Making Computing Cost estimate for the HPS runs

at the HPS weekly
30 August 2012

H. Neal

Costing

- Processing cost mostly go into general infrastructure costs and somewhat adjustable by the time allowed for the processing to complete and use of multiple sites/clouds/GRID. Greatest issue may simply be the availability of cycles.
- Storage costs: usually directly charged to the project
  - Tapes
    - Generally all data is stored on tape
  - Disk
    - Stage pool large enough to prevent thrashing the tape drives
  - Infrastructure
    - Data movers, xrootd servers

Data to be processed and distributed in the FULL 180 day HPS run:

Raw Detector data:

290 Tbytes
(assuming 26->31 KHz trigger rate, 5KByte event size, 1/8th data flow passing L3)

Simulated data:

435 Tbytes
(assuming 1.5x event size but only 10% events in comparison to detector data)

Foresee storing two passes of the data and simulation output for only the equivalent of 10% of the # data events one gets **0.62 PB** of storage

T10Kc tapes it will cost about $50/TB or $31K

We could survive with 1/5th of this on disk ==> 150 TB of disk (90 TB server costs ~$45K)

Staged Approach or Not

- Is it best to prepare for the commissioning run and then the full run or go ahead and prepare for the full run starting with acquisitions in 2014?
  - Tapes don't all need to be bought at once nor does having them all help with preparation of the final computing model
  - Disk servers need further consideration. Acquire a server with space for adding in complement ultimately needed. New SLAC computing model may avoid this detail of planning; we ask the CD … they provide the solution.

Storing it all on Disk

On Wed, Aug 29, 2012 at 01:53:12PM -0700, Homer wrote:
> Could you please give me a rough estimate for how much it would cost
> to purchase 0.5 PB of storage next year.

Hi Homer,

If you use 2TB disks and put only one array on each server, you'd need 6 building blocks for ~$246K.

If you use 3TB disks and put only one array on each server, you'd need 4 building blocks for ~$200K.

If you use 3TB disks and put one array + one expansion on each server, you'd need 2 building blocks for ~$178K.

All configs yield about 518TB of usable space. The 2TB config would yield highest performance since it's spread across 6 servers and 6 arrays. The last config is a capacity config since you have all I/O traffic traversing 2 servers.

We can adjust performance and price based on your requirements. If this is NFS space, then we probably want to consider an advanced file system like GPFS (additional cost). For xrootd, you might prefer the usual standalone configs.

Lance

Test run reality:

Transferred LCIO raw data files AND "recon" LCIO files w/raw data included

retransferred at least 3 times

The most time consuming part was just transferring the data.

	File Sizes					#bytes/event	
	evio	slcio	rec		#events	slcio	rec
1342	365416	480196	124011				
1343	233174	424981	99134				
1344	352190	427293	96300				
1345	1261416	470946	122111				
1347	627330	501609	125832				
1349	1974150	495217	123371		43744	11321	2820
1351	1275828	490588	123499				
1353	1991741	520732	121084		37545	13870	3225
1354	1990072	505578	120828				
1358	785890	494833	125832				
1361	47478	211818	53425				
1363	454513	486612	123499				
1364	530909	468358	124139				
1365	134436	466454	124395				

All types were produced at JLAB and transferred to SLAC.
For the full run do LCIO conversion at JLAB and all recon at SLAC/UNH?

File Size by Data Type

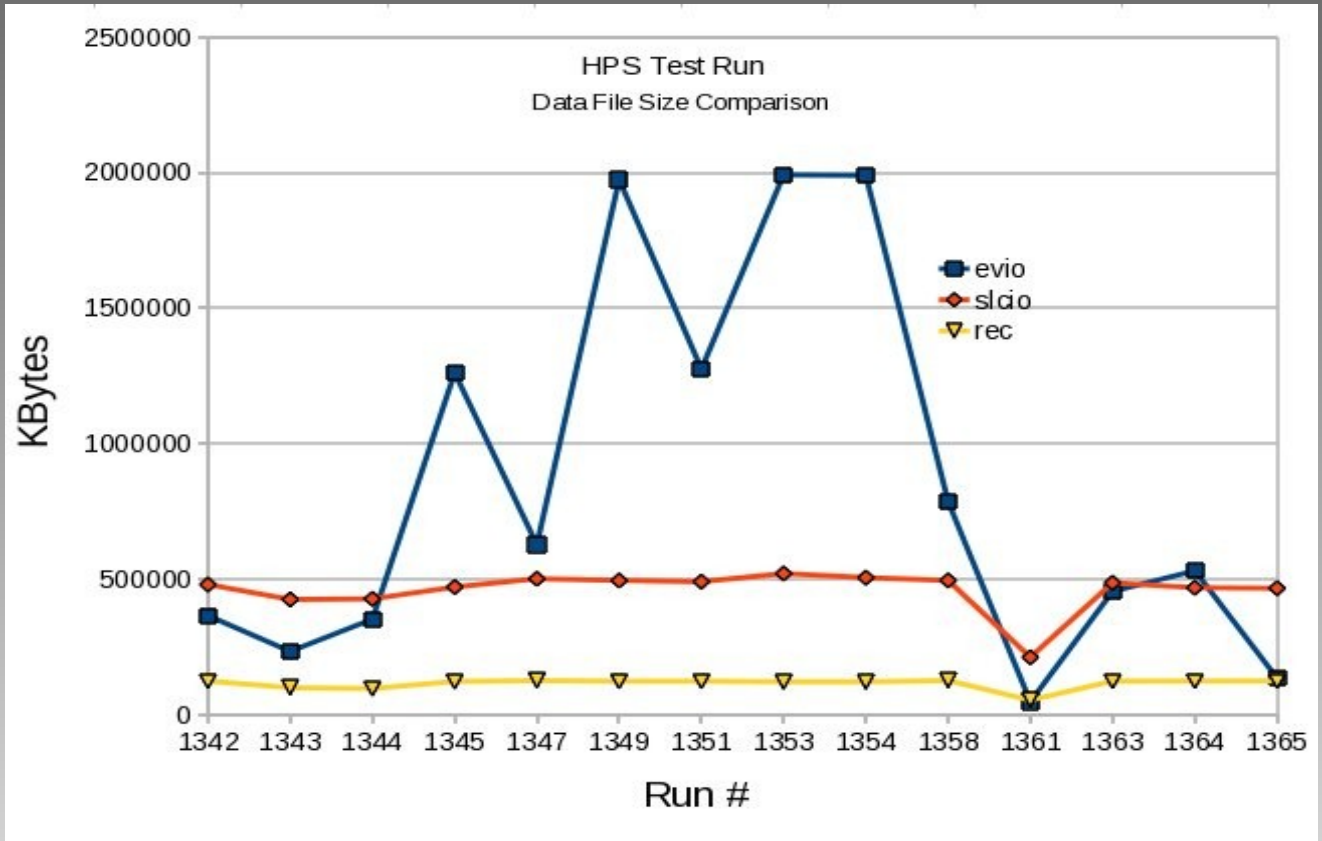

1971936 hps_001349.evio.0
hps_001349.evio.0: 34.9% -- replaced with hps_001349.evio.0.gz
1284332 hps_001349.evio.0.gz

RAW SLCIO Contents

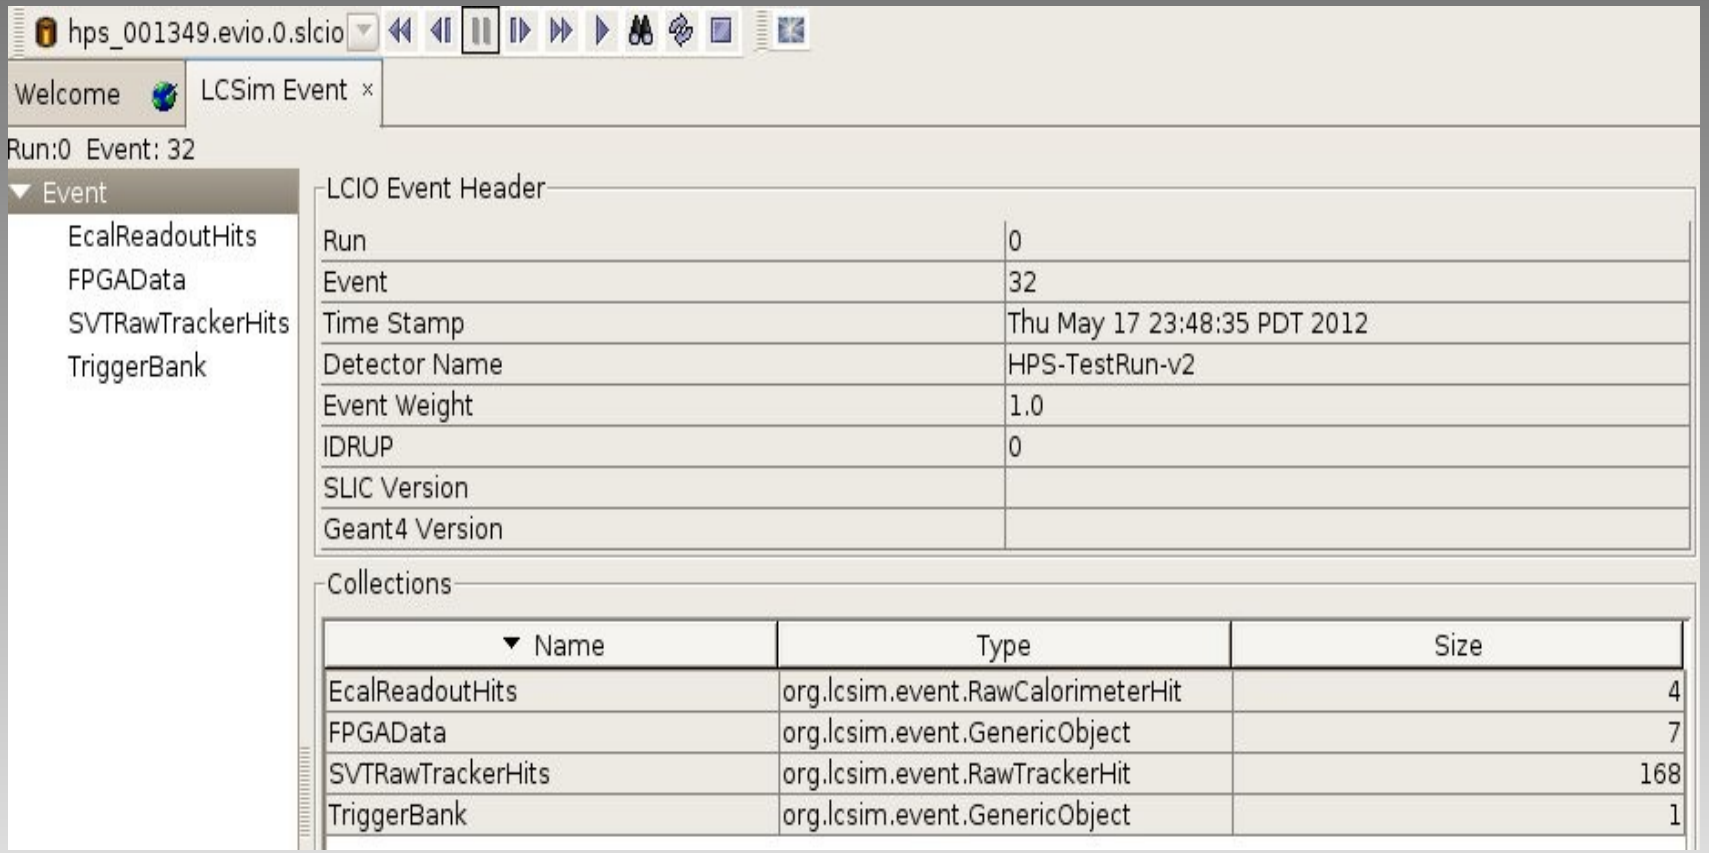

Reco SLCIO Contents

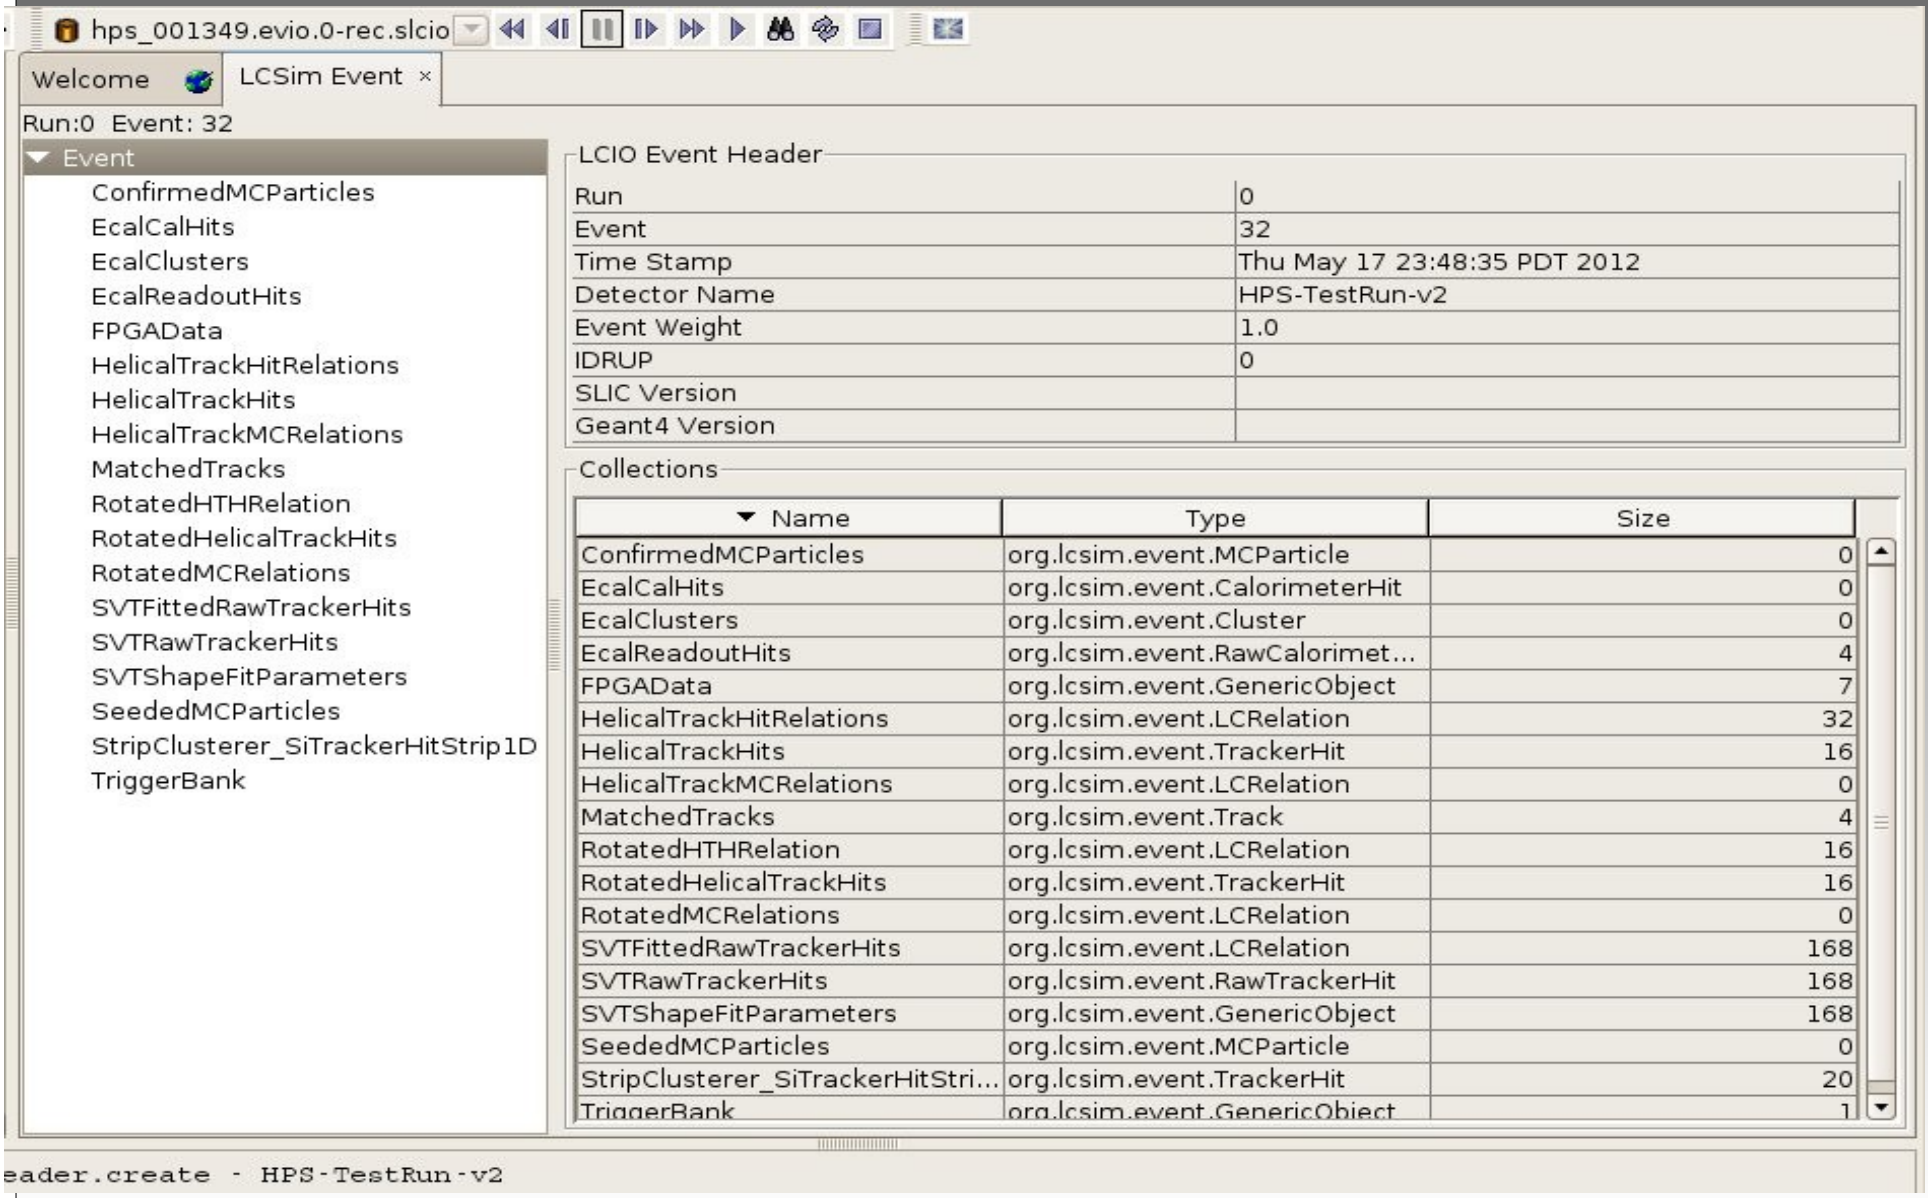

Banks (Tracks and ECAL Hits)

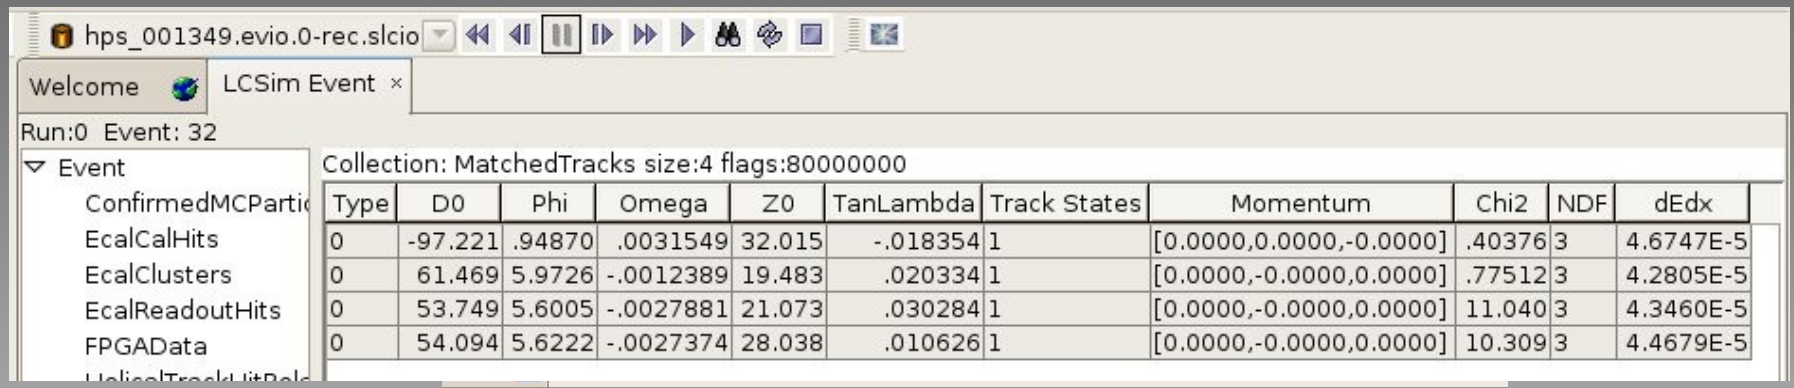

Collection: MatchedTracks size:4 flags:80000000

Type	D0	Phi	Omega	Z0	TanLambda	Track States	Momentum	Chi2	NDF	dEdx
0	-97.221	.94870	.0031549	32.015	-.018354	1	[0.0000,0.0000,-0.0000]	.40376	3	4.6747E-5
0	61.469	5.9726	-.0012389	19.483	.020334	1	[0.0000,-0.0000,0.0000]	.77512	3	4.2805E-5
0	53.749	5.6005	-.0027881	21.073	.030284	1	[0.0000,-0.0000,0.0000]	11.040	3	4.3460E-5
0	54.094	5.6222	-.0027374	28.038	.010626	1	[0.0000,-0.0000,0.0000]	10.309	3	4.4679E-5

Collection: EcalHits size:48 flags:e0000000
CellIDEncoding: system:0:6,layer:6:2,ix:8:-8,iy:16:-6

id: system	id: layer	id: ix	id: iy	raw energy (GeV)	corrected energy (GeV)	X (mm)	Y (mm)	Z (mm)	time (ns)
13	0	-23	-1	.23761	.23761	-300.15	-44.675	1600.9	7.4579
13	0	-23	-3	.0014820	.0014820	-300.15	-74.688	1601.0	7.6507
13	0	-21	-2	.0013452	.0013452	-268.69	-59.678	1601.7	7.7393
13	0	-19	2	8.6486E-4	8.6486E-4	-237.56	59.678	1602.4	8.0538
13	0	-22	-1	.023063	.023063	-284.38	-44.675	1601.3	7.5523
13	0	-21	-1	.0015793	.0015793	-268.69	-44.675	1601.6	7.5985
13	0	-15	3	.0013680	.0013680	-176.10	74.688	1603.5	8.1722
13	0	-23	-2	.0090691	.0090691	-300.15	-59.678	1601.0	7.5609
13	0	-22	-3	5.5893E-4	5.5893E-4	-284.38	-74.688	1601.4	7.6581
13	0	-22	1	.0019683	.0019683	-284.38	44.675	1601.3	7.8842
13	0	-23	1	7.1055E-4	7.1055E-4	-300.15	44.675	1600.9	7.8666
13	0	-21	-3	4.5948E-4	4.5948E-4	-268.69	-74.688	1601.8	7.6653
13	0	-22	-2	.0022185	.0022185	-284.38	-59.678	1601.4	7.6067
13	0	-19	3	4.0505E-4	4.0505E-4	-237.56	74.688	1602.4	8.1803
13	0	-23	2	3.4919E-4	3.4919E-4	-300.15	59.678	1601.0	7.9352
13	0	-20	-2	9.3722E-4	9.3722E-4	-253.09	-59.678	1602.1	7.8634
13	0	12	1	.0014799	.0014799	213.13	44.675	1603.9	9.3254
13	0	10	1	9.5695E-4	9.5695E-4	183.01	44.675	1604.1	9.1863
13	0	9	1	6.2061E-5	6.2061E-5	168.01	44.675	1604.2	9.1506
13	0	-23	-5	.0013485	.0013485	-300.15	-104.74	1601.1	7.8733
13	0	-23	-4	9.0808E-4	9.0808E-4	-300.15	-89.708	1601.1	7.8229
13	0	-22	-5	4.5937E-4	4.5937E-4	-284.38	-104.74	1601.5	7.8758
13	0	3	4	.15542	.15542	78.508	89.708	1604.5	7.3390
13	0	4	3	.0018326	.0018326	93.378	74.688	1604.5	7.4323
13	0	-1	2	.0018187	.0018187	33.745	59.678	1604.3	7.5732

FPGA/ECAL RO/Trigger Banks

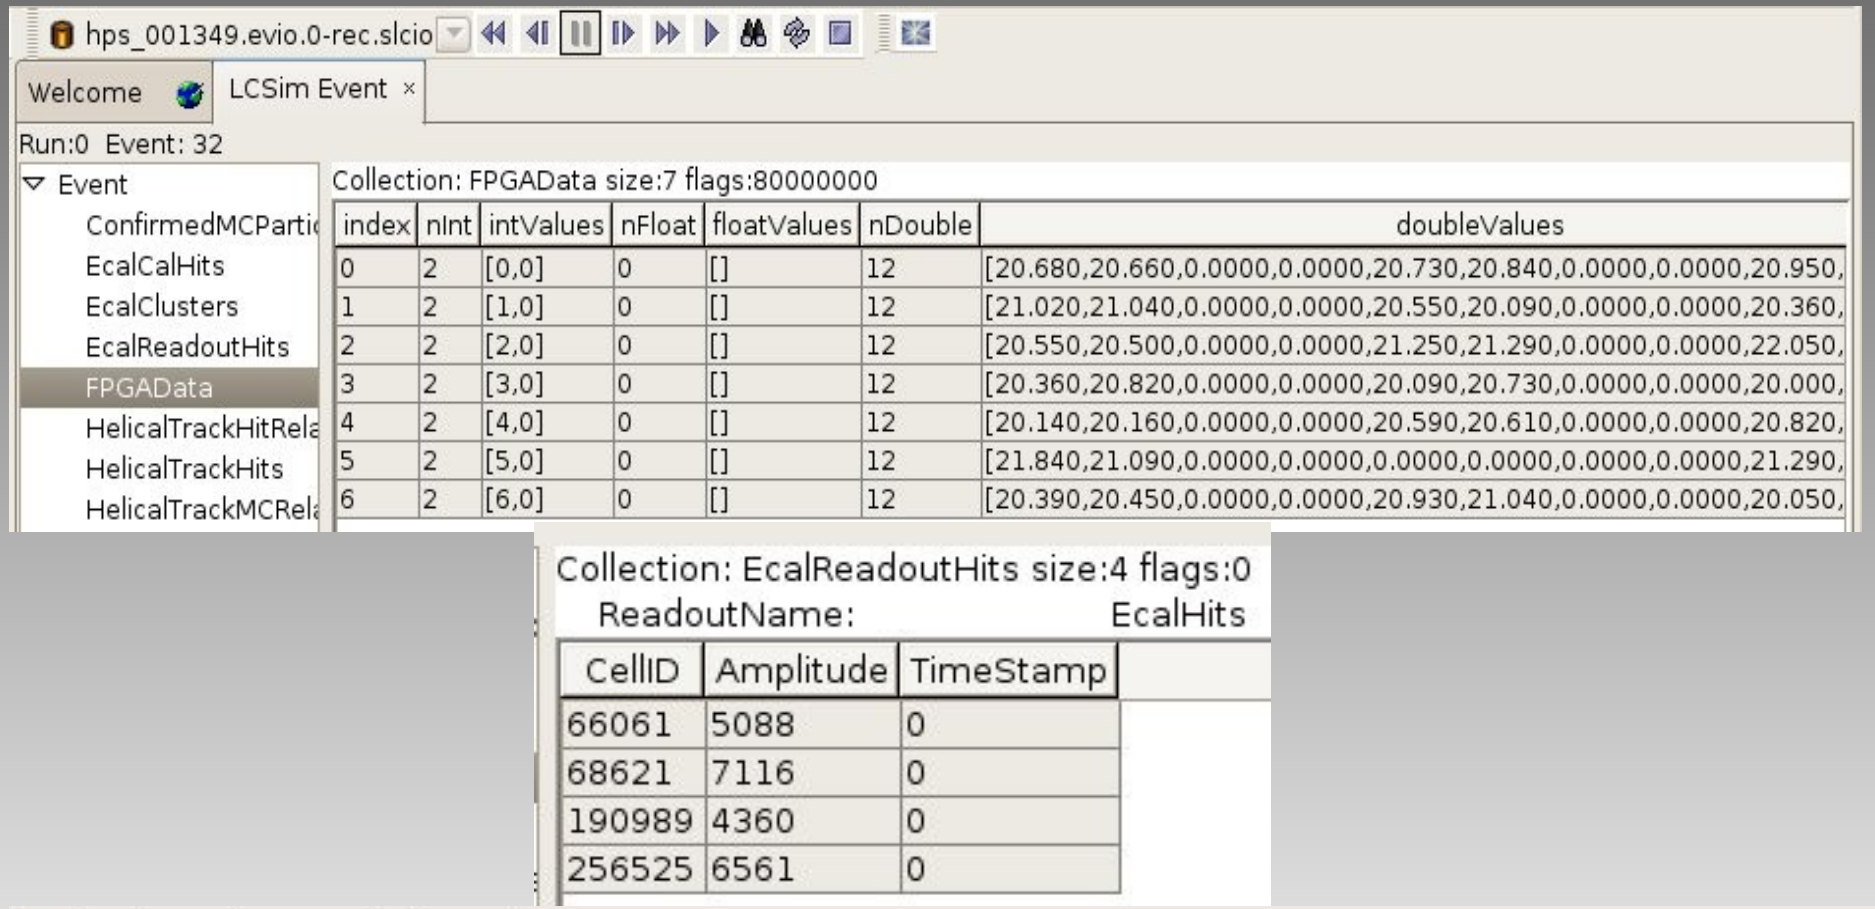

Collection: EcalReadoutHits size:4 flags:0

ReadoutName: EcalHits

CellID	Amplitude	TimeStamp
66061	5088	0
68621	7116	0
190989	4360	0
256525	6561	0

Collection: TriggerBank size:1 flags:80000000

index	nInt	intValues	nFloat	floatValues	nDouble	doubleV
0	8	[-2013265889,-1744830076,10824731,32768,32768,0,0,1337323715]	0	[]	0	[]

Simulation

-rw-r--r--+ 1 phansson at **1279256** Aug 22 17:32
/u/ey/homer/hps2/hps_data/simulation/testrun/egs_5.5gev_0.016x0_500mb_90na_HPS-TestRun
v2/egs_5.5gev_0.016x0_500mb_90na_08_HPS-TestRun-v2_filtered250.**evio**

-rw-r--r--+ 1 phansson at **339743584** Aug 22 17:32
/u/ey/homer/hps2/hps_data/simulation/testrun/egs_5.5gev_0.016x0_500mb_90na_HPS-TestRun-
v2/egs_5.5gev_0.016x0_500mb_90na_08_HPS-TestRun-v2_filtered250.**slcio**

What to expect? Most events in the above file are empty.

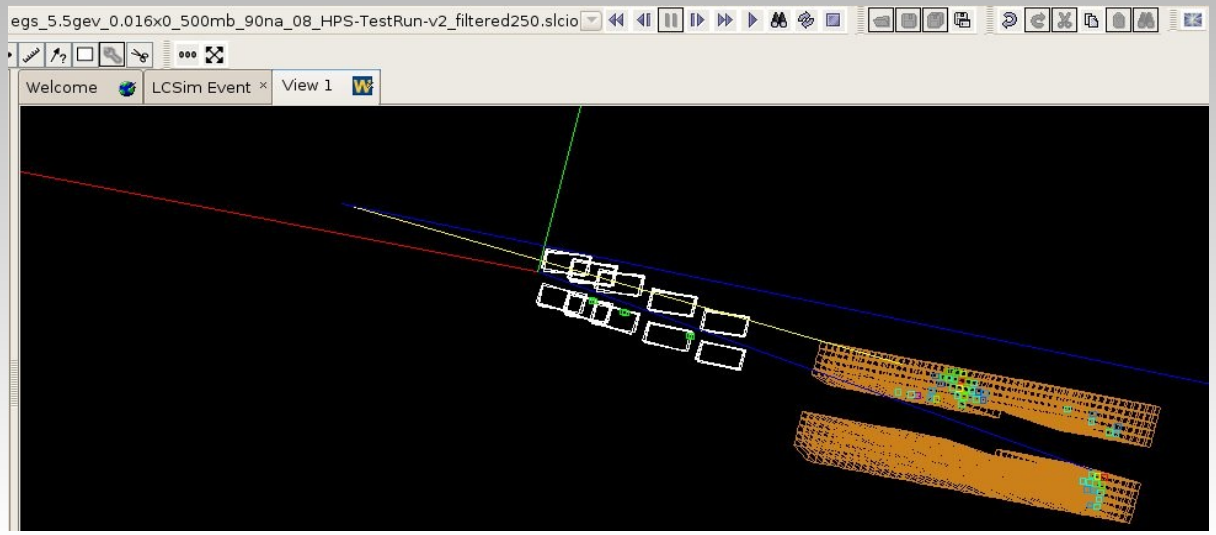

SIM SLCIO Contents

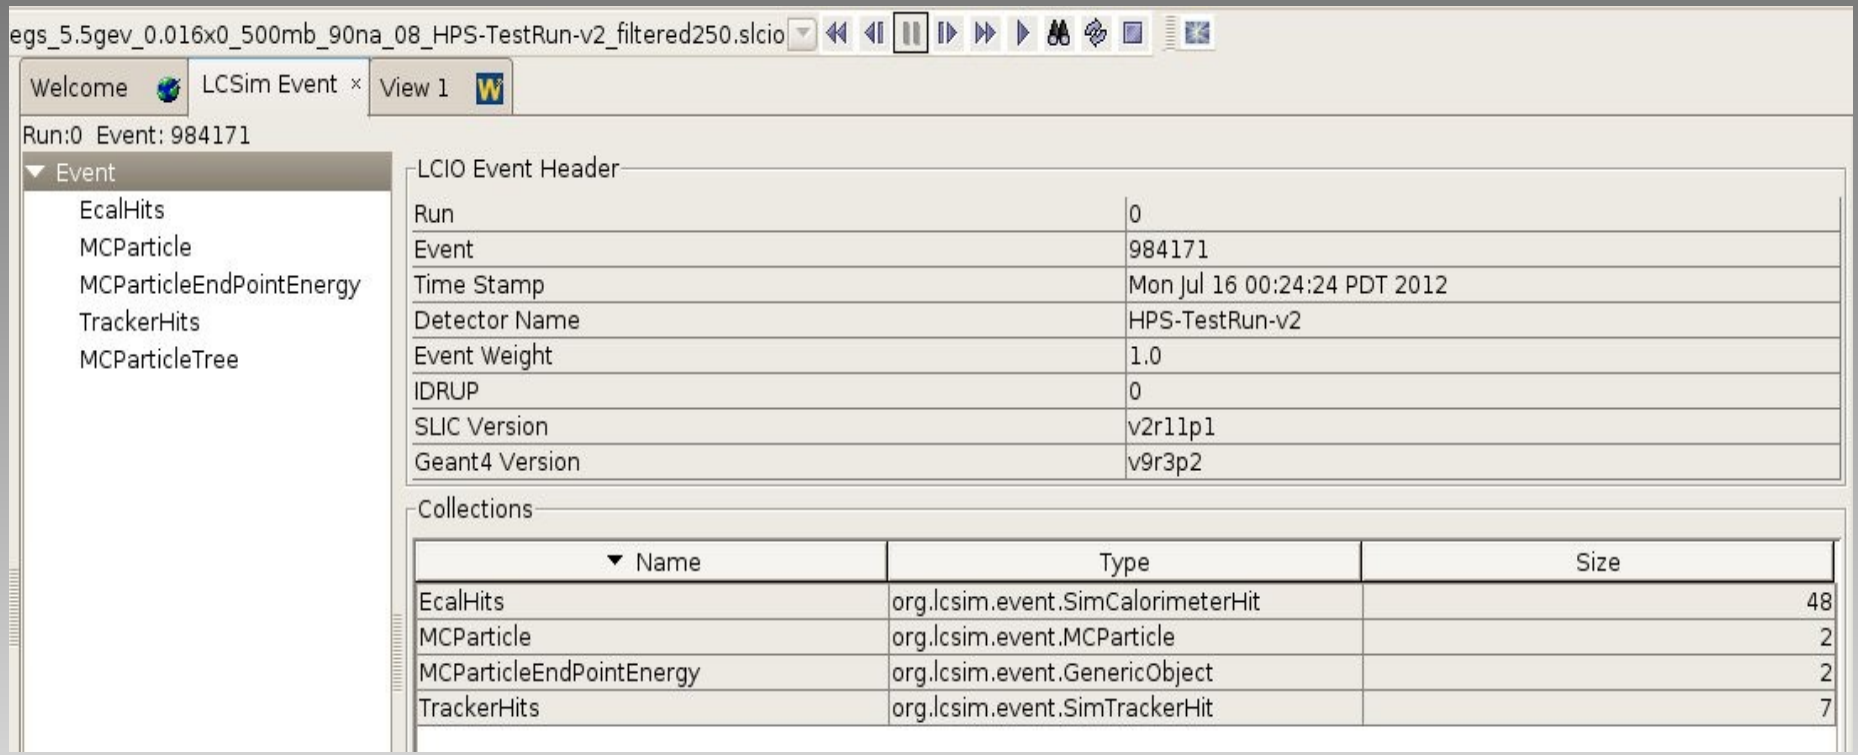

Questions

- Does slcio reconstruction need to be done at JLAB?
  - Decided by resource availability/politics?
- Is the the L3 filter data rate reduction reasonable?
- Use the GRID as for SiD?
  - Many negotiations would be involved.
- Resources?
  - What computing site contributions can one count on?
- Data storage needed?
- How much simulation will be needed for the the full run?
  - Is the 10% of real event count justifiable?
  - Will special large signal or special background samples be needed?

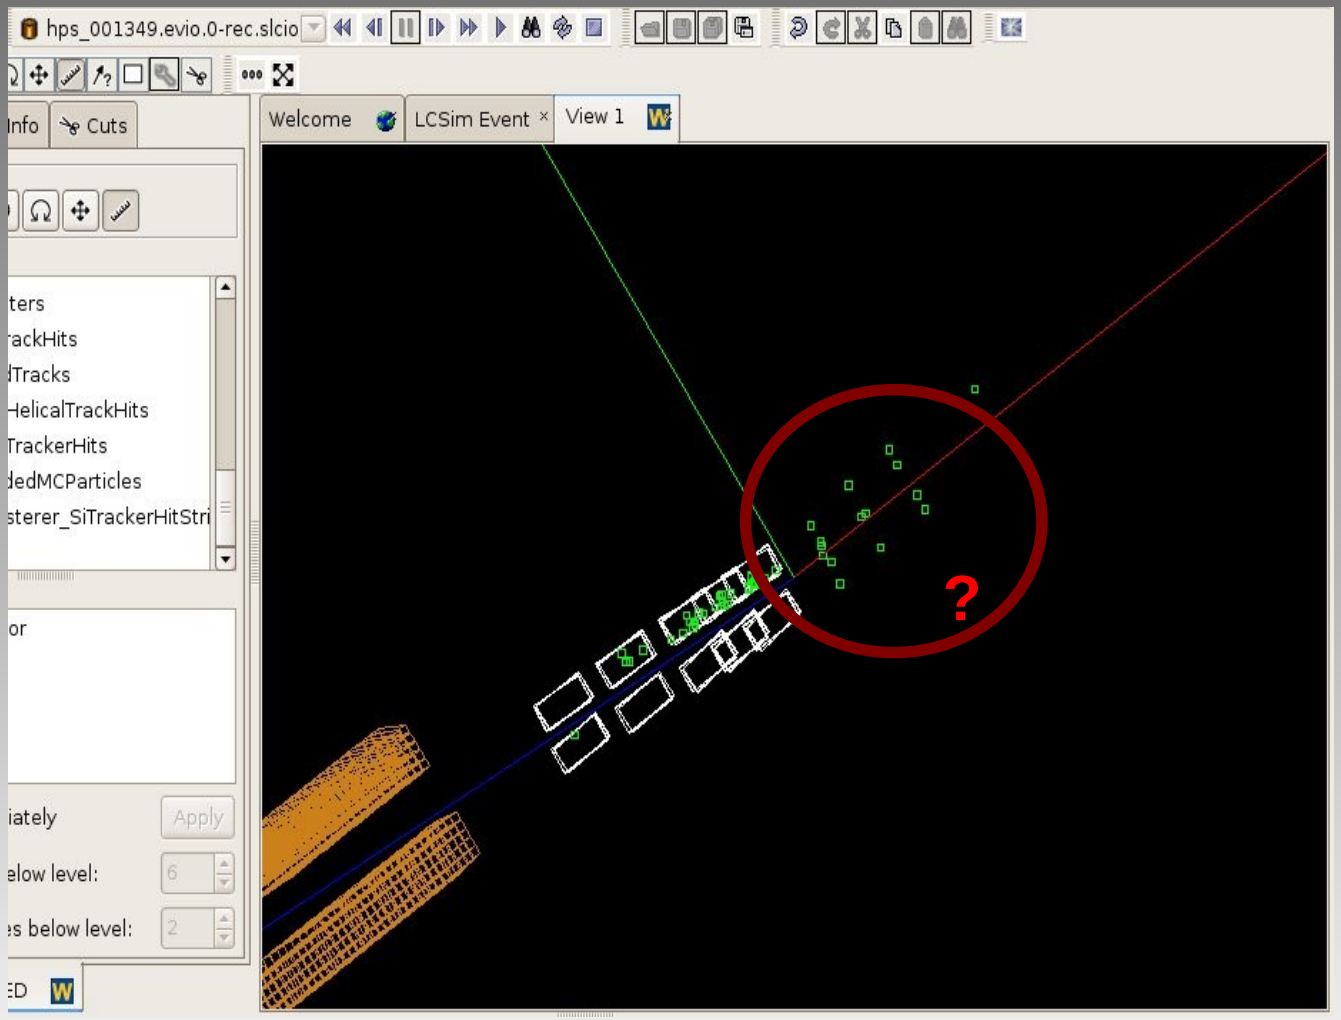
hps_001349.evio.0-rec.slcio
Info
Cuts
Welcome
LCSim Event
View 1
ters
rackHits
lTracks
HelicalTrackHits
TrackerHits
dedMCParticles
sterer_SiTrackerHitStri
or
iately
Apply
elow level:
6
s below level:
2
?