Offline Monitoring Report

February 1, 2015

Kei Moriya

Launch from 2015-01-23 (ver09)

- Ran over all (~7300) data files from last fall run
- Output:
  - detector plugins
  - REST
  - EventStore files
- Goals:
  - Test computing capabilities for processing ~7300 files
  - Provide REST for all files
  - Search for errors

Processing

- Processing very quick (finished most jobs by Sunday night) on new CentOS65 machines
- Request 6 cores, use 6 threads/job
- For each node, 32GB of RAM, 1TB of HDD, 42 threads
- Copying evio files from tape to node to cut down on accesses to /lustre disk

Errors

- Half of all files are completely error-free
- Most remaining are EVIO parsing errors

Output

- Detector plugin webpages being updated
- skim, REST files are in subdirs skims, REST of

  /volatile/halld/RunPeriod-2014-10/offline_monitoring/ver09/

- Disk usage:
  - plugin ROOT files : 56GB
  - REST : 113GB
  - idxa : 705MB
- Ran EventStore for 2-track, 3-track, 4-track, 5-track, 2-track+π^0, 3-track+π^0, 4-track+π^0, 5-track+π^0,

Statistics of Jobs

- Mark created mysql database from job IDs and info within SciComp's database: contains info on when each job entered which stage

- submit → dependency → pending → staging in → active → staging out → complete

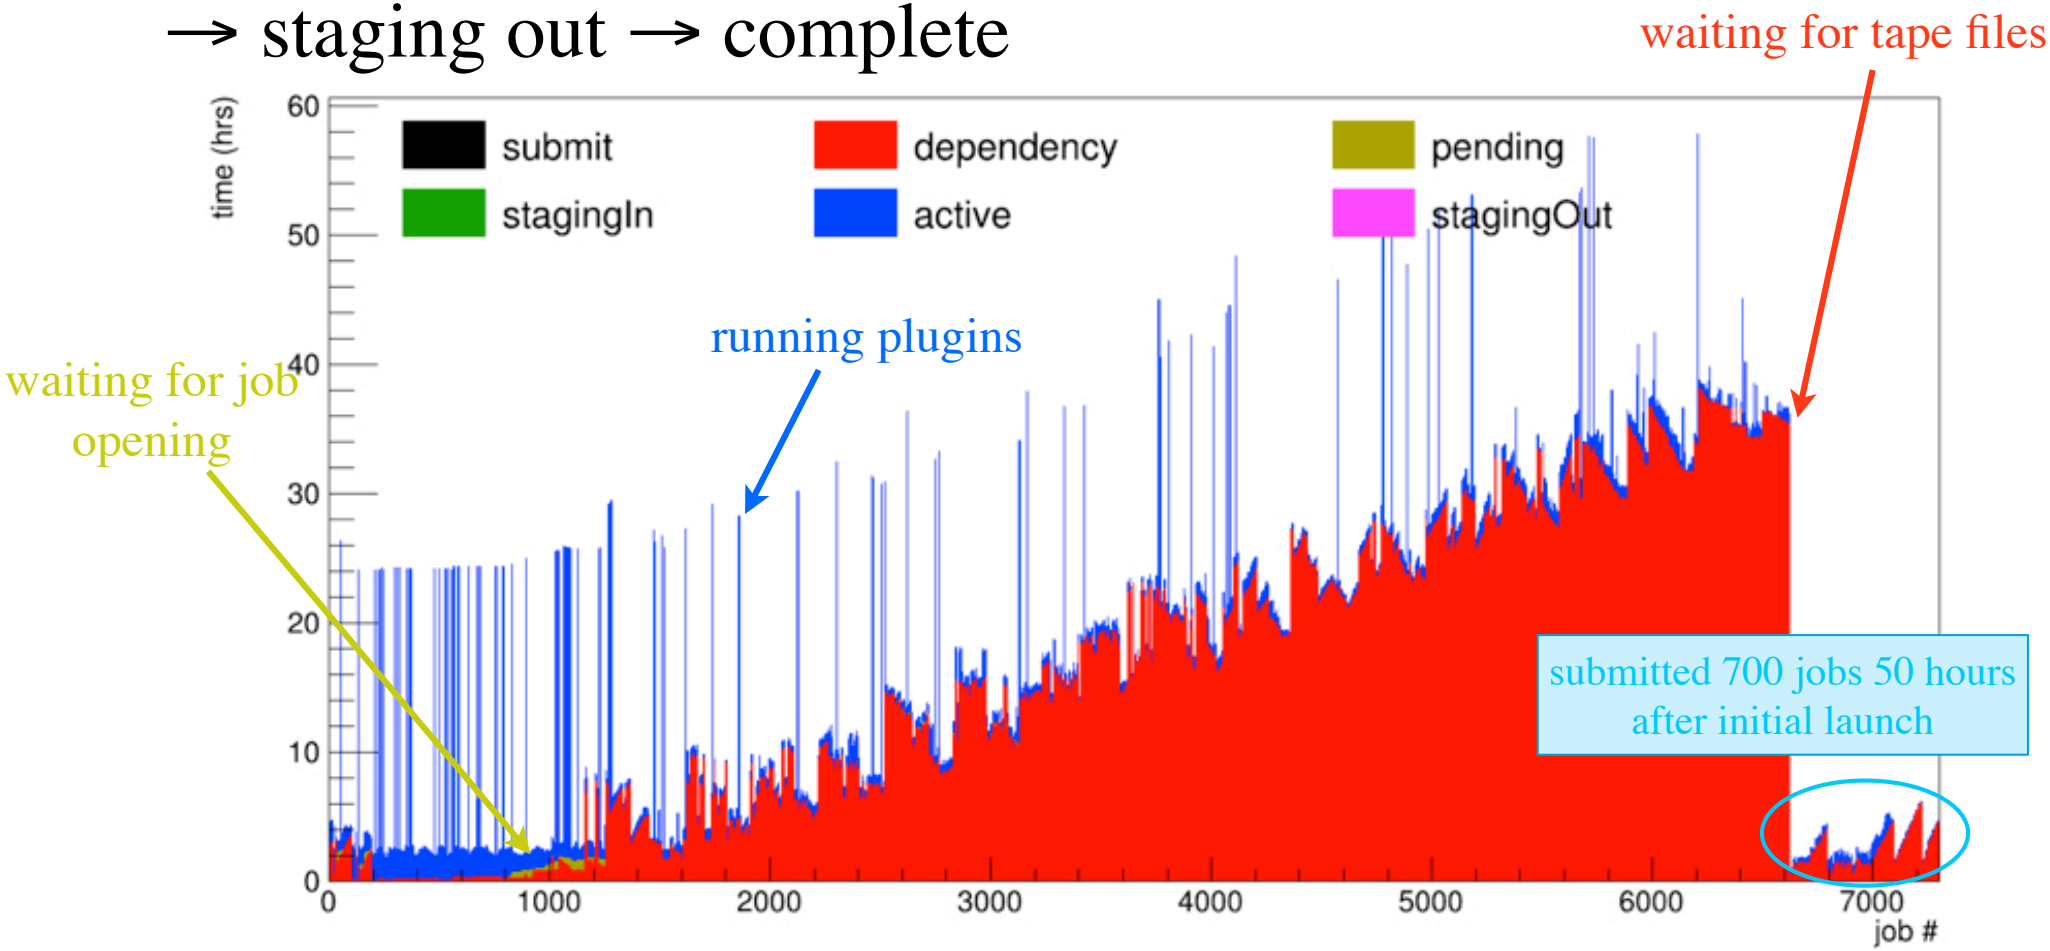

Screenshots

- From http://scicomp.jlab.org/scicomp/#/

Saturday 23:30

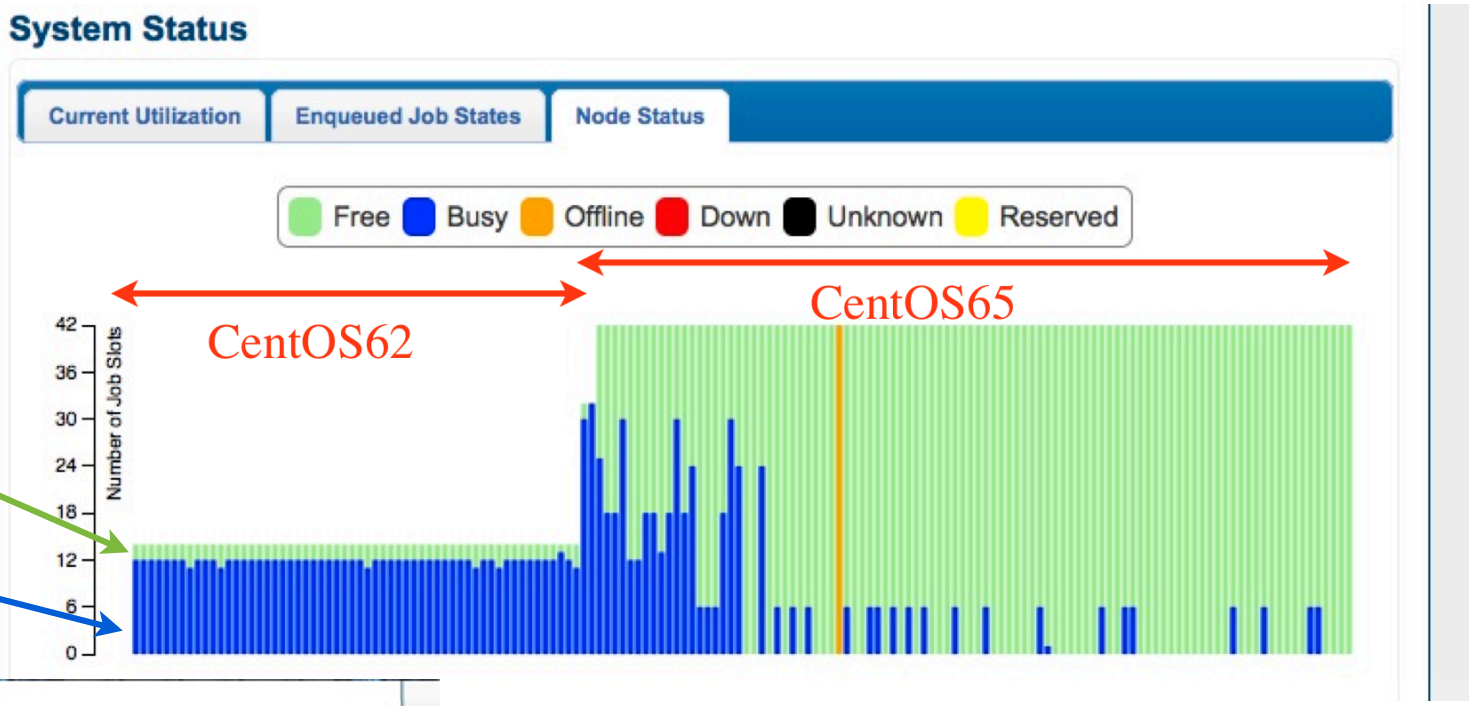

Batch Farm Jobs

Additional information is available through custom queries.

Recently Completed Jobs

User	Success	Failed	Cancelled	Timeout	Over Limit	Total
all	14840	216	5	250	32	15343

Job Queue

Search:

User	Account	Depend	Pending	Staging	Running	Staged	Total
crede	clas	733	67	0	41	0	841
fangguo	halic	0	0	0	19	0	19
gxproj1	halld	1830	0	0	84	0	1914
igorko	clas	47	0	0	0	0	47
primex	clas	0	0	0	3	0	3
rsholmes	halla	0	0	0	10	0	10
tianye	clas	0	2313	0	680	0	2993
whit	halic	0	0	0	3	0	3
All Users		2610	2380	0	840	0	5830

many jobs in “dependency” (waiting for tape files), only small number of running jobs

Time Since Start

- Track how many jobs finished since launch

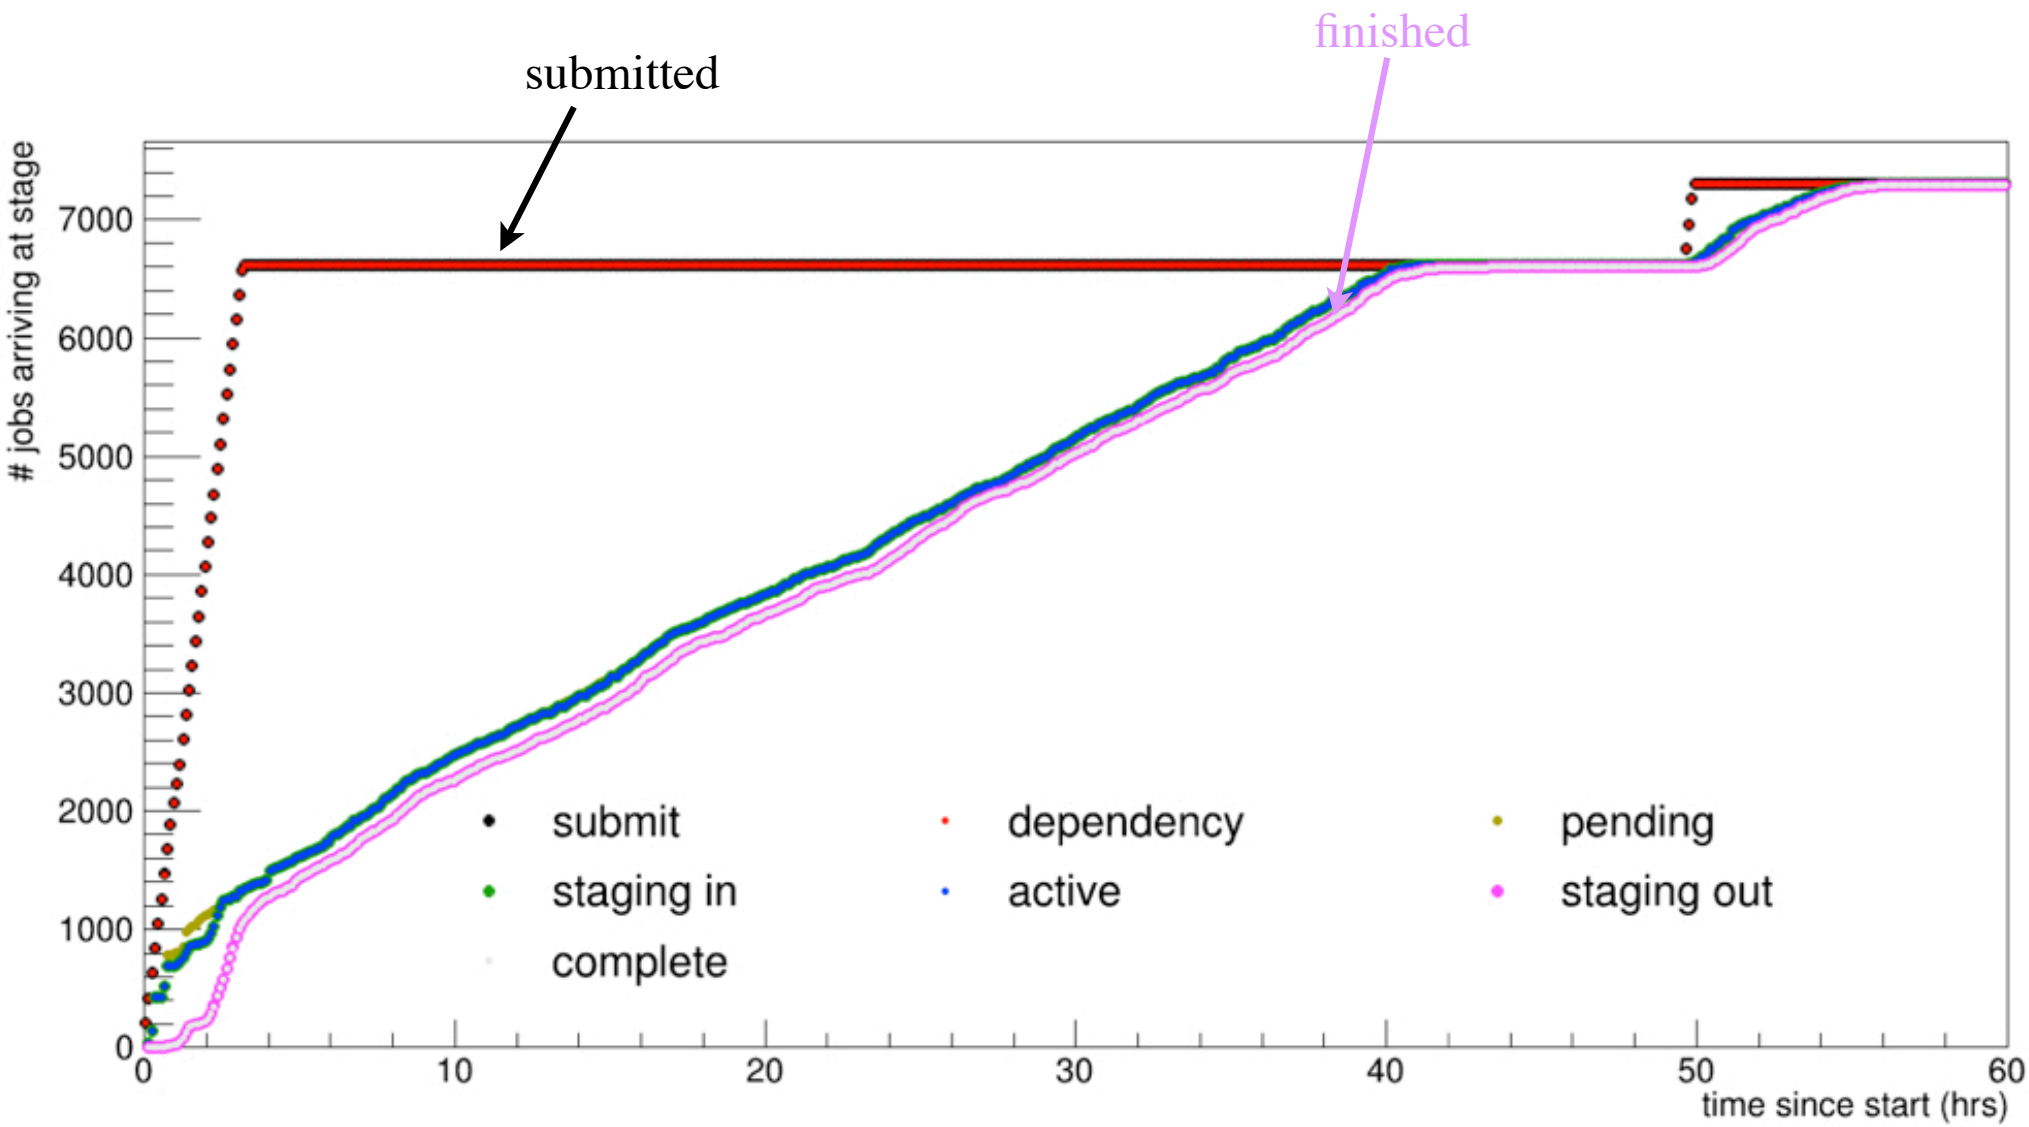

Active Jobs

- # of jobs running at a given time

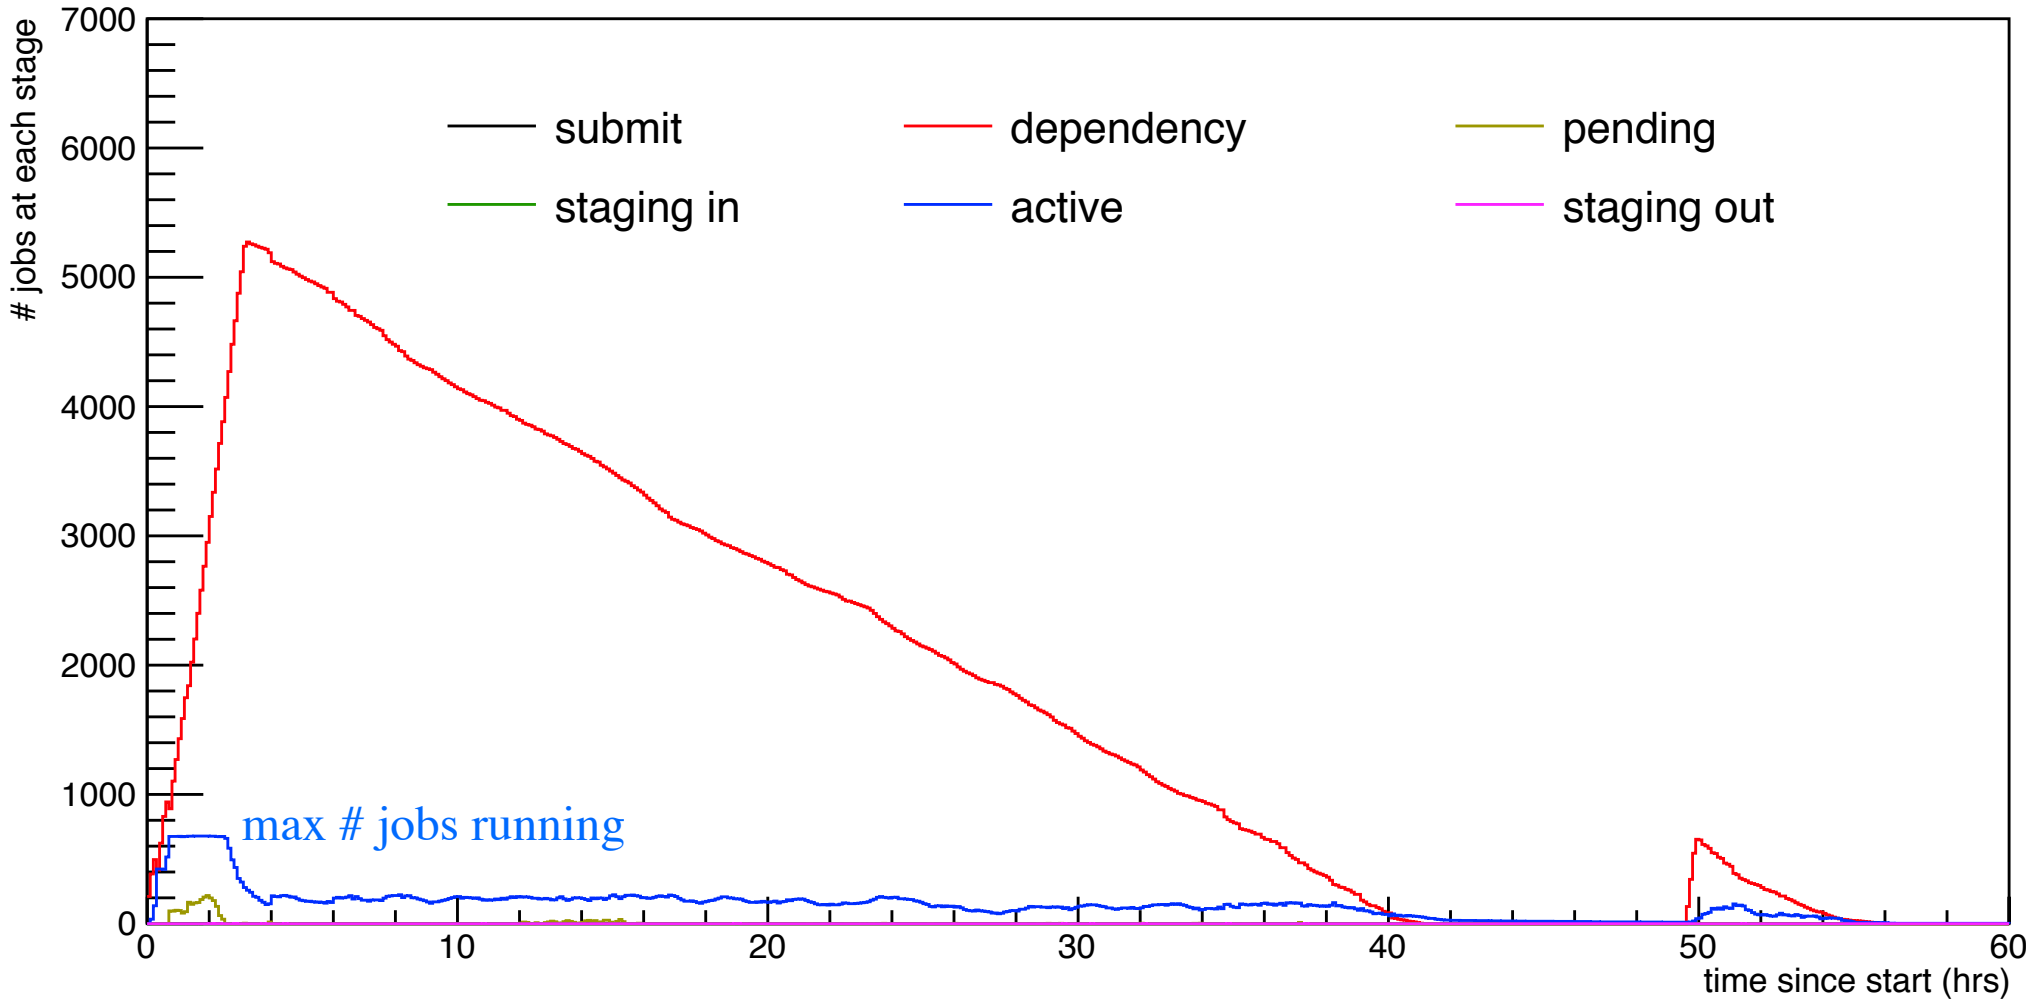

CPU Time vs Wall Time

- Requested 6 cores for each job
- Includes time to copy input file from cache

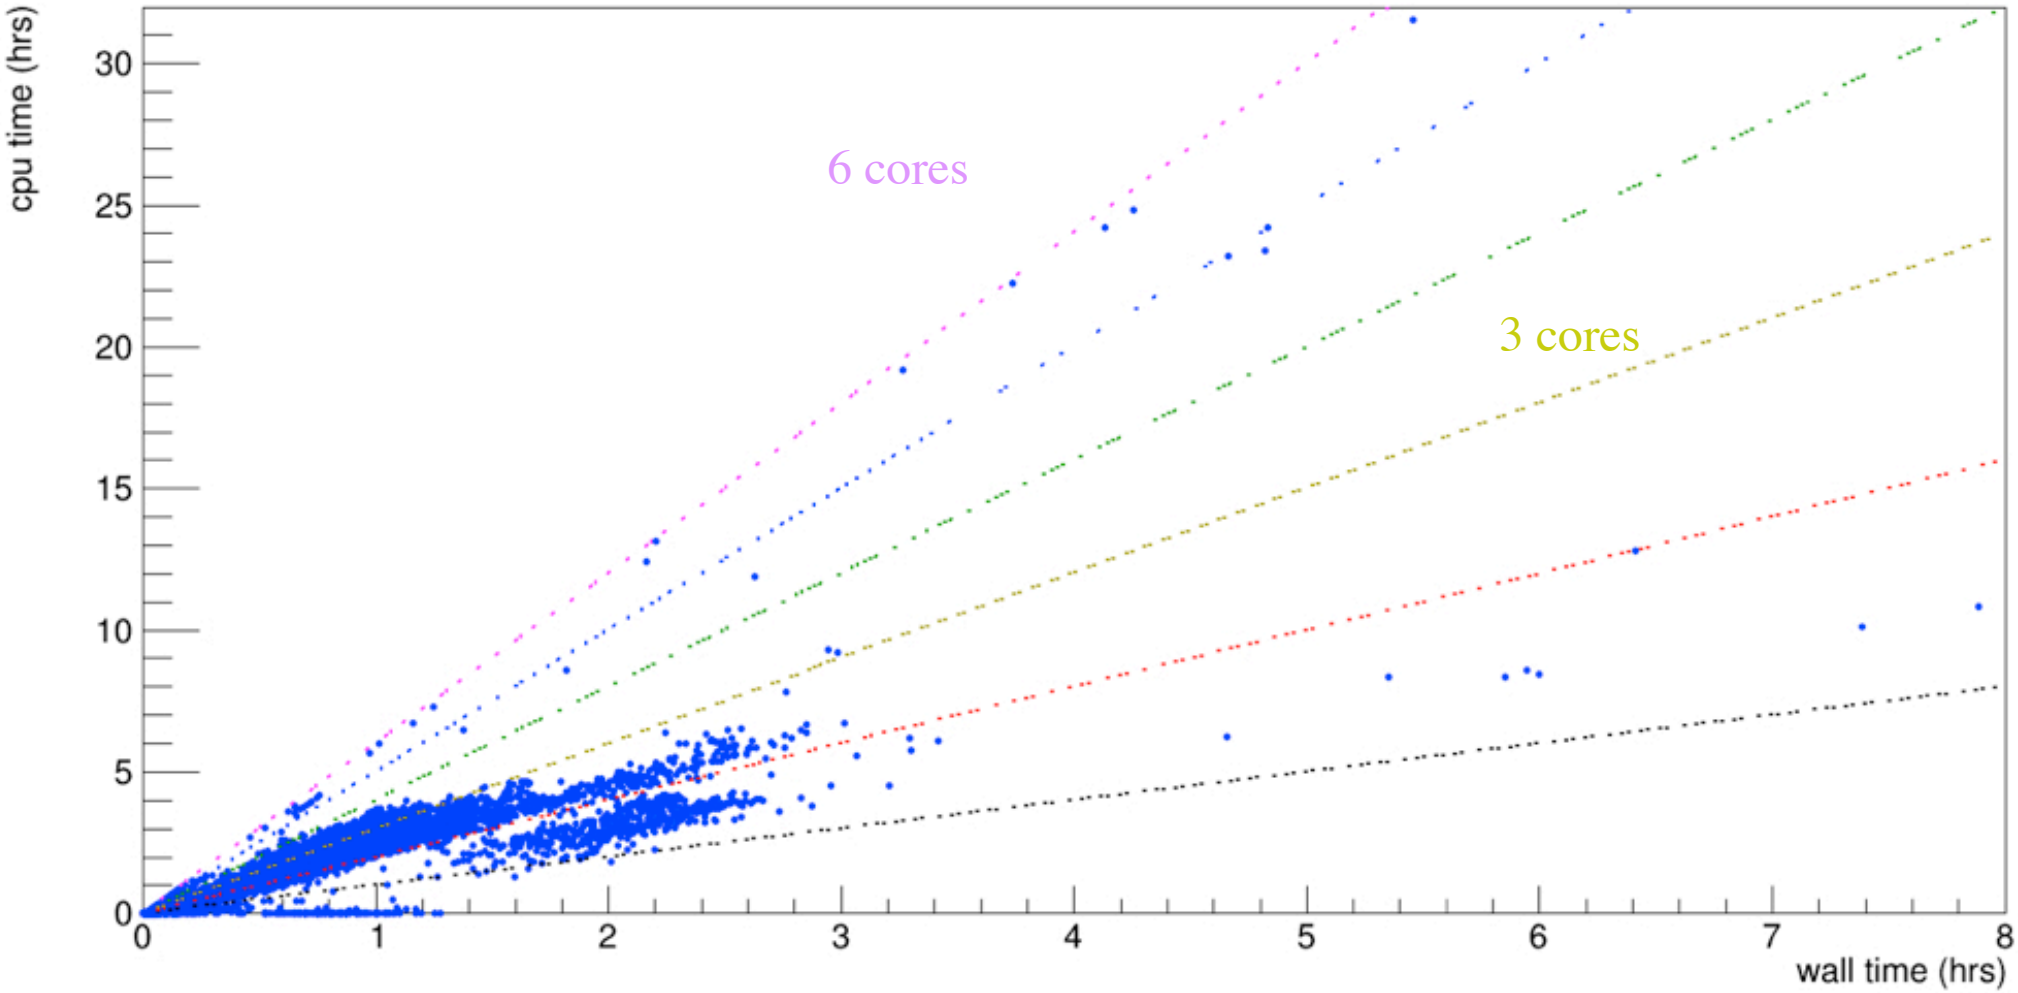

CPU Time vs Plugin Time

- Requested 6 cores for each job
- Use time only to run job

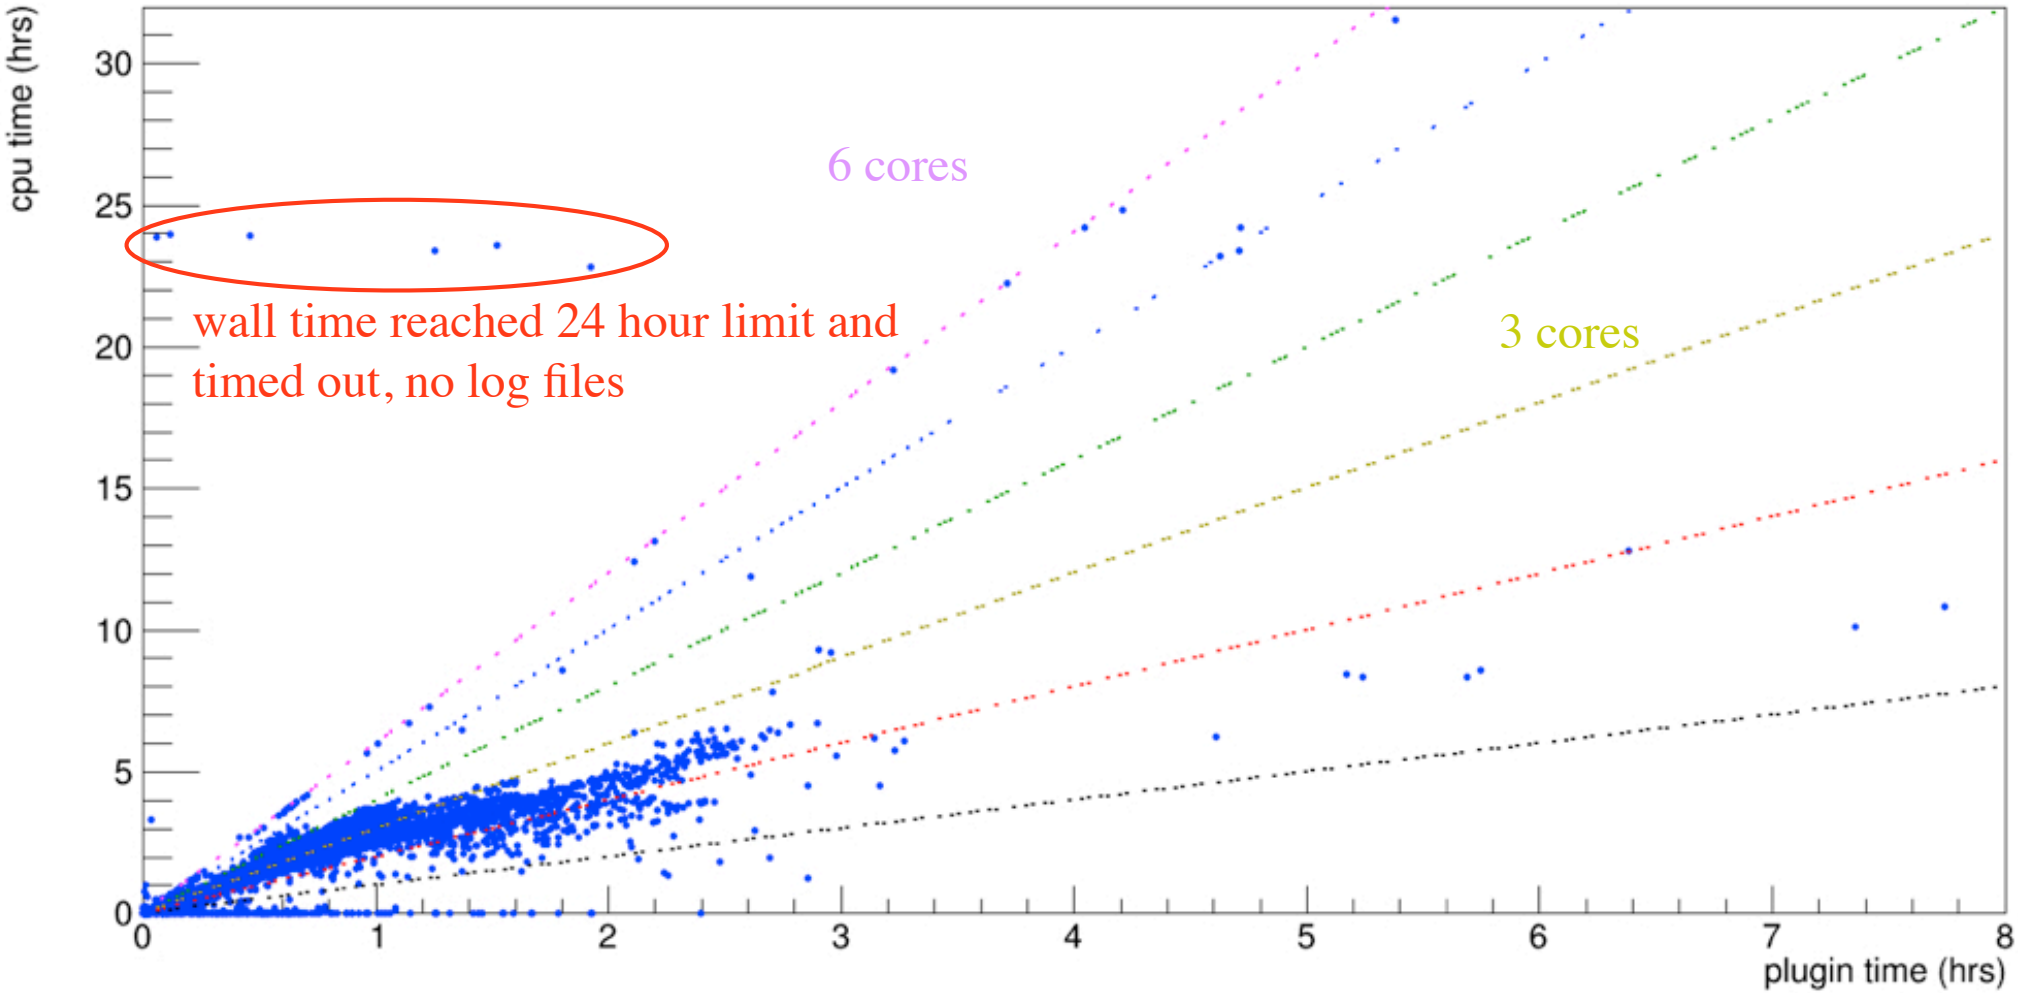

Process Time Per Event

- Divide plugin time by # of events
- Does not take into account time for file transfers, etc.

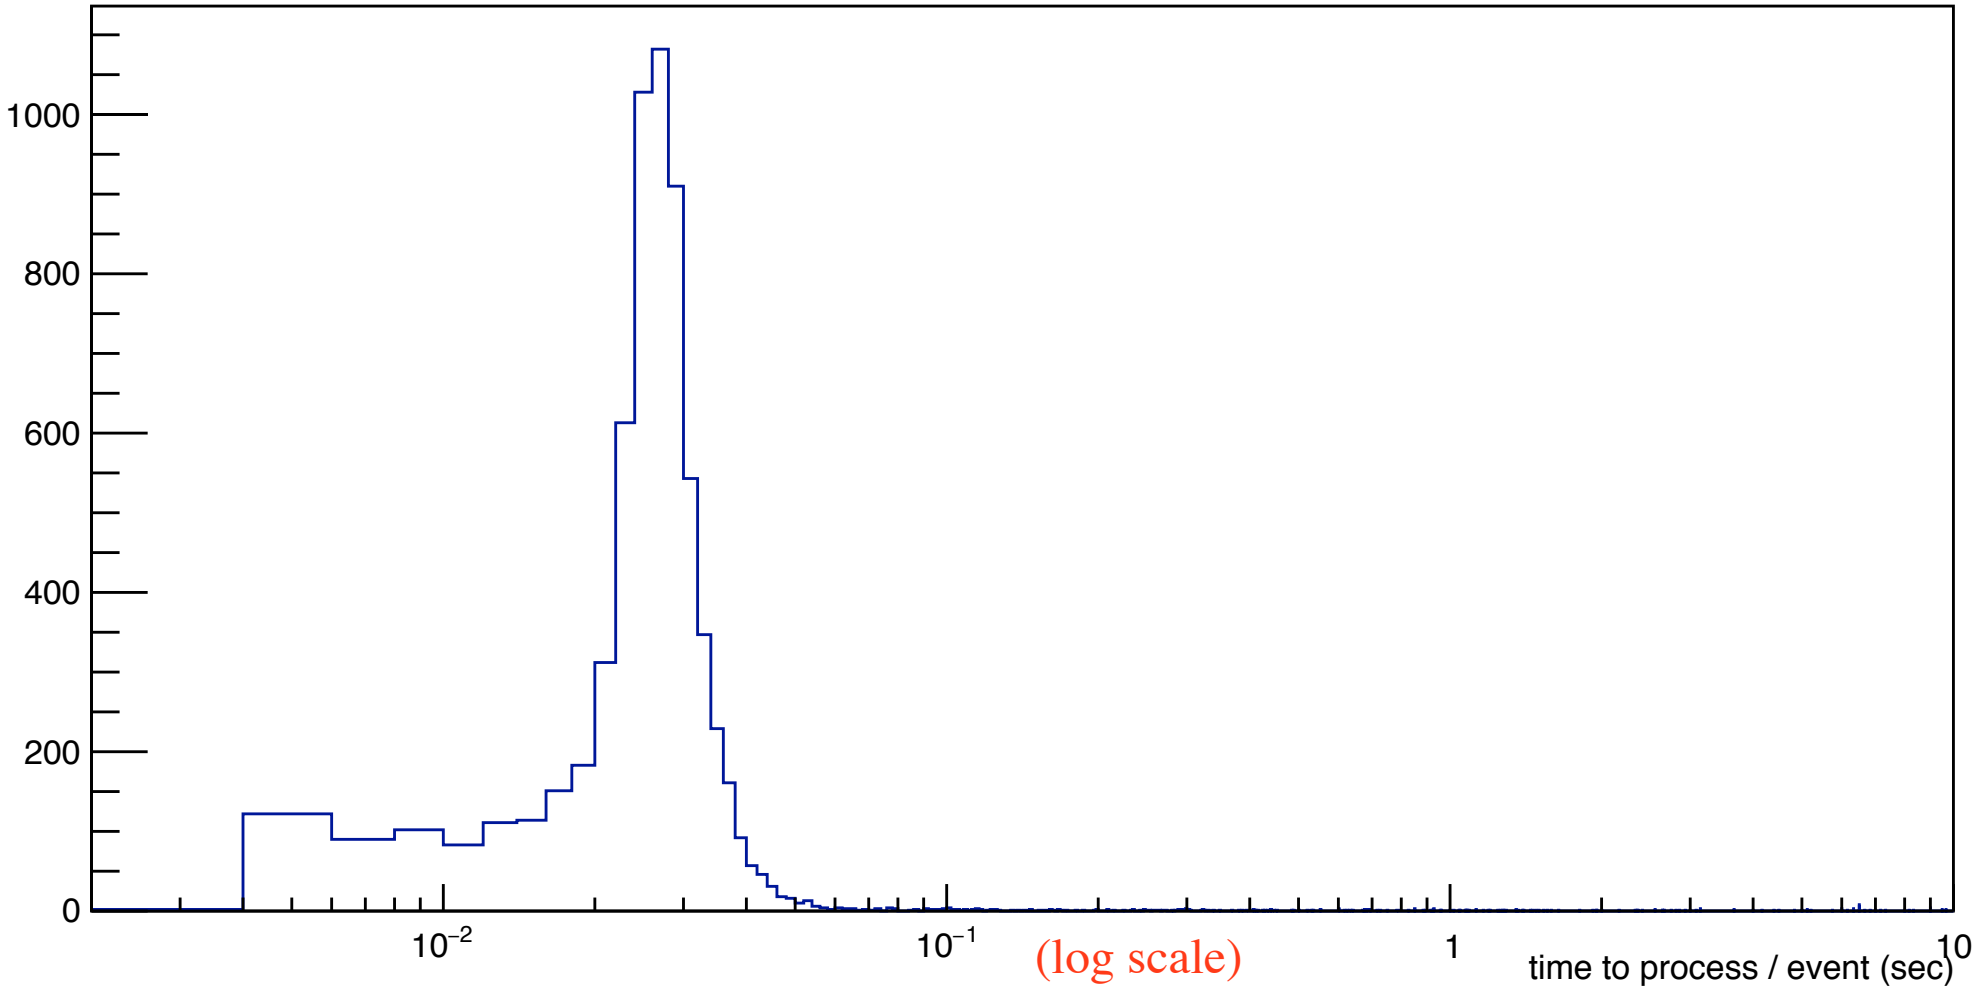

Conclusions

- Total size of EVIO files in initial launch: 107 TB, all copied within 40 hours → upper limit of 0.76 GB/s transfer rate (some files may have already been on tape)
- Automate this analysis to run on each launch

BACKUP
(Additional Information)

Memory Usage

- Requested 4.5 GB for each job
- 265 jobs killed due to lack of resources

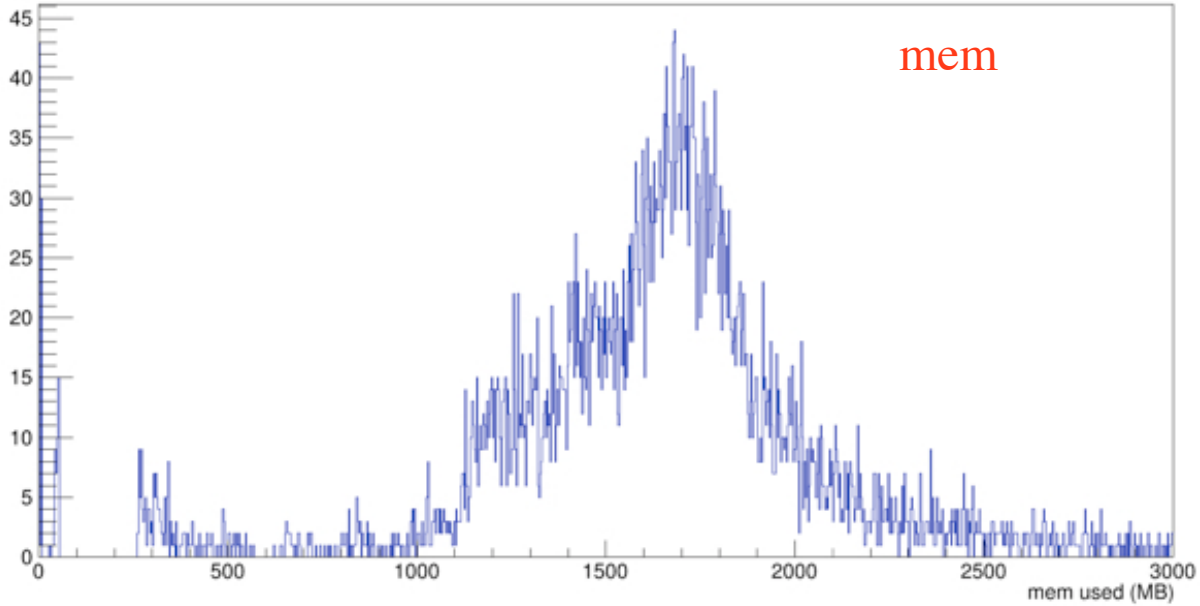

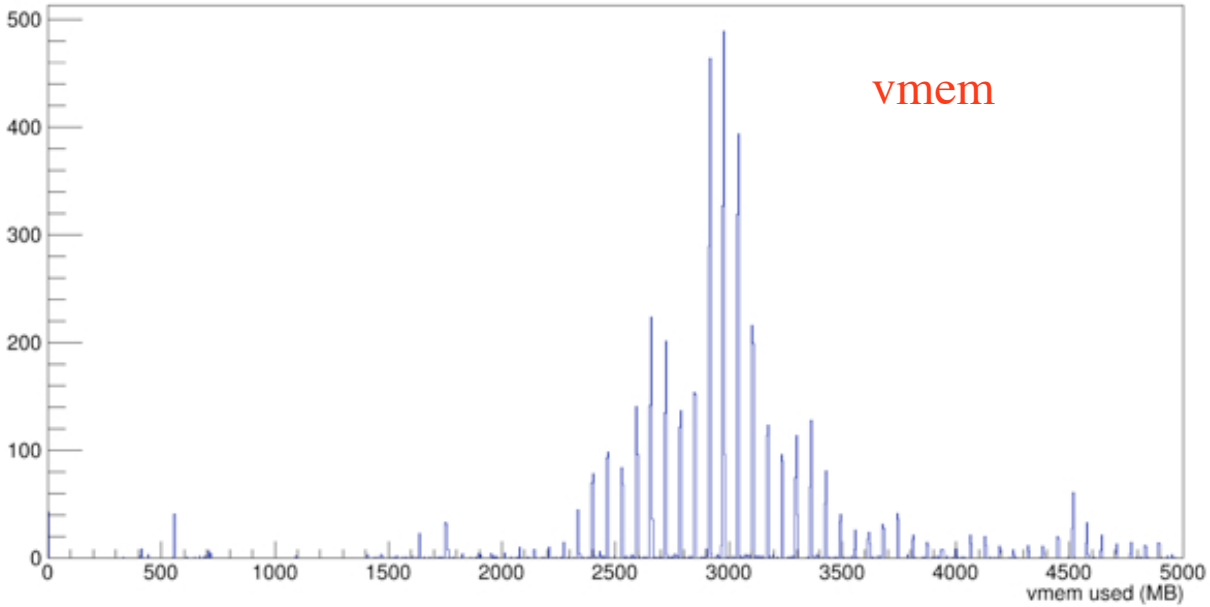

Lengths for Each Stage

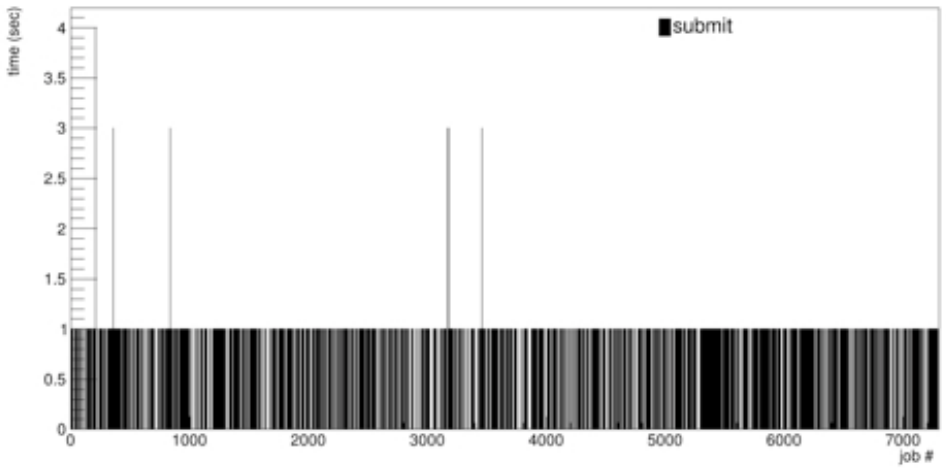

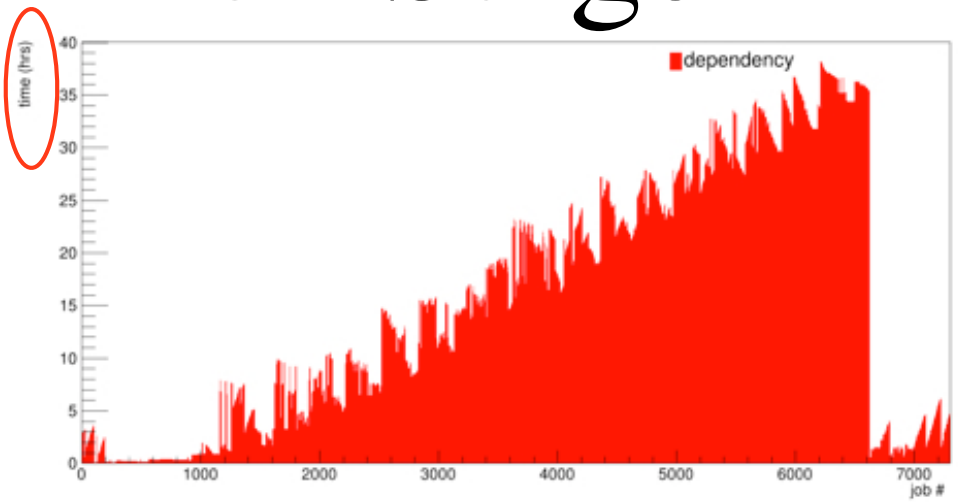

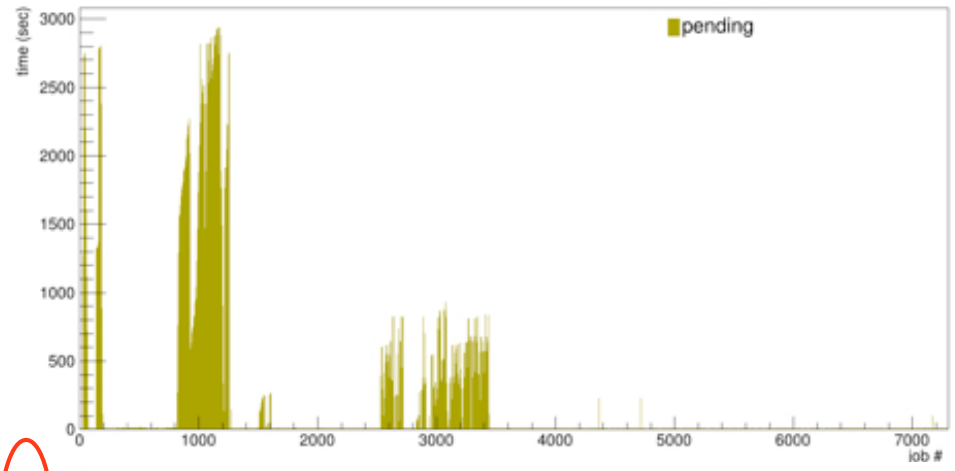

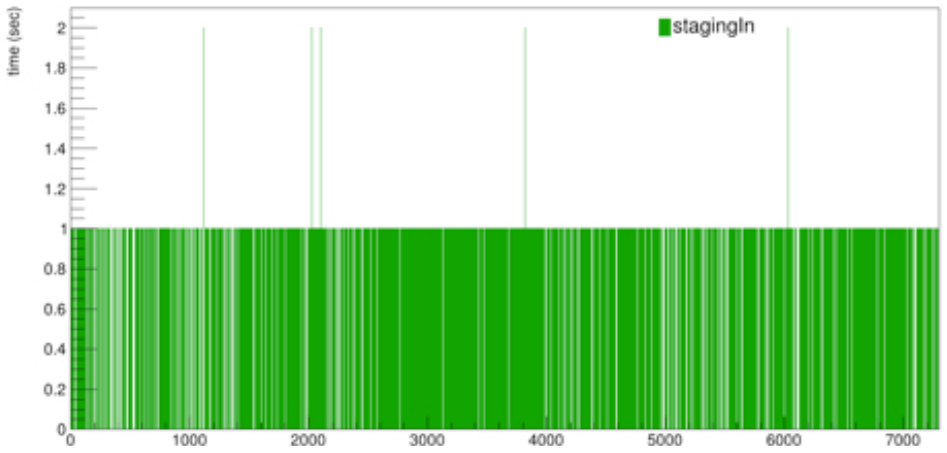

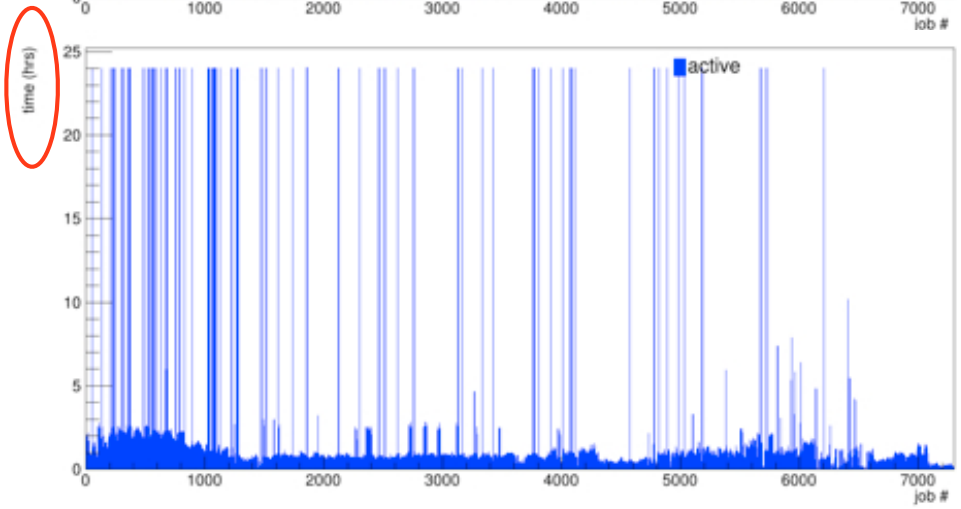

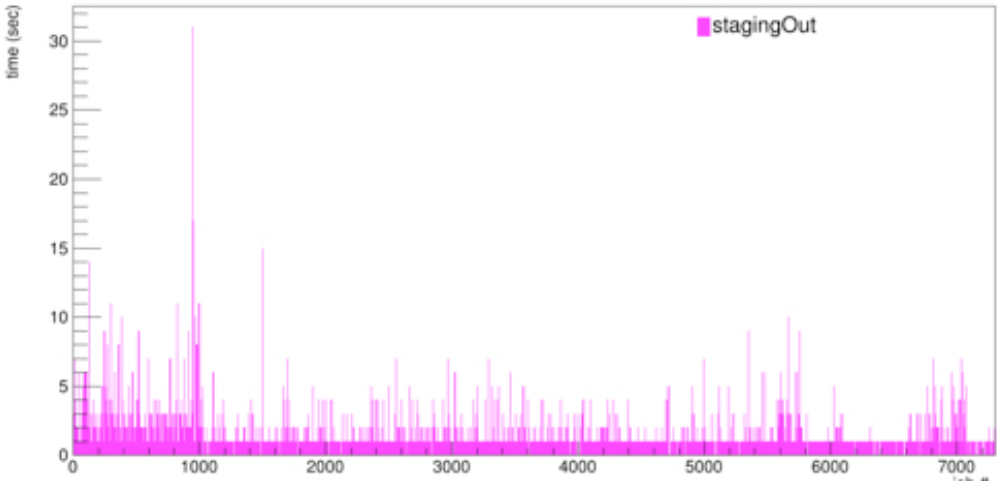

Active Time

- Time to run plugins - usually CPU time is several times this
- Some jobs do not get any CPU time
- Many jobs for run 2439 required more memory