

Offsite Processing of GlueX Raw Data in 2019

David Lawrence

October 3, 2019

Introduction

This document summarizes use of offsite computing facilities for processing GlueX experimental raw data during calendar year 2019. Specifically the National Energy Research Scientific Computing Center (NERSC) and the Pittsburgh Super-computing Center(PSC). This does not include resources used on the Open Science Grid (OSG) for simulation. Actual values are given where available and estimates provided for ongoing or anticipated usage during the remainder of the calendar year. Additionally, estimates on needs for calendar year 2020 are given.

1 Reconstruction Software

The GlueX reconstruction software is continually evolving with frequent changes leading to new version tags. This motivates us to use a system that can easily distribute newly compiled versions which can then be run on all production platforms, regardless of the host OS. This is done using a Docker container to provide cross-platform uniformity and CVMFS for binary distributions. Both NERSC and PSC support the use of containers. NERSC supports SHIFTER¹ while PSC supports Singularity². Both container types can be easily generated from Docker containers. The common Docker container used by GlueX for both of these is:

```
docker:markito3/gluex_docker_devel
```

Both NERSC and PSC support CVMFS. The CVMFS volume used is `/cvmfs/oasis.opensciencegrid.org/gluex`. This mirrors sections of the `/group/halld` disk on the JLab CUE. Thus, new builds at JLab are automatically distributed to offsite facilities via CVMSFS. The Calibration Contstants DataBase (CCDB)

¹<https://docs.nersc.gov/development/shifter/how-to-use>

²<https://sylabs.io/docs/>

is also distributed via CVMFS in the form of a SQLite file. Similarly for “resources” like the magnetic field maps.

It is worth noting that while simulation is outside of the scope of this document, most GlueX simulation jobs are run on the OSG using the same Docker container (via Singularity) and CVMFS file system.

2 GlueX Data Metrics

2.1 Data Volumes

In the Spring of 2018 GlueX produced a raw data volume of nearly 2PB. In the Fall of 2018 another 1.2PB was produced. Early 2019 saw the PrimEx experiment run using a subset of the GlueX detector producing only 0.4PB of raw data. These values are summarized in table 1. Also shown in the table are values from the GlueX computing model used to estimate computing resources based on beam time and running conditions. (See appendix for summaries of the model outputs for parts of 2019 and 2020 running).

For reconstruction jobs, one 20GB raw data file will result in 7GB of output files. For offsite processing, this means an equivalent of about 1/3 of the raw data volume transferred to the remote site for processing will be transferred back to JLab.

	2016	2017	2018	2019	2020
actual (raw data only)	0.624	0.914	3.107	0.400*	
model (raw data only)		0.863	3.172	1.56	6.06
actual (production)	0.55	1.256	1.206*	0.62*	
model (production)		0.607	3.084	1.94	4.34

Table 1: GlueX Data volumes by calendar year. All values are in petabytes(PB). Most years include two run periods. *Marked values indicate partial numbers that are current at the time of writing, but are expected to increase.

3 NERSC

Significant portions of the RunPeriod2018-01 and RunPeriod2018-08 reconstruction pass were done using NERSC. Table 2 shows the facilities used for reconstruction of each batch.

	RunPeriod2018-01 recon v02	RunPeriod2018-08 recon v02
batch01	NERSC	NERSC
batch02	NERSC	JLab SciComp
batch03	NERSC	NERSC
batch04	NERSC	PSC + JLab SciComp
batch05	NERSC	
batch06	NERSC	
batch07	JLab SciComp	

Table 2: Locations where reconstruction was performed for 2018 GlueX data. For each run period, the data was broken up into batches for processing that are roughly equal in size.

3.1 Allocation Award

For allocation year 2019 we were awarded 35M “NERSC hours”. This is equivalent to about 42k GlueX jobs on the Cori II regular queue. This was about 1/3 of our request. During the course of the year, CLAS12 transferred most of their allocation to us amounting to about 14.5M NERSC hours. A final transfer of 4M NERSC hours was awarded to us in late August from allocations recovered by NERSC from awards to other projects that had not used them. The total allocation amount for GlueX at NERSC for calendar year 2019 was 58.5M NERSC hours or 70.5k jobs on Cori II regular queue. By the end of the RunPeriod2018-08 recon ver02 campaign, 5M NERSC hours were remaining on the GlueX account. Some of this is anticipated may be used for RunPeriod2019-11 data depending on the state of calibrations.

3.2 Data Transfer

Data transfer from Jefferson Lab to NERSC is done using globus. Jefferson Lab has a 10Gbps connection to ESnet which limits the maximum rate data may be transferred. Figure 1 shows the data transfer for the first 4 batches of RunPeriod2018-01 data. The average rate for each of these was limited largely due to contention for the Lustre file system where files are staged from tape before being transferred to NERSC. Prior to starting the first batch of the RunPeriod2018-08 campaign a dedicated data transfer node (DTN) was installed and the software configured to stage directly to it rather than Lustre. Other network parameters were tuned and the throughput was improved substantially. Figure 2 shows the transfer rate for some test jobs just before the main batch01 campaign. This demonstrated the ability to sustain 10Gbps rate to NERSC for multiple hours.

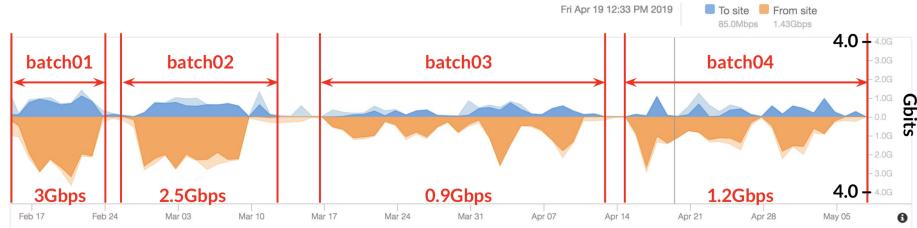


Figure 1: Data transfer rate to NERSC during processing of first 4 of 7 batches of RunPeriod2018-01 data. (~ 55% of data from run period). Transfer rates during this time period were not stable largely due to contention for the Lustre file system used to stage files read from tape before transferring offsite.

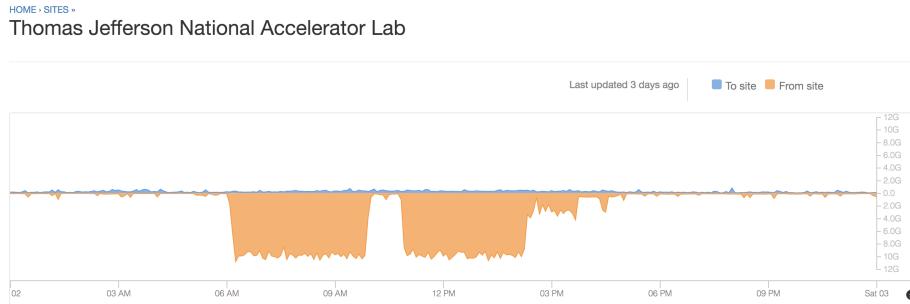


Figure 2: Data transfer rate to NERSC just prior to processing RunPeriod2018-08 data. The installation of a dedicated DTN (Data Transfer Node) which files are staged directly to along with some other optimizations significantly improved the sustained transfer rate to NERSC.

3.3 Job Rates

Figure 3 shows the processing rate vs. number of threads for a typical GlueX reconstruction job on NERSC Cori II. Note that on Knights Landing (KNL) processors, there are 4 hardware schedulers for each core. This results in a progressively slower rate increase as each additional hardware thread assigned to a core becomes active. The plot shows that some benefit is still observed all the way up to 268 threads.

Figures 4 and 5 shows the NERSC performance for the batch01 and batch03 data sets of RunPeriod2019-08 respectively. The batch01 data set was processed using the “regular” queue on Cori II while batch03 used the “low” priority queue. In the bottom plot of figure 5 there are a number of jobs that occurred on the 2nd and 8th days where the time to process took over 7.5 hours. These correlate to issues at NERSC with the file system that had been occurring since the upgrade to CL7 just before the RunPeriod2018-08 campaign began. The

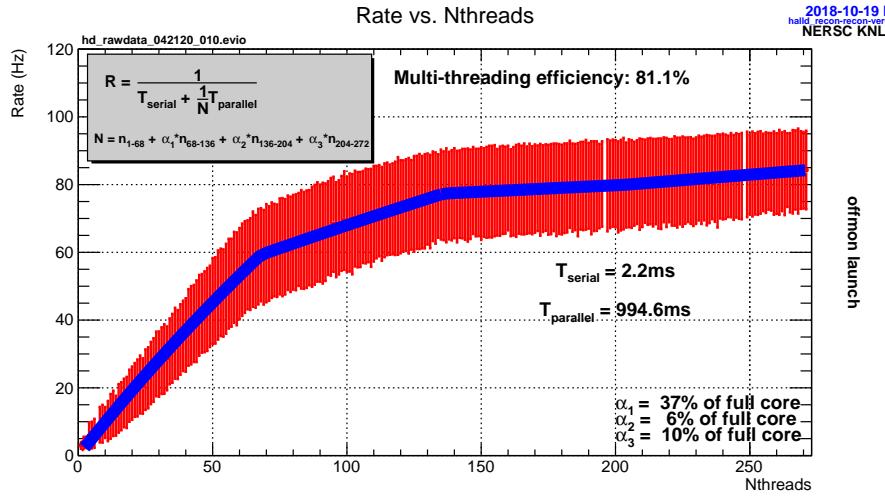


Figure 3: Thread scaling for NERSC job on Cori II (KNL).

system was patched twice with the last one being on 9/9 which seemed to have finally resolved the issue.

3.4 Haswell queue

We have not used Cori I (haswell) for production runs due to the greater demand and fewer resources on the system compared to Cori II (KNL). This comes at a cost of 2.4 times as much of our allocation if running on the Cori II regular queue. (Only 1.2 times if running on the low priority queue). A brief test was done in mid-August 2019 where 10 jobs were submitted to the Cori I queue. Figure 6 shows the results. The first jobs took nearly 20 hours to start. Each job would run for approximately 2.5 hours. Except for a brief period of time where 4 jobs were running, only 2 jobs would simultaneously run. This is believed to be primarily due to the NERSC scheduling policy of starting a maximum of 2 jobs based on priority with additional jobs started via “back filling”. The back filling of jobs is done after all jobs have been scheduled based on priority and there are holes in the schedule that lower priority jobs can fit into. GlueX jobs (single node for 8hrs) tend to be small compared to others on Cori II which allows them to be queued via the back filling mechanism easily. This may not be the case on Cori I and this test seems to support that. Based on this, we expect the effective throughput we could expect on Cori I would be quite small compared with what has been achieved with Cori II.

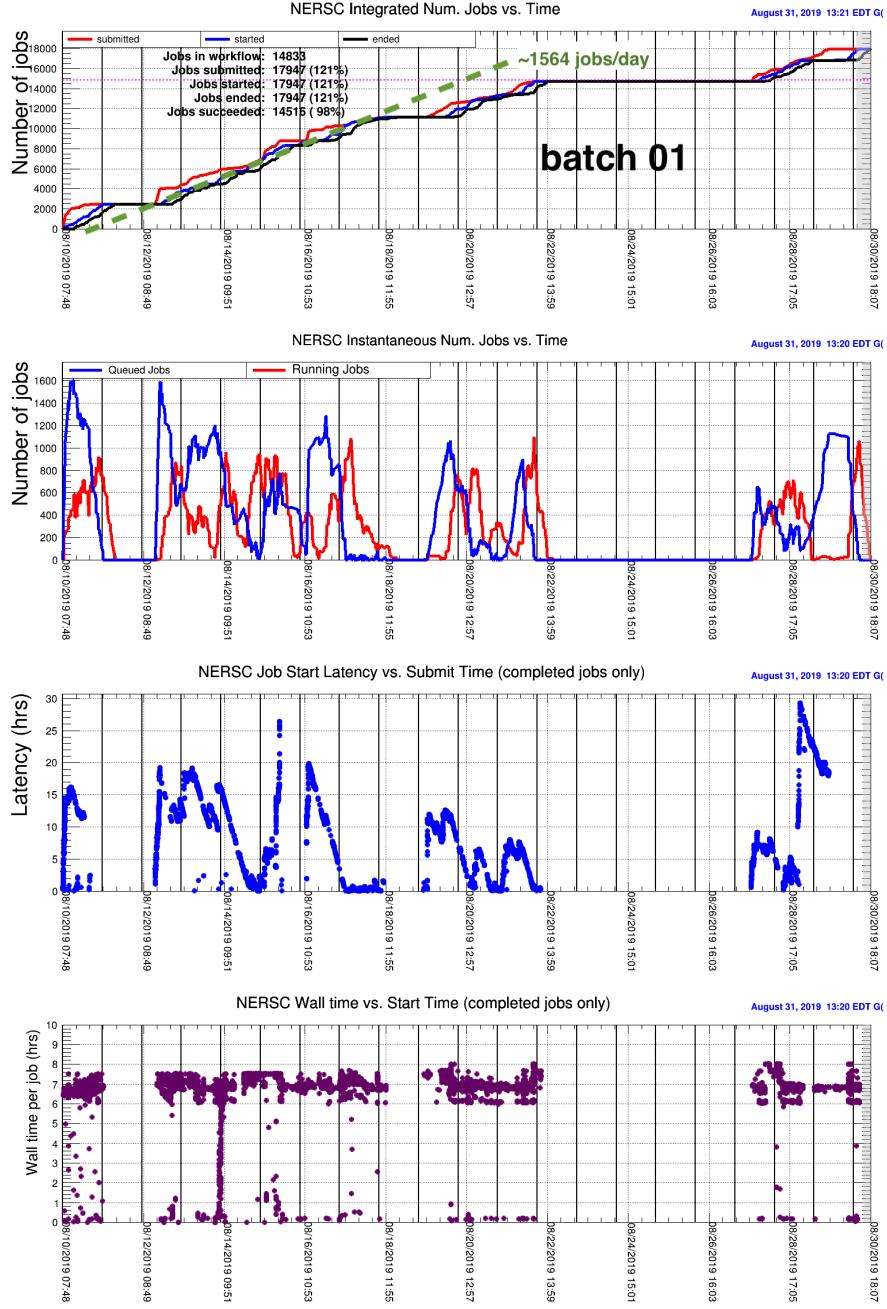


Figure 4: RunPeriod2018-08 data processing at NERSC for batch 01 using the “regular” priority queue. The top plot shows the integrated number of jobs vs. time while the second one shows the instantaneous jobs vs. time. The 3rd plot shows the latency between job submission and job start times while the 4th shows wall clock time for each job.

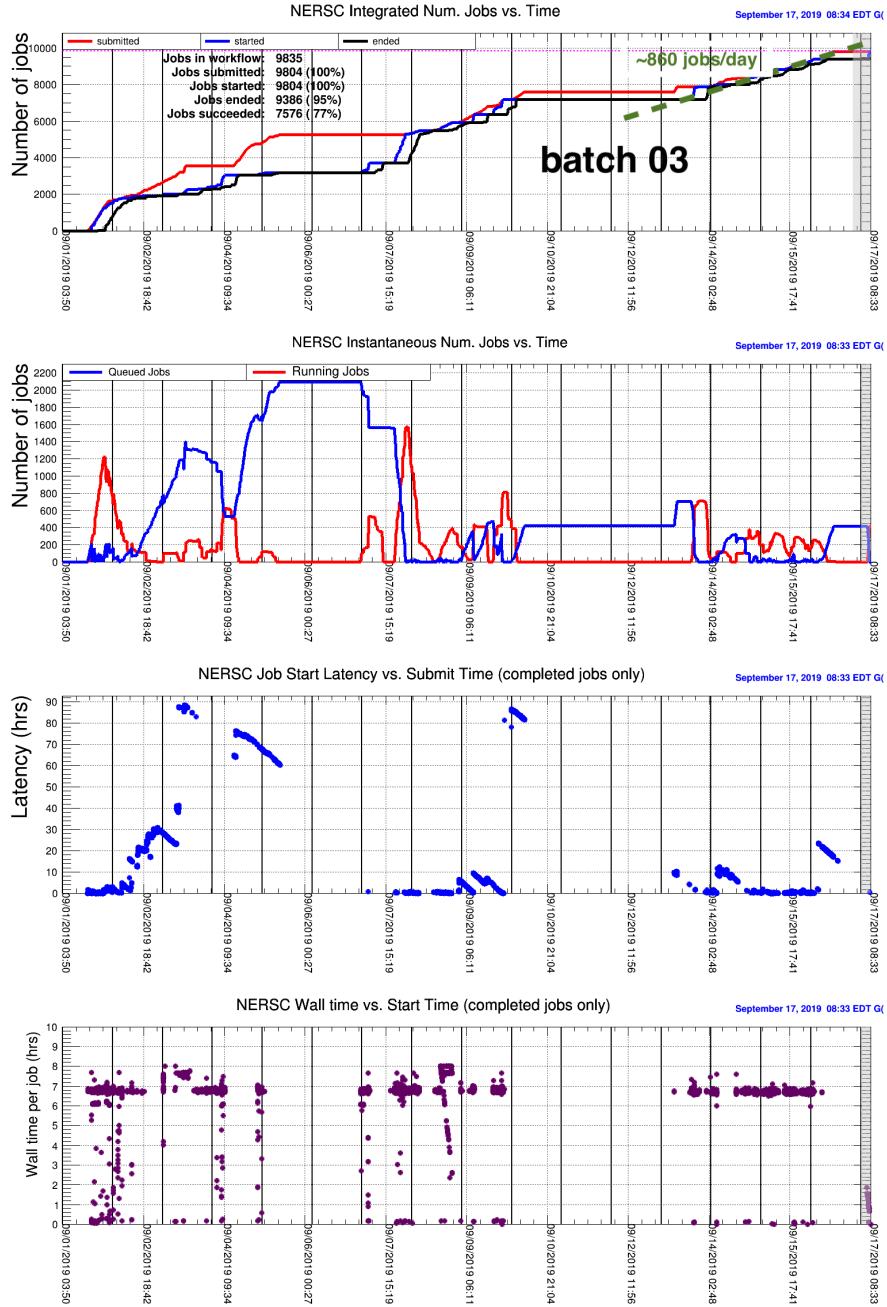


Figure 5: RunPeriod2018-08 data processing at NERSC for batch 03 using the “low” priority queue. The top plot shows the integrated number of jobs vs. time while the second one shows the instantaneous jobs vs. time. The 3rd plot shows the latency between job submission and job start times while the 4th shows wall clock time for each job.

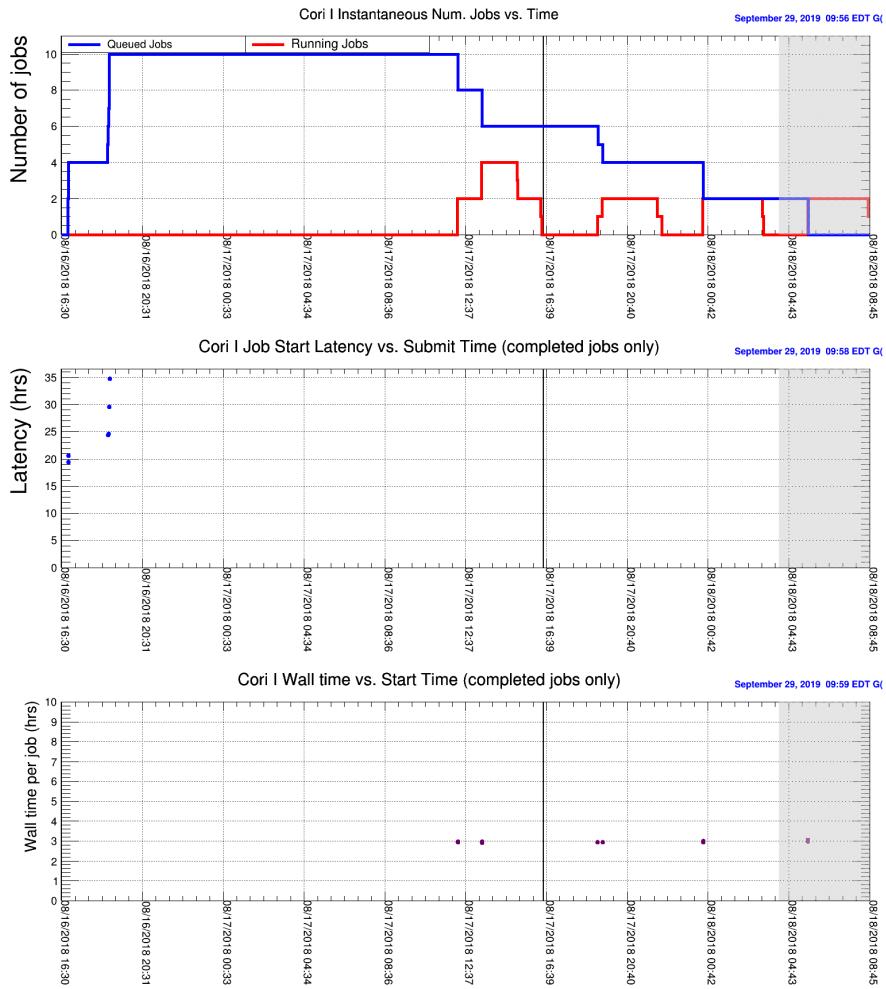


Figure 6: Jobs from short 10 job test using Cori I (haswell). Top: number of jobs vs time. Bottom: Latency between submit time to Cori I queue and start time of job.

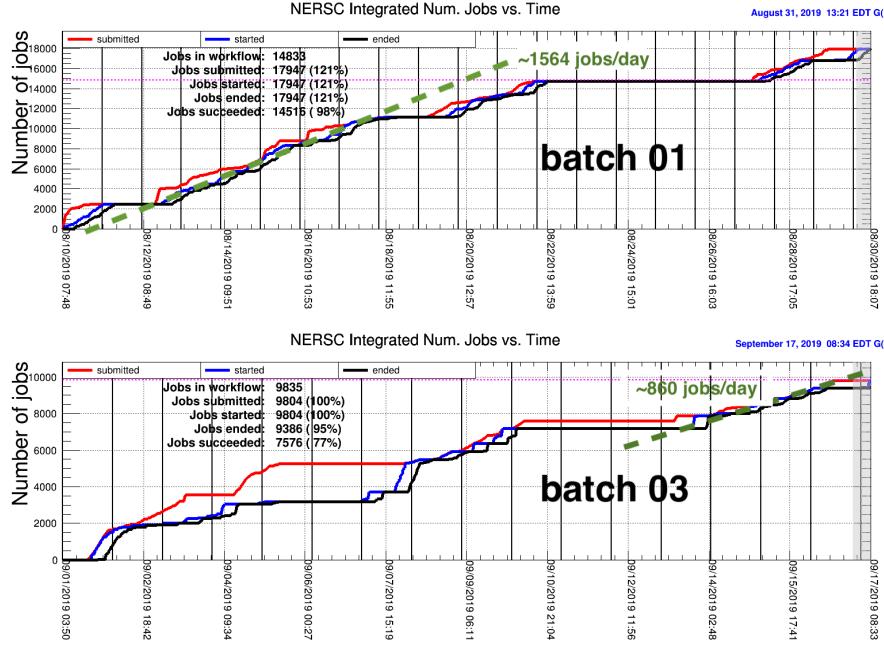


Figure 7: RunPeriod2018-08 data processing at NERSC for batches 01 and 03. The top plot (batch 01) used the “regular” queue on Cori II while the bottom plot (batch 03) used the “low” priority queue on Cori II. The green dotted lines are estimates of the job processing rates for regions where a reasonably steady state was achieved.

3.5 Low Priority Queue

We currently believe most jobs run on Cori II are scheduled via the back filling mechanism (see previous section). If this is the case, then using the *low* priority queue instead of the *regular* queue would not significantly affect our throughput in a campaign. Using the *low* queue charges half as much of a NERSC allocation as using the *regular* queue. For the RunPeriod2018-08 recon campaign, we used the *regular* queue for batch01 and the *low* queue for batch03. Figure 7 shows the job throughput vs. time for both of these batches. These indicate that the steady-state throughput of jobs using the regular queue (~ 1564 jobs/day) was roughly twice that when using the low priority queue (~ 860 jobs/day). Note that these rates are taken by looking at the throughput during a series of days when the rate was fairly constant. It is not possible to draw any definitive conclusions from these tests since both are strongly affected by what other jobs were queued on Cori II at the times they were run. They do show, however, that it is possible to run on the low priority queue with a reasonable throughput, albeit reduced.

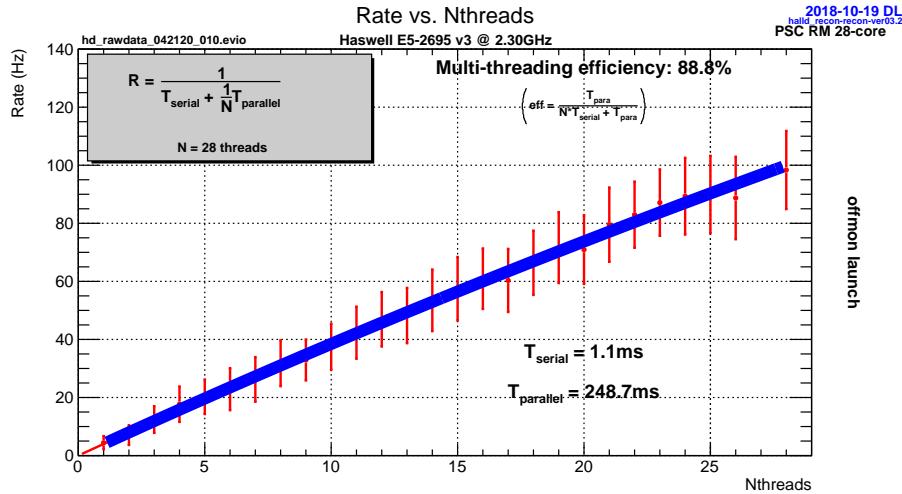


Figure 8: Thread scaling for PSC Bridges RM node.

4 PSC Bridges

This year marked the first ever production processing of JLab experimental data at the Pittsburgh SuperComputing Center (PSC). Figures 8 and 9 show the event processing rate and memory usage respectively as a function of the number of threads. These were run on the PSC “Regular Memory” nodes which contain 28 core Intel x86 CPUs. The nodes are configured to *not* use hyper-threading so the scaling curve in figure 8 has only one slope.

4.1 Allocation Award

We were awarded an XSEDE allocation on PSC for 5.9M SUs for the term starting October 1, 2019. Prior to receiving the award an advance was requested for 0.85M SUs, 10% of the total amount requested in the full proposal. The advance was granted and it was used to process 70% of the RunPeriod2018-08 batch04 data in September 2019. This consisted of 6989 jobs which used 805k of the 850k advance.

4.2 Job Rates

The job throughout at PSC was fairly steady during the course of the campaign. Figure 10 shows the jobs as a function of time. A rate of roughly 300 jobs/day was maintained during a more than 2 week period. The job processing times

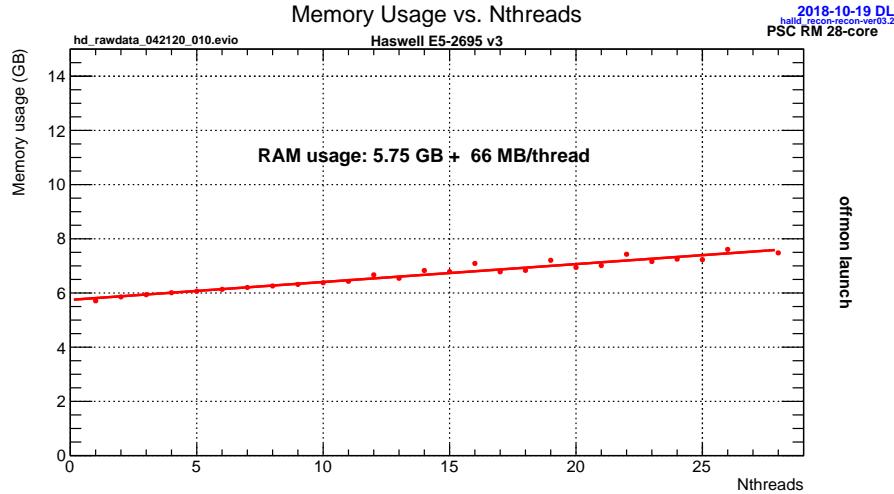


Figure 9: Memory usage vs. number of threads for GlueX reconstruction job on PSC Bridges RM node.

were very steady at about 4.25 hours. There was a brief period of time on 9/5 when jobs started timing out (see bump in bottom plot of figure 10). This was due to an issue at PSC with CVMFS. It was corrected and a portion of our allocation refunded so that those jobs could be re-run without penalty. Aside from that, the overall failure rate for jobs at PSC was only about 0.2%, much lower than what has been observed at NERSC ($\sim 2\%$).

5 Anticipated Resource Needs for 2020

GlueX anticipates collecting 7.3PB of (uncompressed) data during RunPeriod2019-11 (see Appendix C). Of this, 6.2PB is expected to be flagged a “good” runs for processing. The data is split into 20GB files with each file processed as a separate job. Thus, there will be 310k jobs to process. The remaining XSEDE allocation at PSC is enough to process 43k jobs. The remaining 267k jobs will be done using NERSC. Jobs run on Cori II processing a 20GB file take approximately 6.75 hrs to complete. NERSC charges 90 “NERSC hours” for each node hour used. We will therefore need 162M NERSC hours to complete 267k jobs. This is summarized in table 3.

The JLab SciComp Farm has 14k cores (haswell equivalent). Assuming 75% utilization averaged over a year, this is equivalent to 92M core-hours. This capacity though must be shared with the other 3 experimental halls at JLab. Hall-D is allocated 40.5% of this or 37.3Mhr. Hall-D will have significant demand for calibration, monitoring, analysis, and other misc. jobs that are not part of

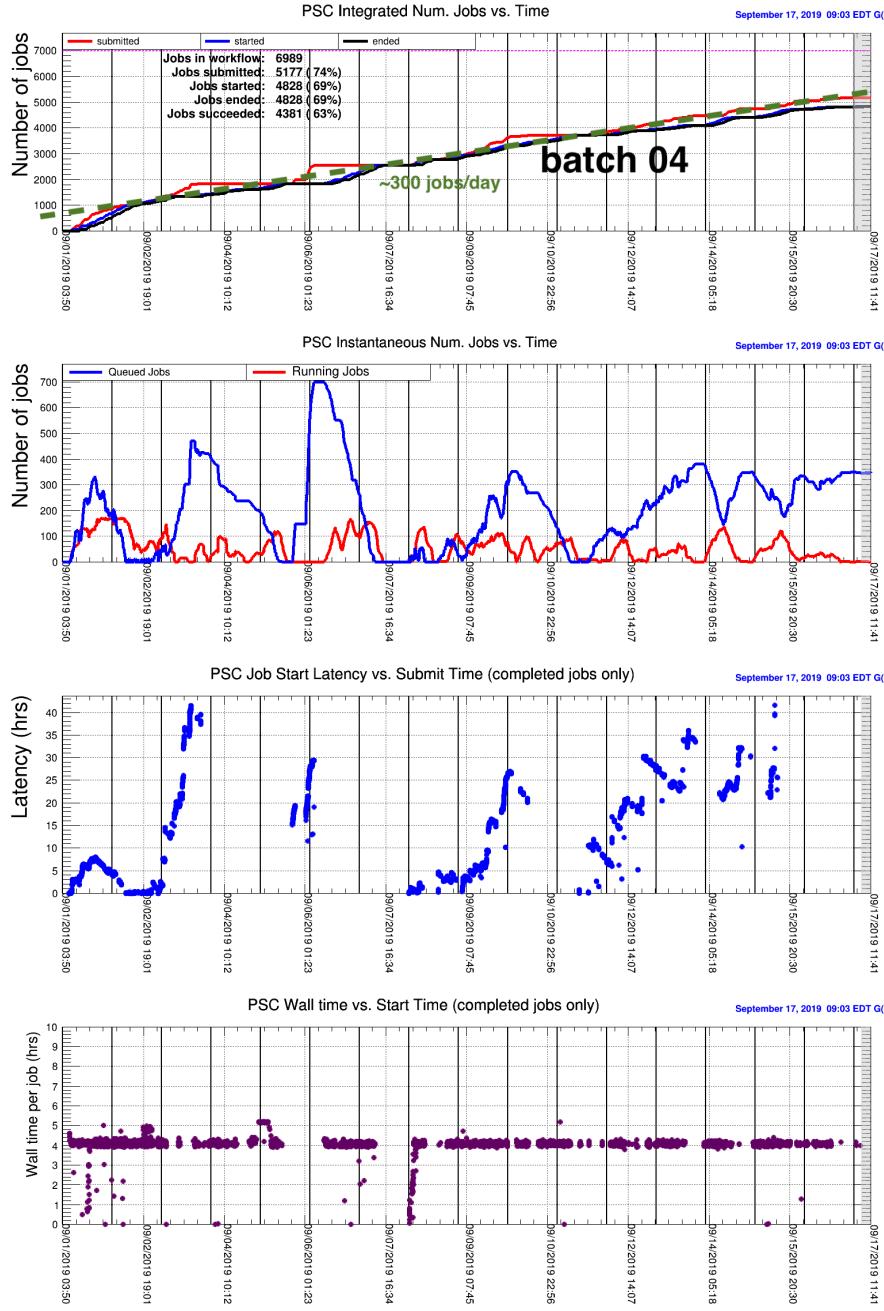


Figure 10: RunPeriod2018-08 data processing at PSC for batch 04. The top plot shows the integrated number of jobs vs. time while the second one shows the instantaneous jobs vs. time. The 3rd plot shows the latency between job submission and job start times while the 4th shows wall clock time for each job.

Total data to transfer to Offsite (NERSC + PSC)	6.2PB (7.3PB×0.85)
Total jobs to be run at PSC	42k
XSEDE units per job PSC Bridges RM	119
XSEDE units required for PSC Bridges RM	5.0M
Total NERSC units required	165M
Total jobs to be run at NERSC	267k
NERSC units per job Cori II	608
NERSC units required for Cori II	162M
Total NERSC units required	165M

Table 3: Estimated NERSC units required by GlueX for AY2020.

the reconstruction pass proper. According to the GlueX computing model, these other jobs are expected to amount to 36.2Mhr. Thus, for 2020, JLab SciComp resources are not expected to contribute significantly to the reconstruction of RunPeriod2019-11 data.

5.1 Disk Space at NERSC

Disk space usage is determined by the maximum number of nodes we may feed simultaneously at steady state. KNL nodes can consume data at 20GB/6.75hr. It was demonstrated in 2019 that we are capable of filling the 10Gbps pipe from JLab for sustained periods of time. At a rate of 1GB/s we would be able to sustain 1,215 nodes at steady state. This means at least 1215 x 20 GB files on disk for the live jobs and another 1 files for each job for queued files. This totals 48.6TB of input raw data. Space for the output will be an additional 33% of that or 16.1 TB. All of these files will need only temporary storage at NERSC so are best suited for the scratch disk. A minimum of 64.7TB is therefore required. An additional small amount would be desirable for use outside of production running (e.g. potential monitoring launch over data temporarily kept at NERSC). Thus, the total request will be for **70TB** of scratch disk space.

A RunPeriod-2019-11 : Calendar year 2019 only

```
=====
          PAC Time: 1.8 weeks
          Running Time: 3.6 weeks
          Running Efficiency: 50%
-----
          Trigger Rate: 90.0 kHz
          Raw Data Num. Events: 82.6 billion (good production runs only)
          Raw Data compression: 1.00
          Raw Data Event Size: 13.0 kB
          Front End Raw Data Rate: 1.20 GB/s
          Disk Raw Data Rate: 1.20 GB/s
          Raw Data Volume: 1.294 PB
          Bandwidth to offsite: 492 MB/s (all raw data in 1 month)
          REST/Raw size frac.: 15.00%
          REST Data Volume: 0.388 PB (for 2.0 passes)
          Analysis Data Volume: 0.547 PB (ROOT Trees for 2.82 passes)
          Total Real Data Volume: 2.2 PB
-----
          Recon. time/event: 182 ms (5.5 Hz/core)
          Available CPUs: 10000 cores (full)
          Time to process: 5.0 weeks (all passes)
          Good run fraction: 0.85
          Number of recon passes: 2.0
          Number of analysis passes: 2.82
          Reconstruction CPU: 8.3 Mhr
          Analysis CPU: 0.863 Mhr
          Calibration CPU: 0.9 Mhr
          Offline Monitoring CPU: 1.7 Mhr
          Misc User CPU: 16.4 Mhr
          Incoming Data CPU: 0.080 Mhr
          Total Real Data CPU: 28.2 Mhr
-----
          MC generation Rate: 25.0 Hz/core
          MC Number of passes: 2.0
          MC events/raw event: 0.40
          MC data volume: 0.155 PB (REST only)
          MC Generation CPU: 0.7 Mhr
          MC Reconstruction CPU: 3.3 Mhr
          Total MC CPU: 4.1 Mhr
-----
          TOTALS:
          CPU: 32.3 Mhr
          TAPE: 2.4 PB
```

B RunPeriod-2019-11 : Calendar year 2020 only

```
=====
          PAC Time: 8.4 weeks
          Running Time: 16.7 weeks
          Running Efficiency: 50%
-----
          Trigger Rate: 90.0 kHz
          Raw Data Num. Events: 386.7 billion (good production runs only)
          Raw Data compression: 1.00
          Raw Data Event Size: 13.0 kB
          Front End Raw Data Rate: 1.20 GB/s
          Disk Raw Data Rate: 1.20 GB/s
          Raw Data Volume: 6.056 PB
          Bandwidth to offsite: 2303 MB/s (all raw data in 1 month)
          REST/Raw size frac.: 15.00%
          REST Data Volume: 1.817 PB (for 2.0 passes)
          Analysis Data Volume: 2.562 PB (ROOT Trees for 2.82 passes)
          Total Real Data Volume: 10.4 PB
-----
          Recon. time/event: 182 ms (5.5 Hz/core)
          Available CPUs: 10000 cores (full)
          Time to process: 23.2 weeks (all passes)
          Good run fraction: 0.85
          Number of recon passes: 2.0
          Number of analysis passes: 2.82
          Reconstruction CPU: 39.1 Mhr
          Analysis CPU: 4.038 Mhr
          Calibration CPU: 4.2 Mhr
          Offline Monitoring CPU: 7.7 Mhr
          Misc User CPU: 16.4 Mhr
          Incoming Data CPU: 0.376 Mhr
          Total Real Data CPU: 71.7 Mhr
-----
          MC generation Rate: 25.0 Hz/core
          MC Number of passes: 2.0
          MC events/raw event: 0.40
          MC data volume: 0.727 PB (REST only)
          MC Generation CPU: 3.4 Mhr
          MC Reconstruction CPU: 15.6 Mhr
          Total MC CPU: 19.1 Mhr
-----
          TOTALS:
          CPU: 90.8 Mhr
          TAPE: 11.2 PB
```

C RunPeriod-2019-11 : Total

```
=====
          PAC Time: 10.1 weeks
          Running Time: 20.3 weeks
          Running Efficiency: 50%
-----
          Trigger Rate: 90.0 kHz
          Raw Data Num. Events: 469.3 billion (good production runs only)
          Raw Data compression: 1.00
          Raw Data Event Size: 13.0 kB
          Front End Raw Data Rate: 1.20 GB/s
          Disk Raw Data Rate: 1.20 GB/s
          Raw Data Volume: 7.350 PB
          Bandwidth to offsite: 2795 MB/s (all raw data in 1 month)
          REST/Raw size frac.: 15.00%
          REST Data Volume: 2.205 PB (for 2.0 passes)
          Analysis Data Volume: 3.109 PB (ROOT Trees for 2.82 passes)
          Total Real Data Volume: 12.7 PB
-----
          Recon. time/event: 182 ms (5.5 Hz/core)
          Available CPUs: 10000 cores (full)
          Time to process: 28.2 weeks (all passes)
          Good run fraction: 0.85
          Number of recon passes: 2.0
          Number of analysis passes: 2.82
          Reconstruction CPU: 47.4 Mhr
          Analysis CPU: 4.901 Mhr
          Calibration CPU: 5.1 Mhr
          Offline Monitoring CPU: 9.4 Mhr
          Misc User CPU: 16.4 Mhr
          Incoming Data CPU: 0.456 Mhr
          Total Real Data CPU: 83.6 Mhr
-----
          MC generation Rate: 25.0 Hz/core
          MC Number of passes: 2.0
          MC events/raw event: 0.40
          MC data volume: 0.882 PB (REST only)
          MC Generation CPU: 4.2 Mhr
          MC Reconstruction CPU: 19.0 Mhr
          Total MC CPU: 23.1 Mhr
-----
          TOTALS:
          CPU: 106.7 Mhr
          TAPE: 13.5 PB
```