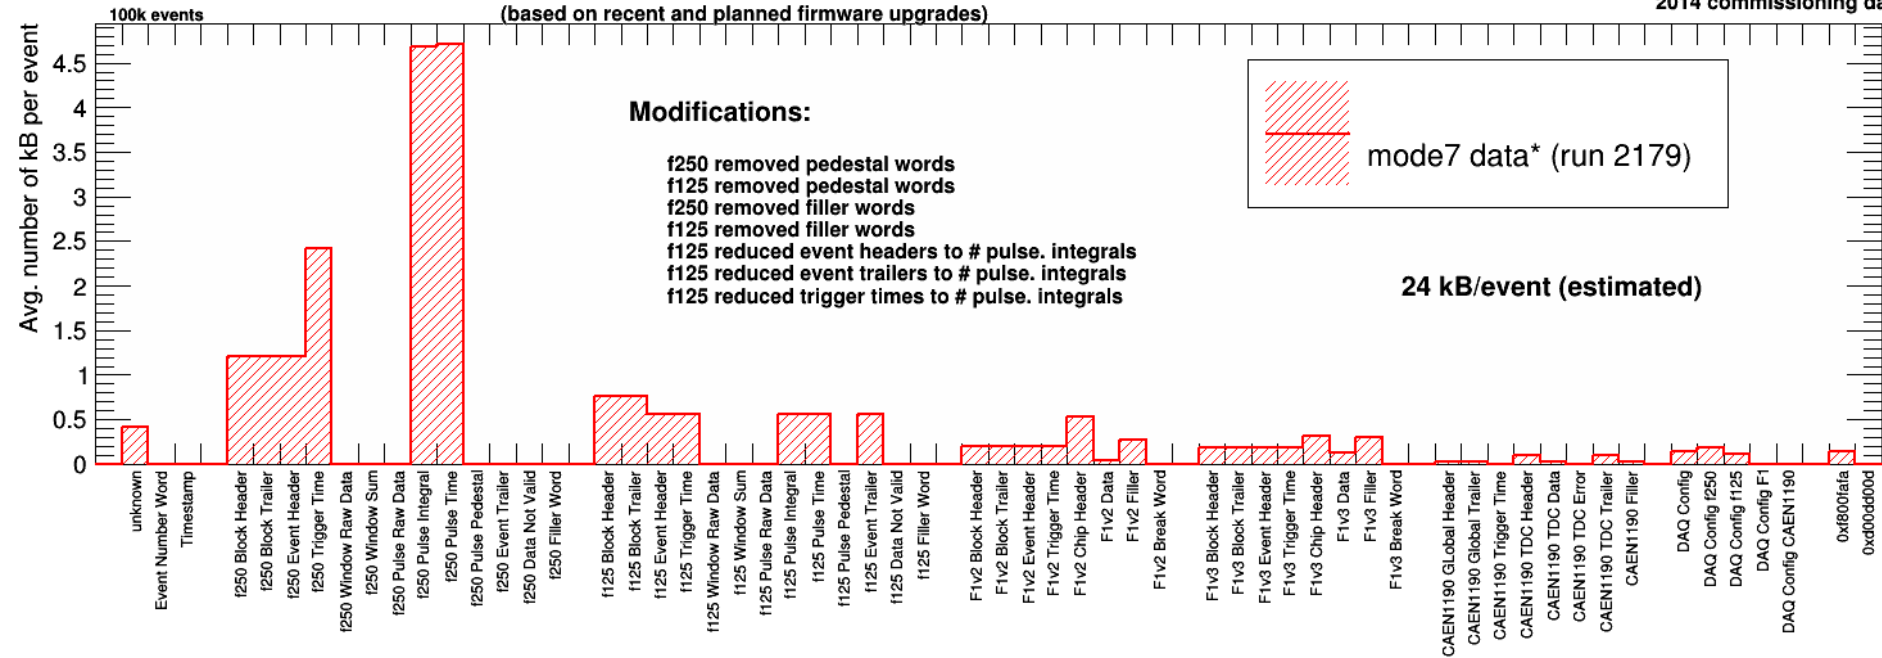
# mode7(integrals) and mode8(full samples) EVIO words/event



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

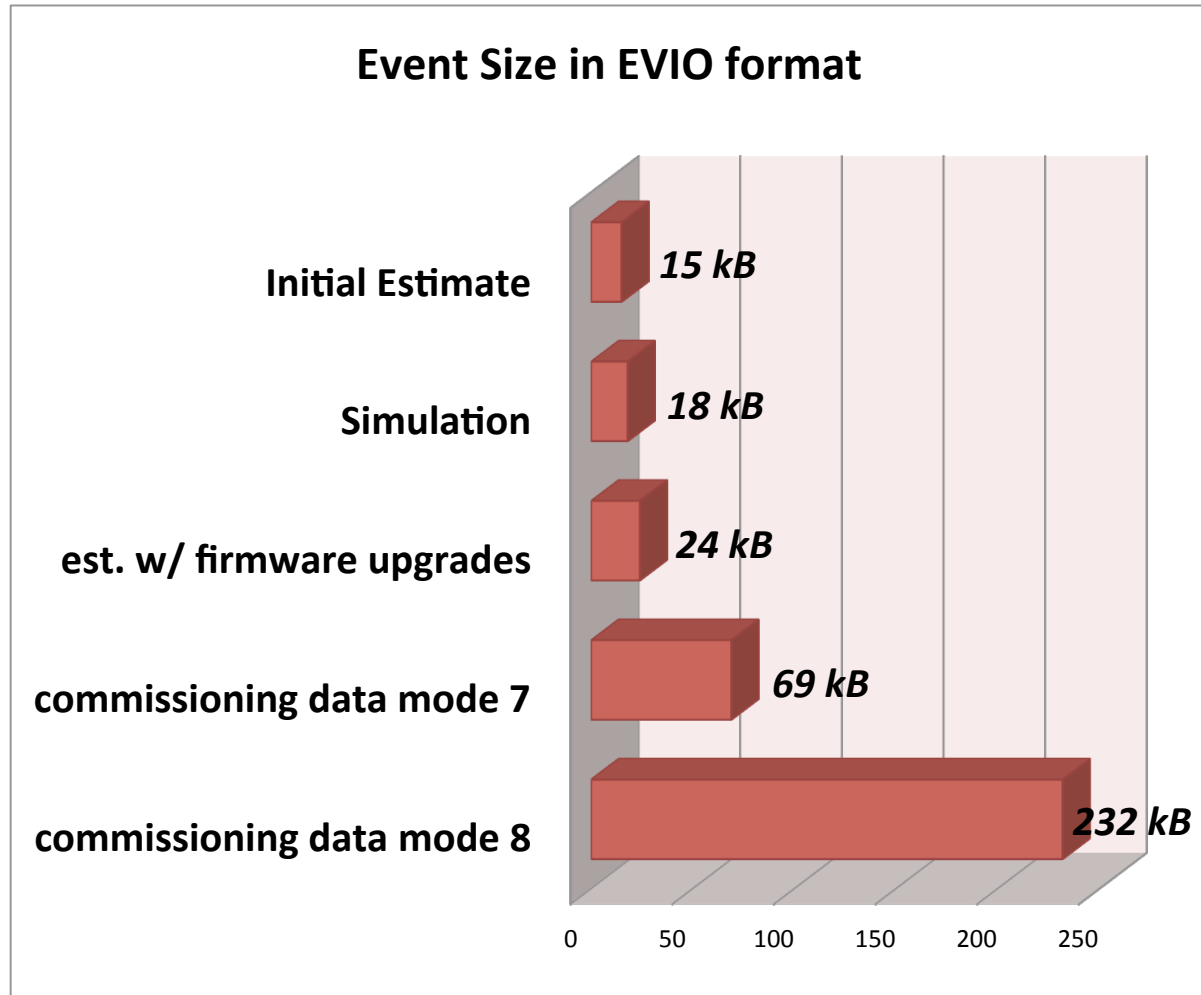
### Estimate of future production running EVIO words/event (based on recent and planned firmware upgrades)



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

# Event Size

- Simulation was consistent with initial estimate of event size
- Actual data was more than x4 larger
- Much of the data was taken in “raw” mode where fADC samples were saved





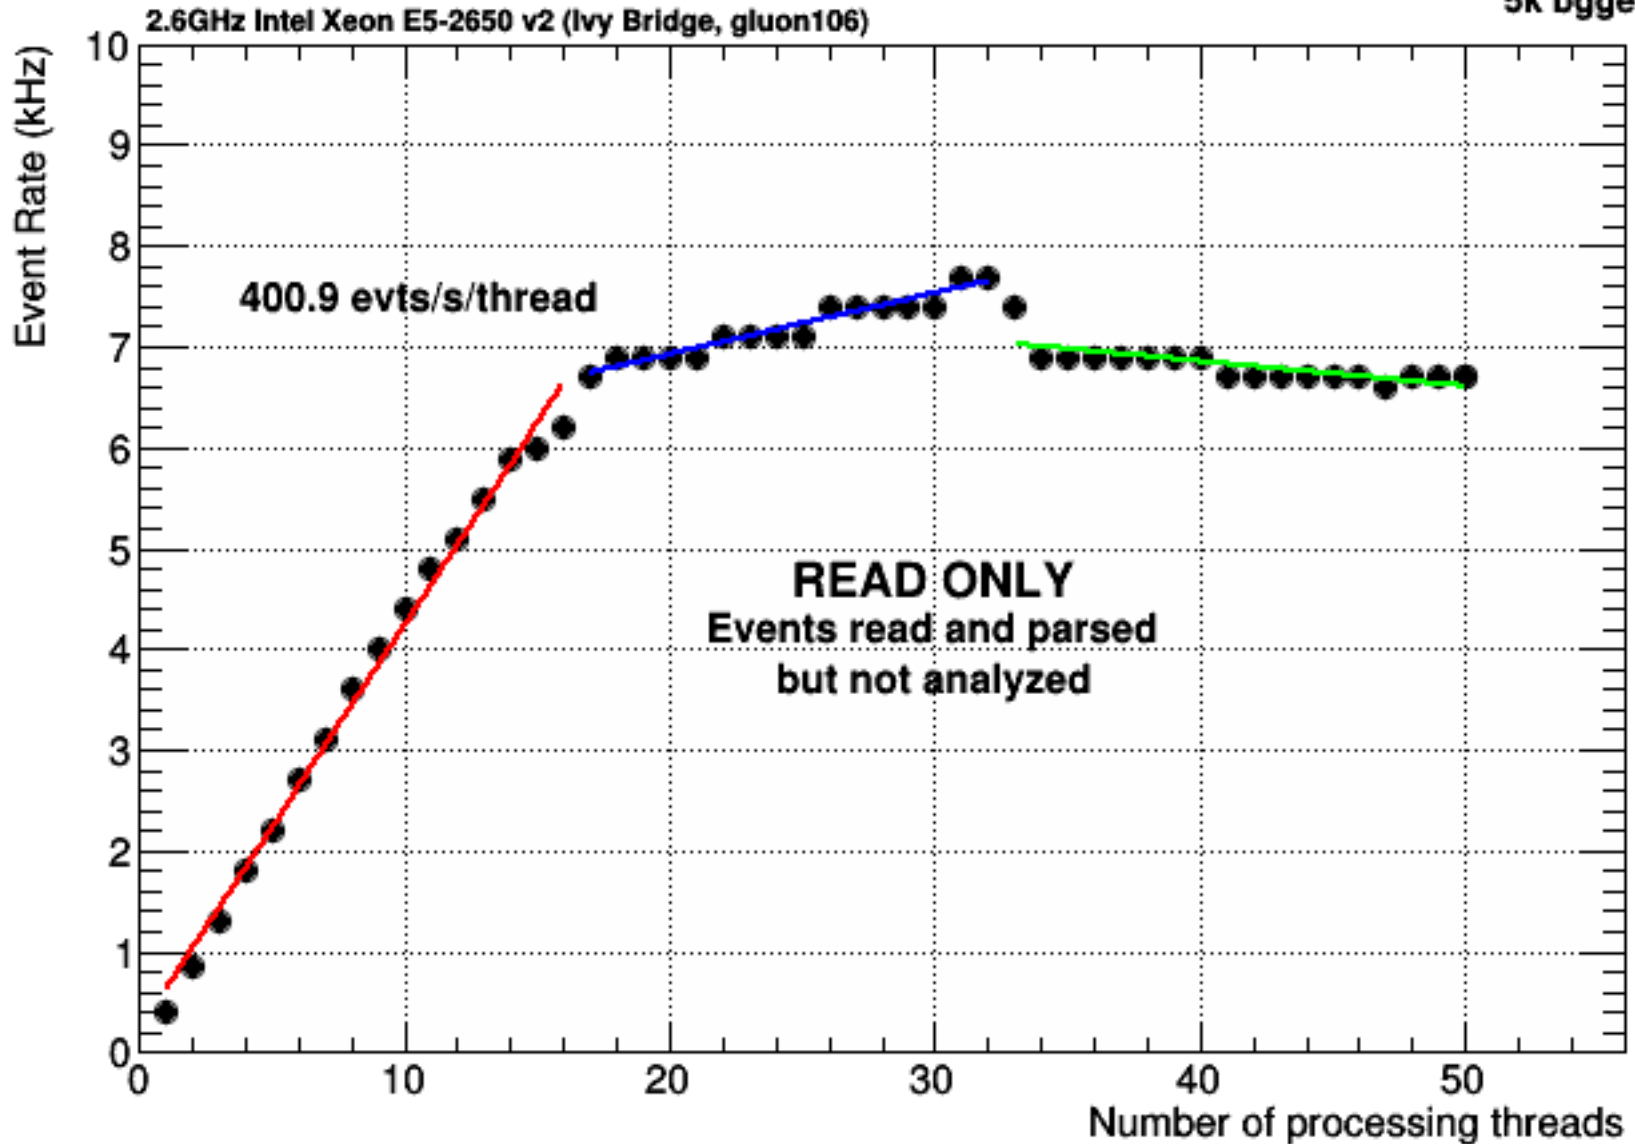
# EVIO Formatted Raw Data Files

- File format specified in detail by CODA group (<https://coda.jlab.org/drupal/system/files/coda/onlineFormat/eventbuilding.pdf>)
- Some corrupted events encountered
  - Problem due to race condition and only occurs for high rates. Has since been fixed in CODA.
  - Wrote new EVIO parser code
    - Error recovery (detects and skips bad blocks/events)
    - Mechanism to efficiently grow buffer size
- Event parsing including disentangling parallelization

```
EVIO Statistics for /cache/mss/hallD/RunPeriod-2014-10/rawdata/Run001715/hd_rawdata_001716_000.evio :  
-----  
  Nblocks: 53919  
  Nevents: 95302  
  Nerrors: 136  
Nbad_blocks: 76  
Nbad_events: 60
```

# Integrated rate vs. number of threads

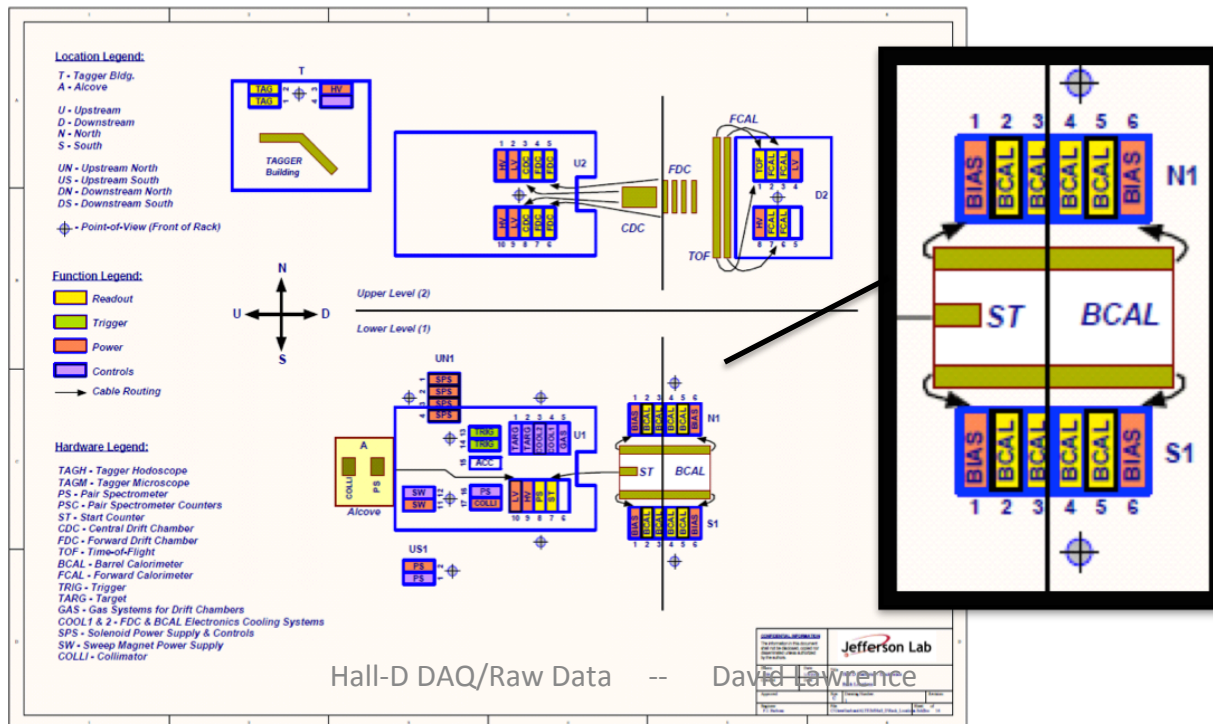
January 23, 2015 DL  
sim-recon: svn 17000, JANA svn 2115  
5k bggen events




sim-recon-commissioning bggen data in EVIO

# DAQ to Detector Translation Table


- The Translation Table is used to convert from DAQ system coordinates (rocid, slot, channel) into detector-specific coordinates (e.g. BCAL module, layer, sector, end)
- ~23k channels defined in SQLite DB file
- Stored in CCDB as XML string for offline analysis with complete history:
  - /Translation/DAQ2detector



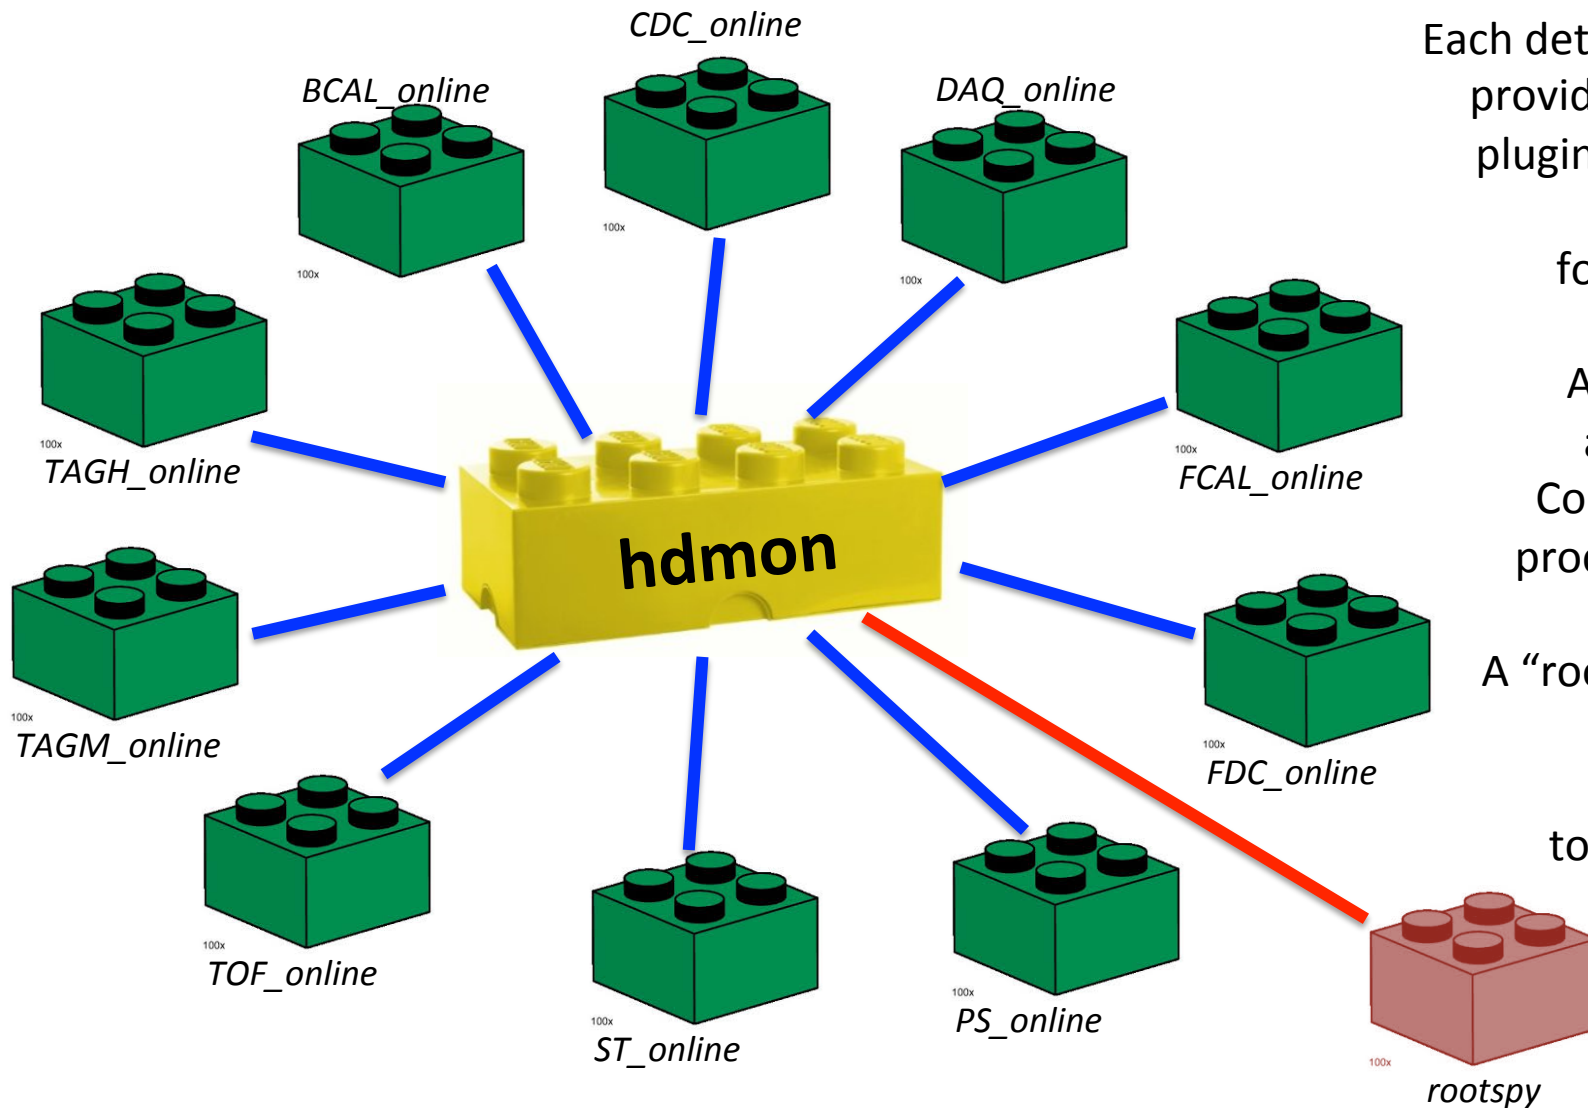
### hits per event by VME crate for run 2179



### kB per event by VME crate for run 2179



# Monitoring Plugins



Each detector system provides 1 or more plugins that create histograms for monitoring

All plugins are attached to a Common DANA process (*hdmon*)

A “rootspy” plugin publishes all histograms to the network

# Summary

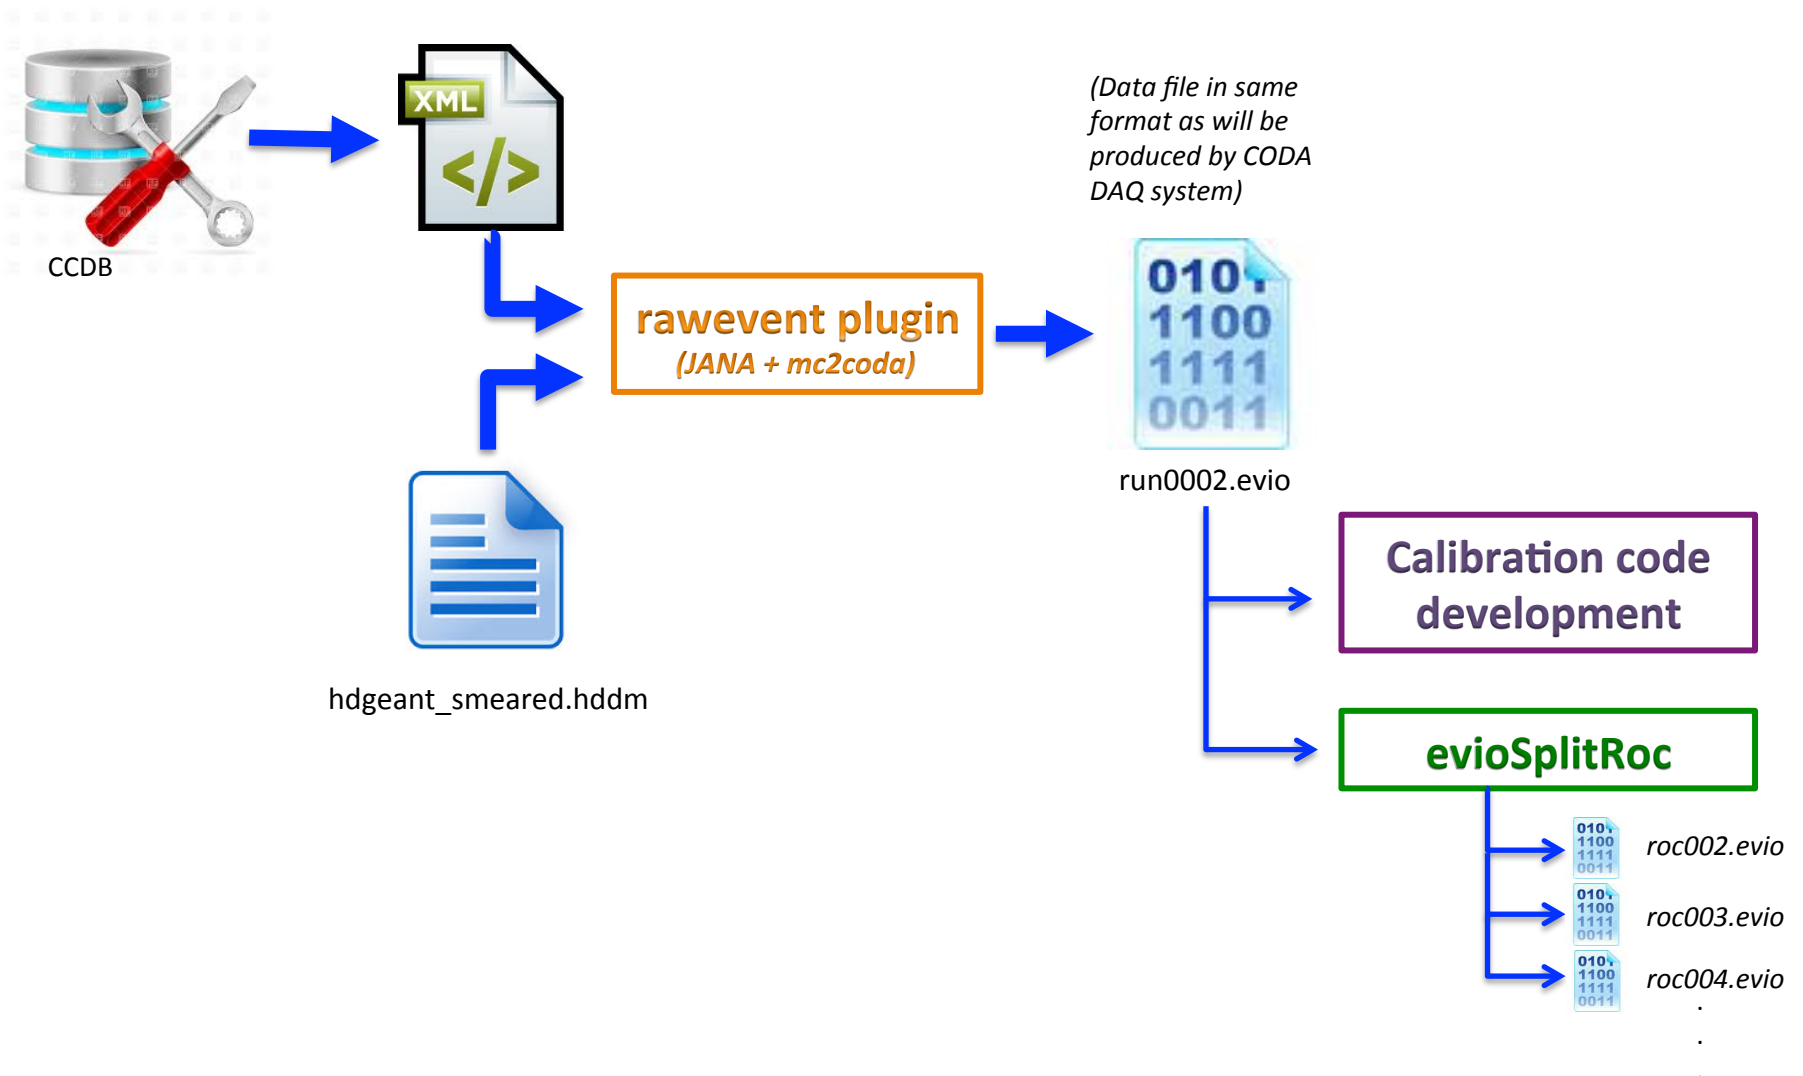
- 120 TB written to RAID and copied to tape
- 600 MB/s written to RAID from DAQ while taking data
- 450 MB/s copy from RAID to tape
- Event size increased from original estimate, but not all compression options exhausted
- Event parsing and application of Translation Table implemented and in routine use by collaboration

# Backup Slides



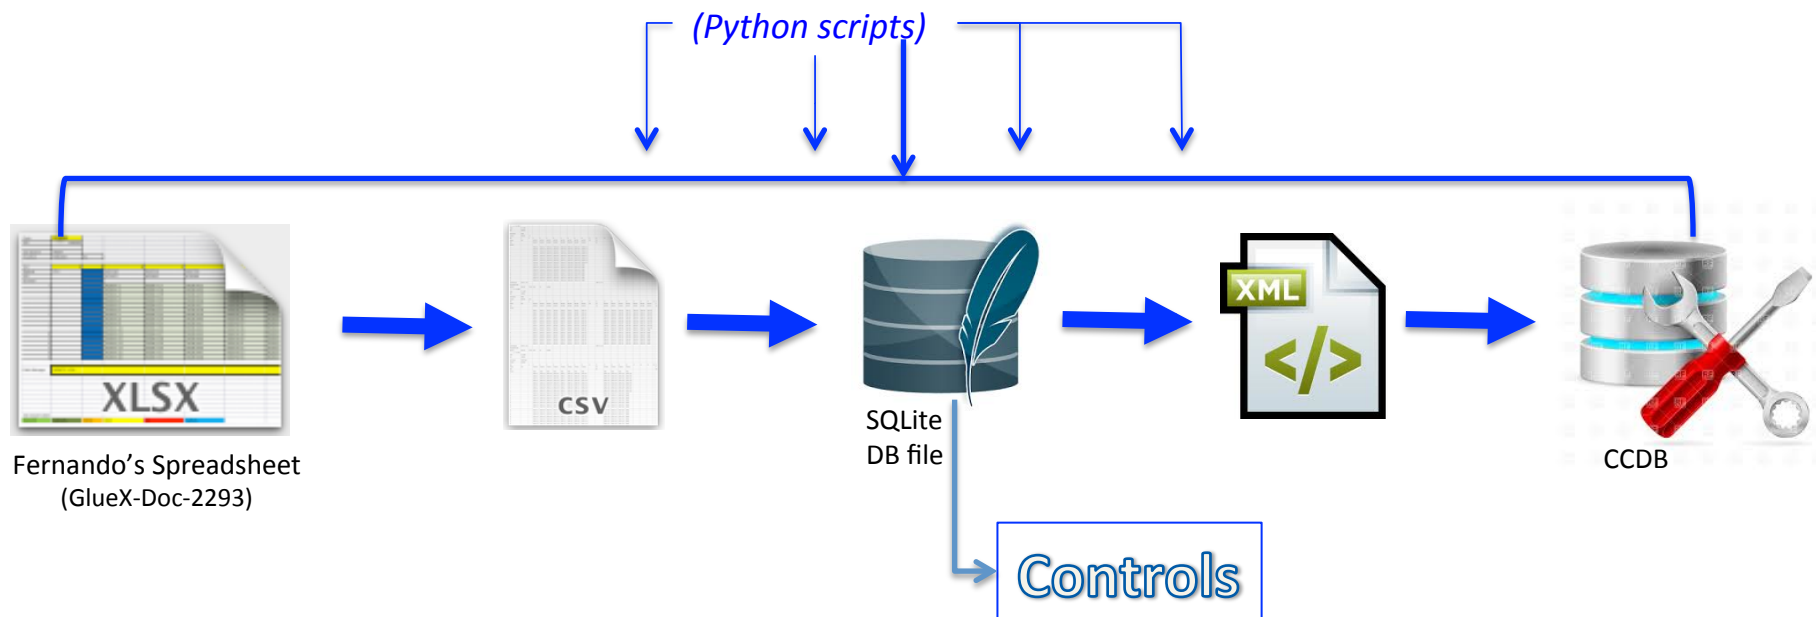
# Raw Data Formatted Files

*(from simulated data)*

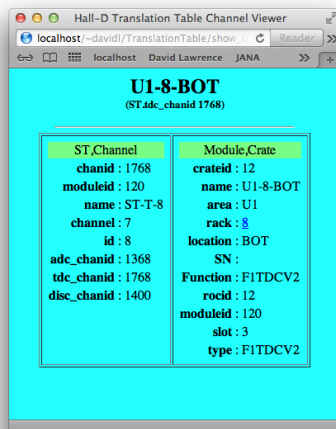




# Translation Table



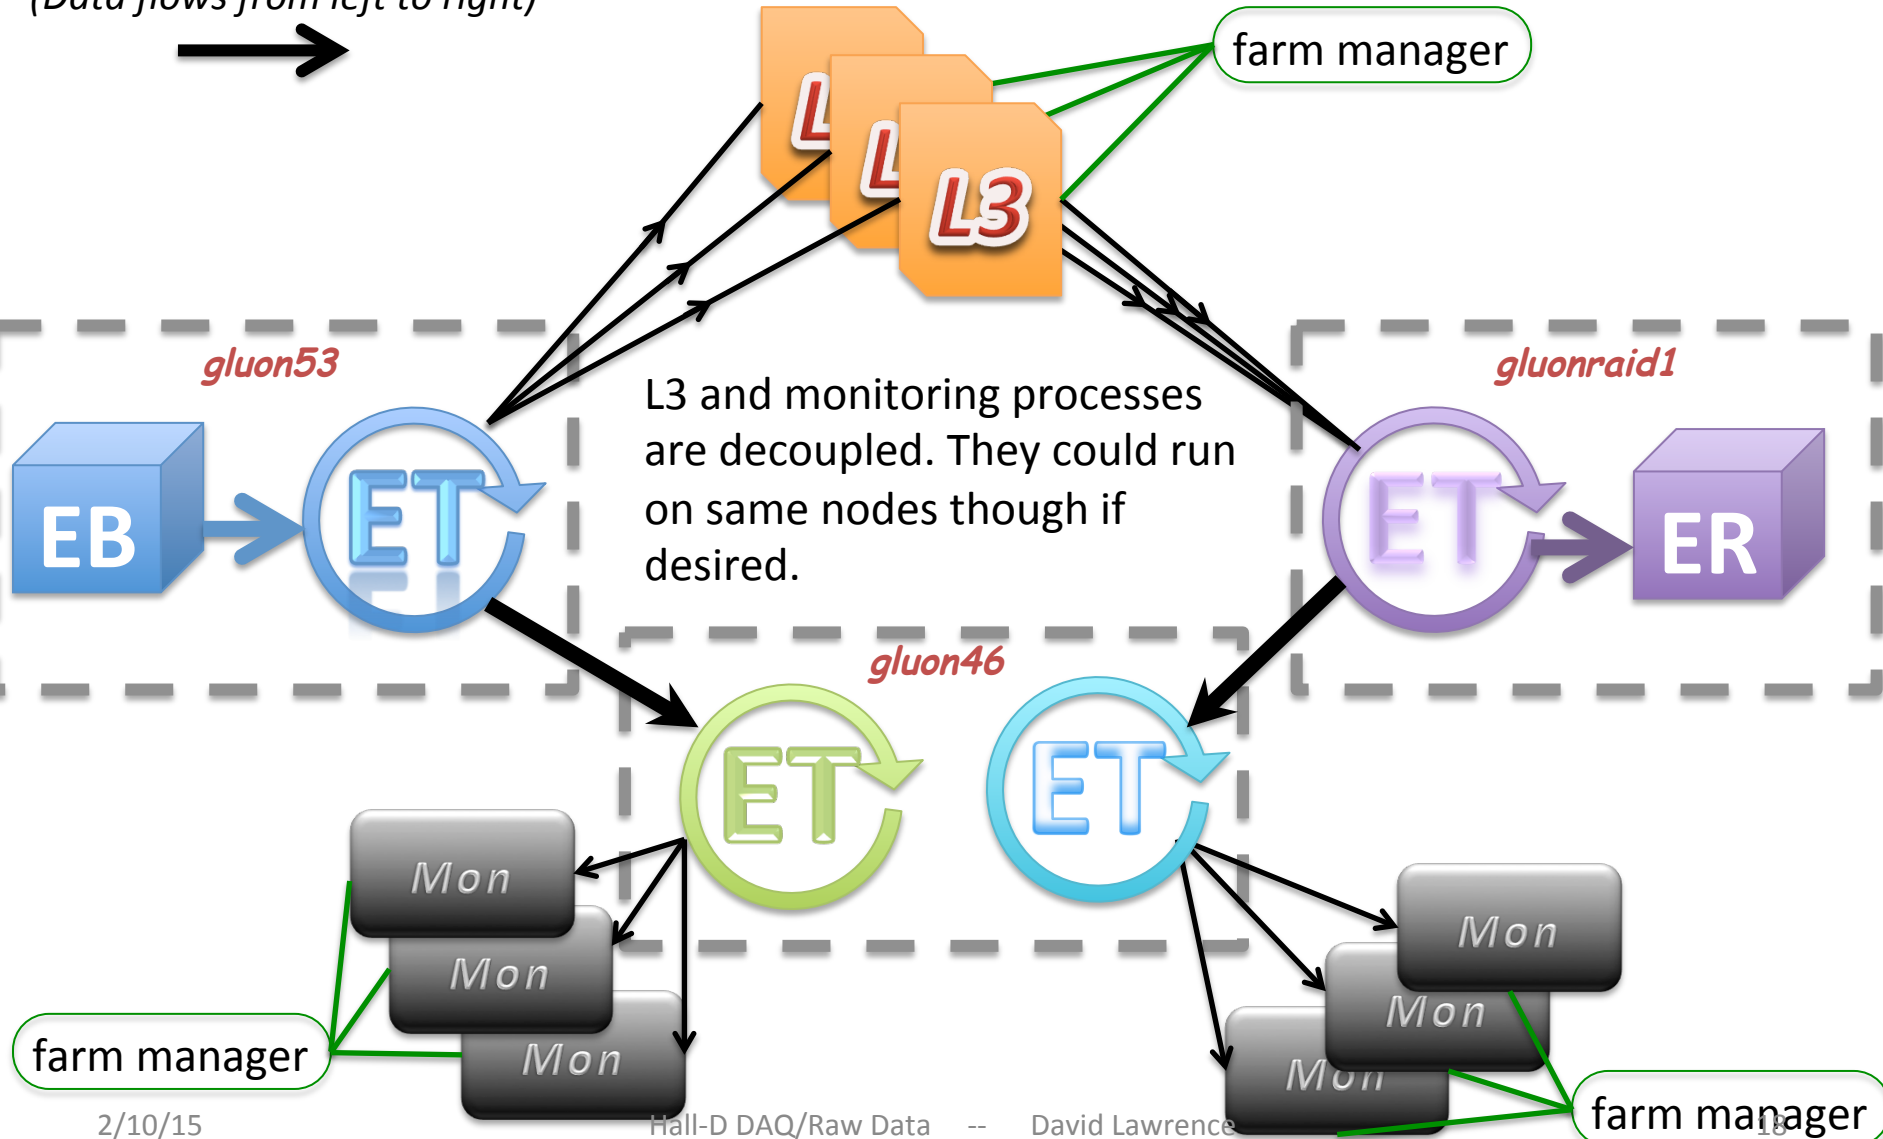
- PHP web pages set up to browse the SQLite DB and draw some simple graphics to display the rack/crate/module layout



<https://www.jlab.org/Hall-D/test/TranslationTable>

# L3 and monitoring architecture

(Data flows from left to right)



# hdmongui

## Hall-D Data Monitoring Farm Status

UDL: cMsg://gluondb1/cMsg/janactl

name	node	Nthr	Nevents	rate (Hz)
gluon100.jlab.org_1064	gluon100.jlab.org	32	153946	675.7
gluon101.jlab.org_16938	gluon101.jlab.org	32	202377	930.9
gluon102.jlab.org_16954	gluon102.jlab.org	32	145899	683.8
gluon105.jlab.org_27603	gluon105.jlab.org	32	58545	264.9
gluon106.jlab.org_9217	gluon106.jlab.org	32	58579	273.7
gluon107.jlab.org_22081	gluon107.jlab.org	32	58563	272.2
gluon108.jlab.org_5593	gluon108.jlab.org	32	145867	655.7
gluon109.jlab.org_4445	gluon109.jlab.org	32	198906	904.5
gluon110.jlab.org_21543	gluon110.jlab.org	32	58539	262.8
gluon111.jlab.org_23740	gluon111.jlab.org	32	58567	273.8
gluon104.jlab.org_11666	gluon104.jlab.org	32	118622	686.0

*processes run  
multi-threaded*

Totals

Total Rate: **5884.2 Hz**

Total Events: 1,258,410

Total Nodes: 11

Total Threads: 352

Node Details

Name: gluon102.jlab.org\_16954

Node: gluon102.jlab.org

Threads: 32

Event Source: ET:/tmp/  
et\_sys\_davidl:MONTEST:gluonraid1  
(type=JEventSource\_EVIO)


command line: hdmmon --config=/home/davidl/HallID/ONLINE/daq\_pro\_vers/  
daq/config/monitoring/hdmonTEST.conf ET:/tmp/  
et\_sys\_davidl:MONTEST:gluonraid1

Event Source:  COOL User specified: ET:/tmp/et\_sys\_davidl:MONTEST:gluonraid1

Level: TEST   
*multiple "levels" supported*

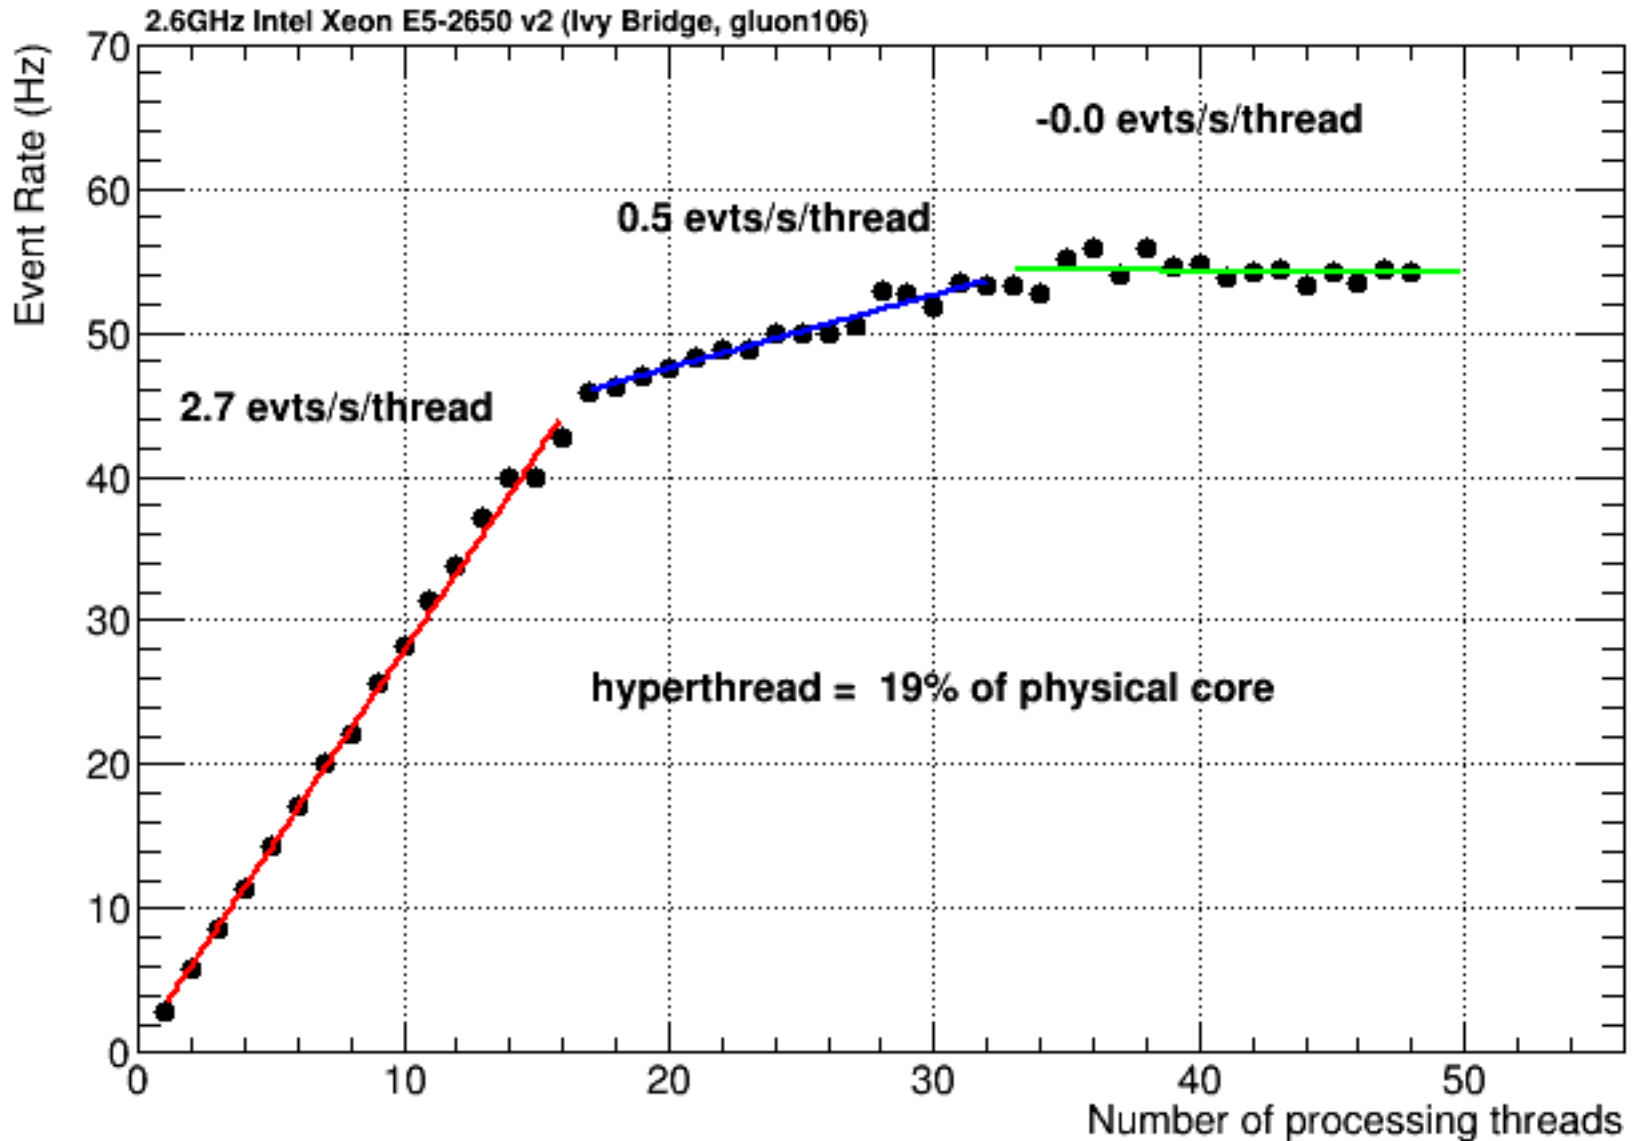
# Memory Usage vs. Num. Threads

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



# Integrated rate vs. number of threads

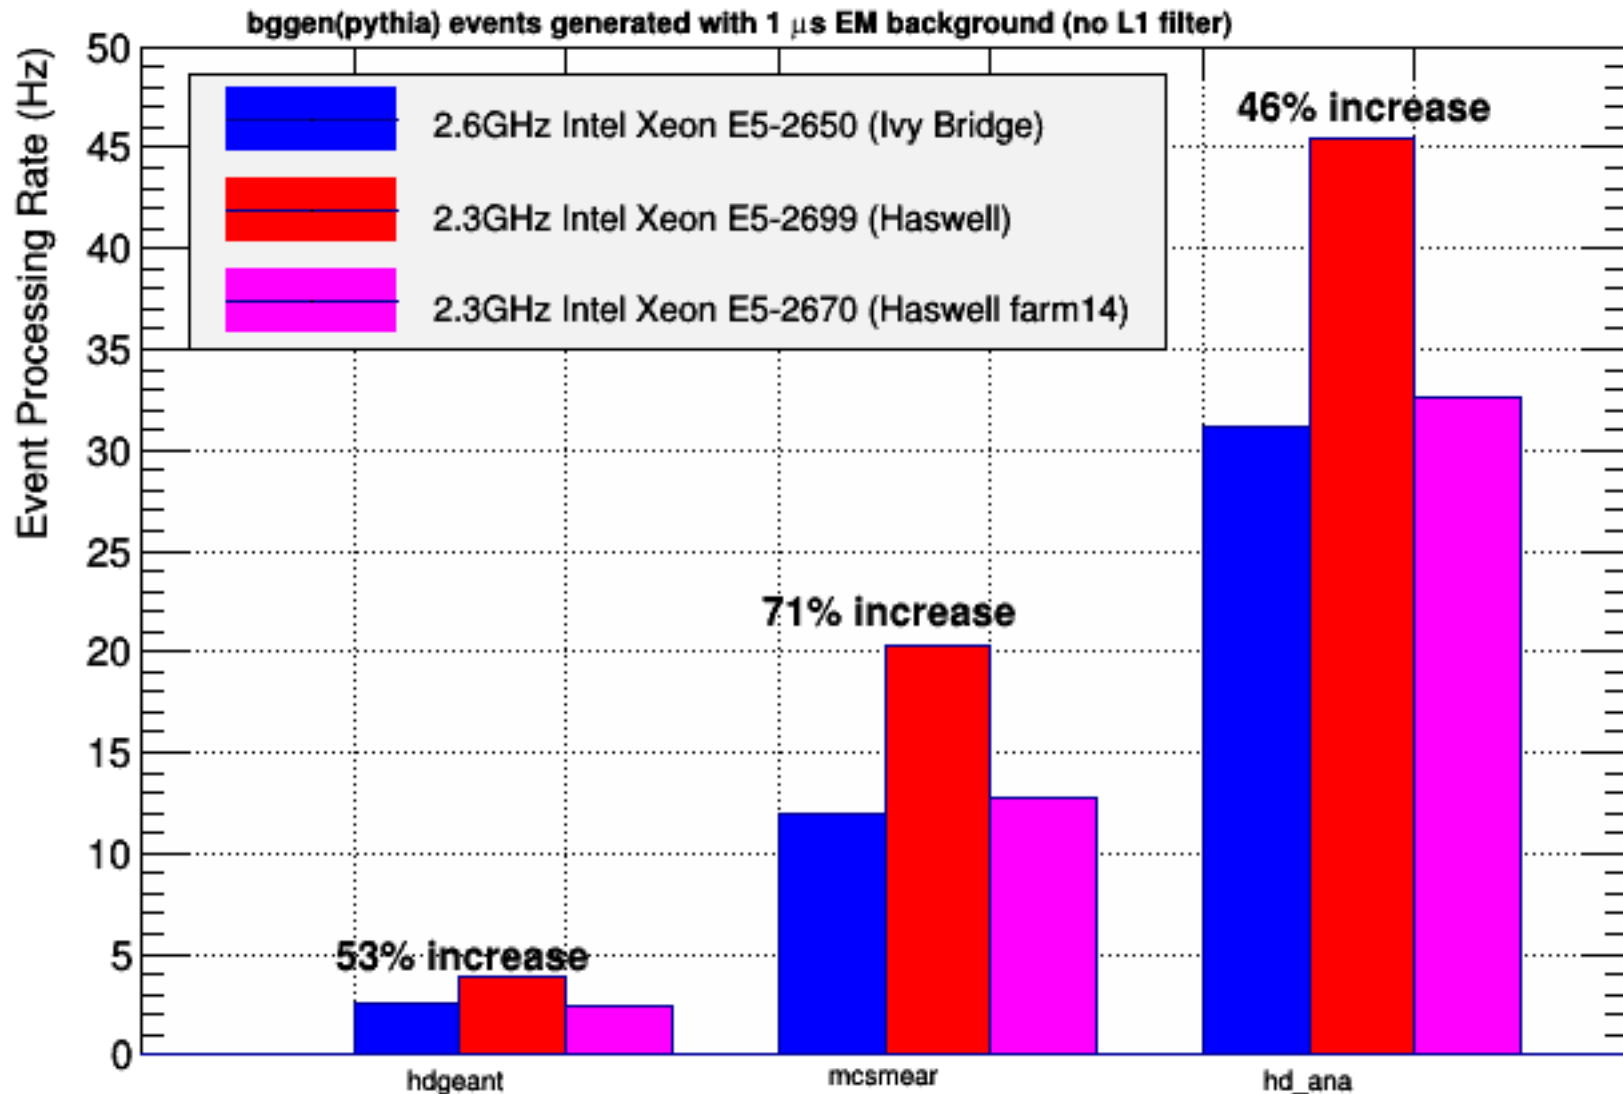
January 22, 2015 DL  
JANA svn 2115 (pre0.7.3)



JANA TestSpeed with 100k GOVERNOR\_ITERATIONS

# Single Thread Processing Rates

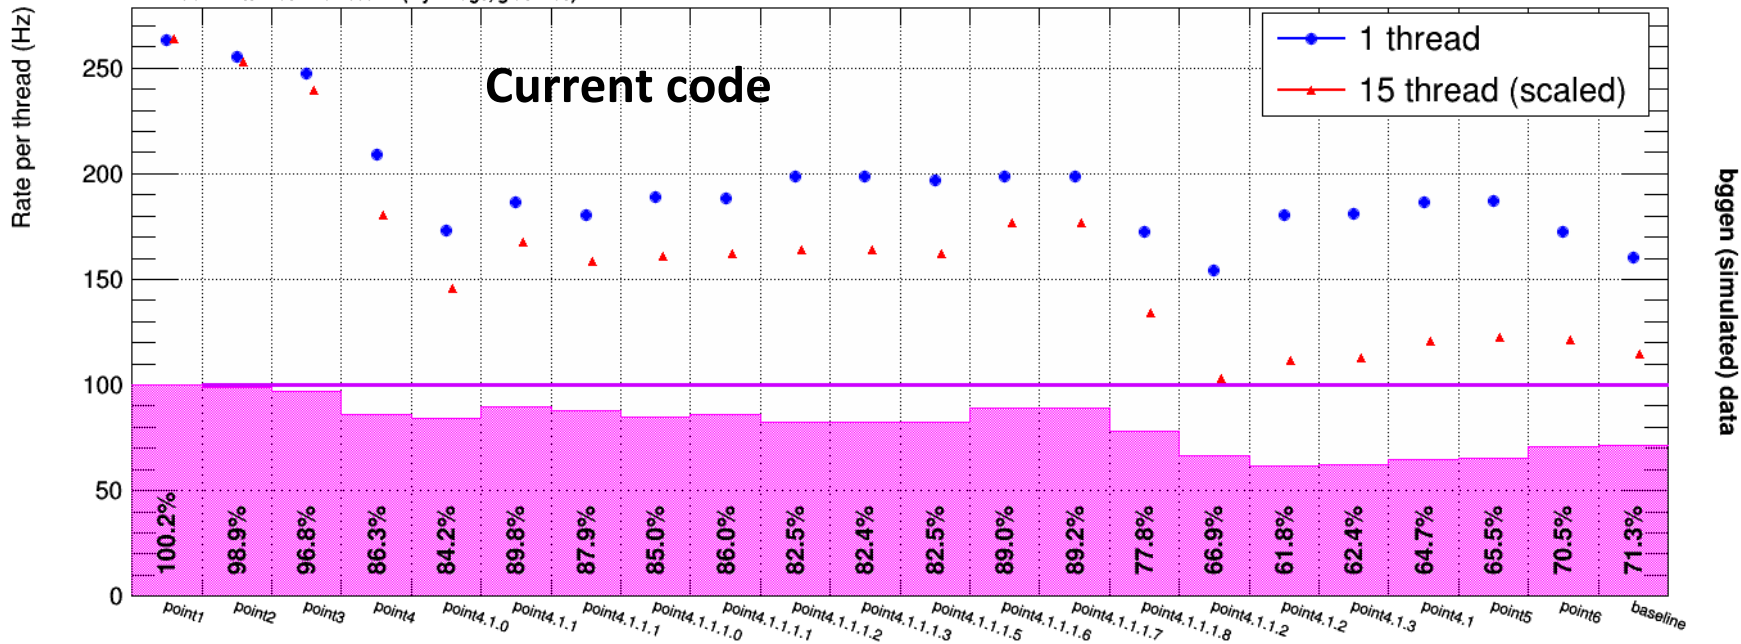
January 16, 2015 DL  
svn revision 14114



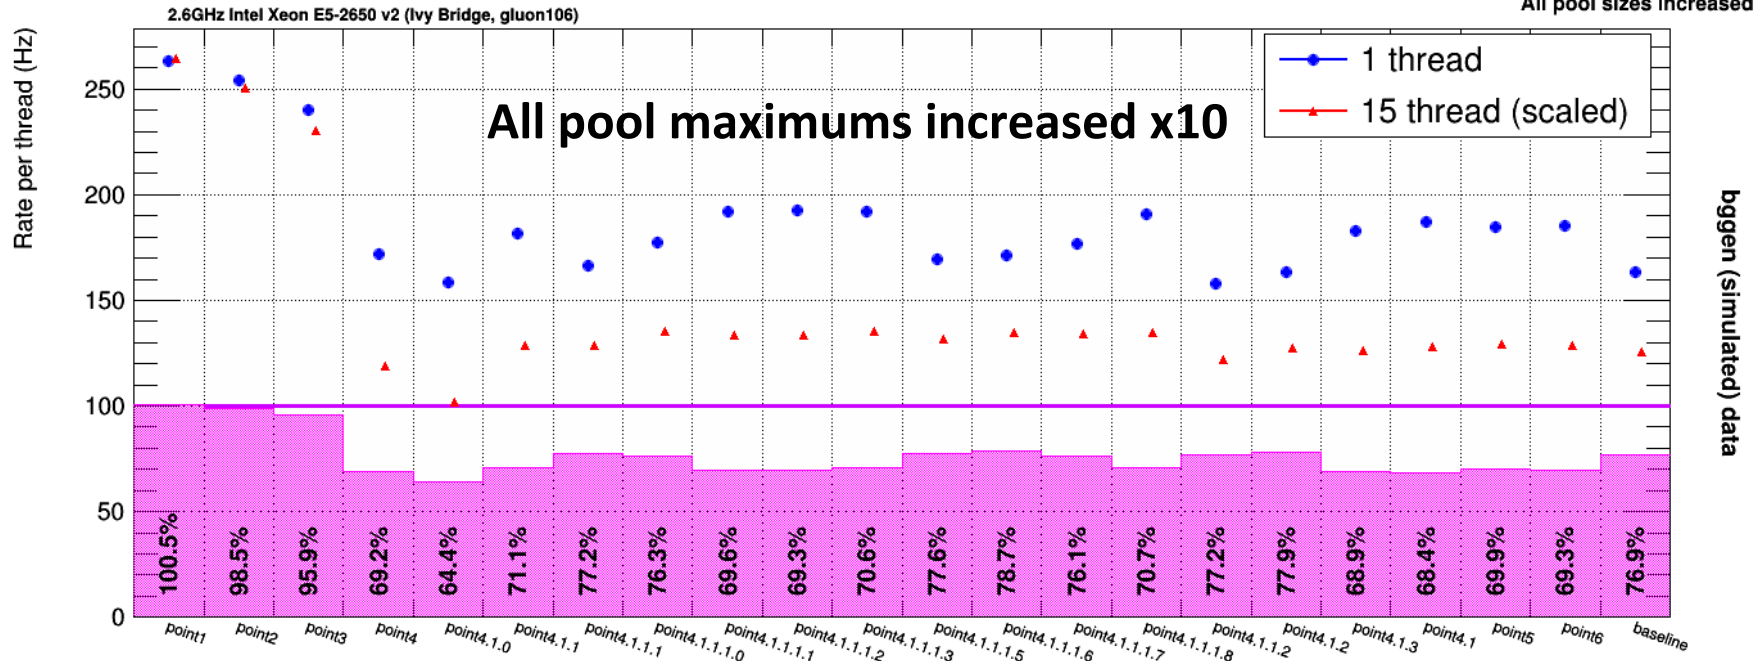
# DTrackCandidate:CDC rate vs. point in code

January 28, 2015 DL  
 sim-recon: svn 17000, JANA svn 2115  
 All pool sizes nominal

2.6GHz Intel Xeon E5-2650 v2 (Ivy Bridge, gluon106)



# DTrackCandidate: CDC rate vs. point in code



# DTrackCandidate: CDC rate vs. point in code

