Computing & Software, A.I. Report

C. Fanelli, D. Lawrence

ECCE Bi-weekly Meeting
June 21 2021

Farm fractions

SDCC RACF: 50k cores https://www.bnl.gov/compsci/SDCC/infrastructure.php

JLab SciComp: 13.9k cores https://scicomp.jlab.org/scicomp/index.html#/farmNodes

assumption: we will utilize 75% of the 2k cores averaged over 4 months
i.e. running jobs 75% of time and not running jobs 25% of time

	Total Cores	Pledged Cores (ECCE only)	% of farm (4 mo.) (ECCE only)	% of farm (4 wk.) (ECCE only)
BNL	50k	2k*0.75/2 = 0.75k	1.5%	6.4%
JLab	13.9k	2k*0.75/2 = 0.75k	5.4%	23%
OSG	N/A	4k/2 = 2k	2k-cores	8.6k-cores

assume ATHENA will use similar, but may or may not overlap with our compressed schedule

Transferring data between BNL and JLab

Transferred SIDIS 18x100 and 18x275 generated files directories to JLab via globus:

/~/gpfs02/eic/DATA/YR_SIDIS/ → /~/expphy/volatile/eic/davidl/DATA/YR_SIDIS/

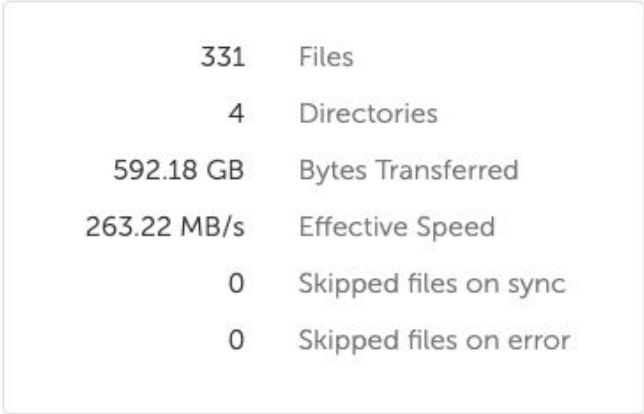

331	Files
4	Directories
592.18 GB	Bytes Transferred
263.22 MB/s	Effective Speed
0	Skipped files on sync
0	Skipped files on error

Available via xrootd:

```
> singularity shell /cvmfs/eic.opensciencegrid.org/singularity/rhic_sl7_ext.simg
> export LD_PRELOAD=/usr/lib64/libXrdPosixPreload.so
> ls root://sci-xrootd.jlab.org//osgpool/eic/DATA/YR_SIDIS/ep_18x100
```

Multiple sources now configured and available for production

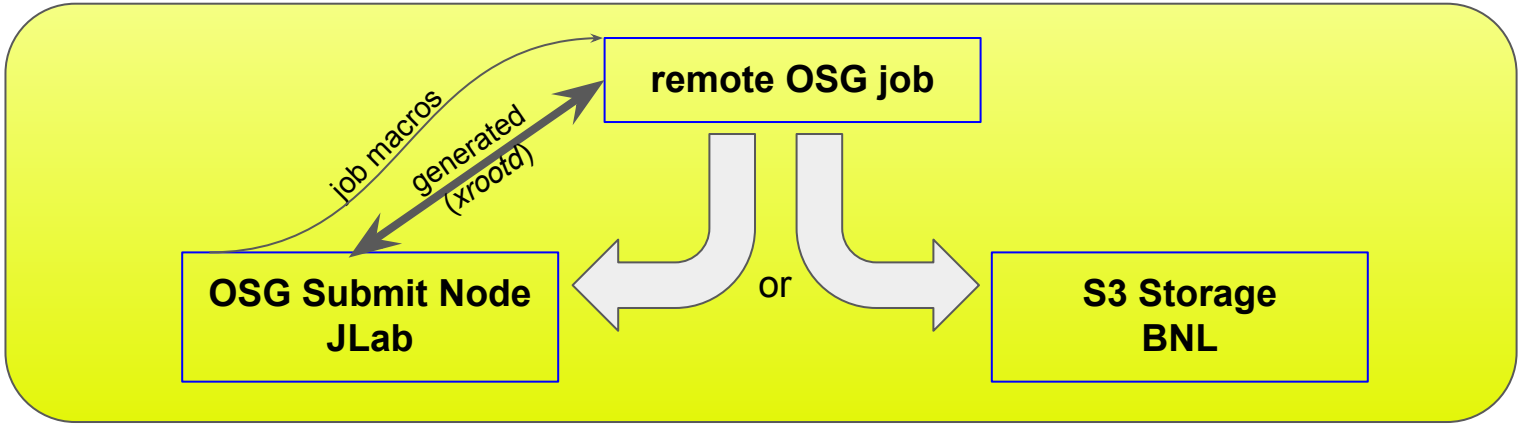

JLab
SLURM
/work/eic2 0.5PB
/work/eic3 0.5PB

BNL
condor
S3 (object store) 1PB

MIT/Bates
opportunistic

S3 Storage
BNL

Test Production Jobs on OSG and JLab SciComp

(HTCondor) *(SLURM)*

SIDIS 18x100 (pythia6) 100k events (100 jobs @1k each)

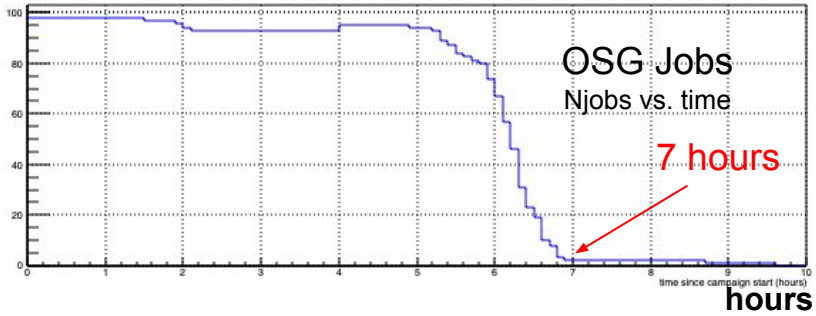

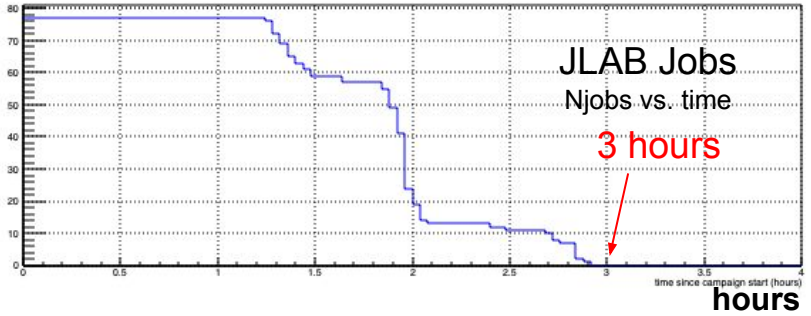

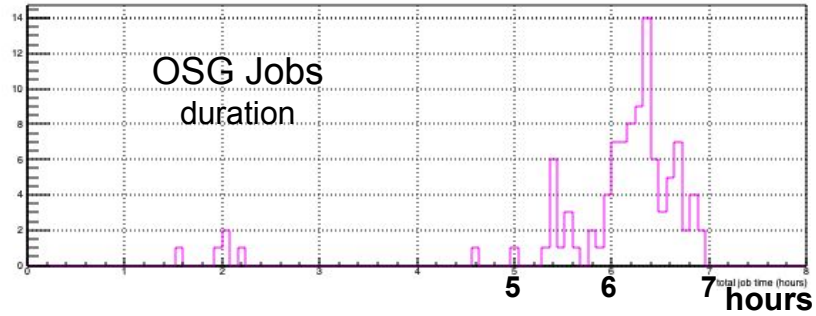

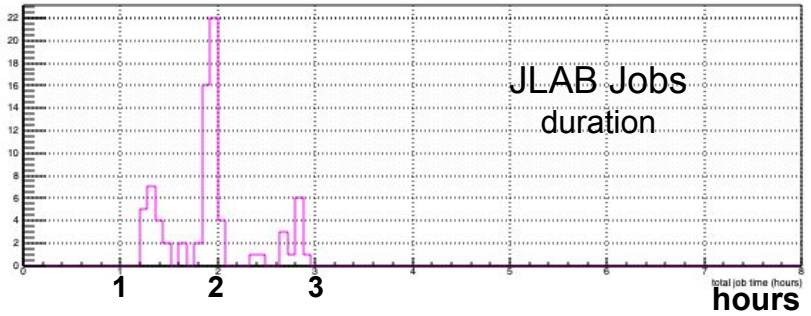

Results and Summary

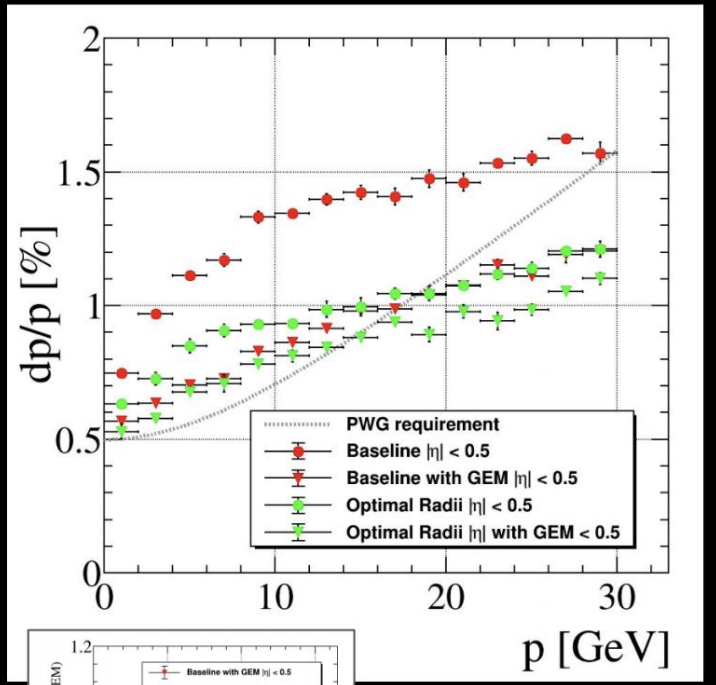

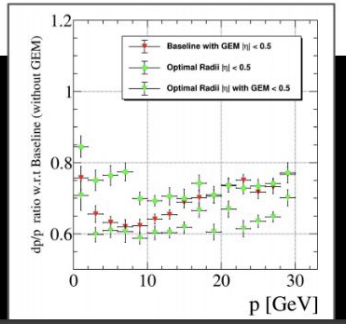

- Preliminary results from AI optimizing simultaneously momentum, polar and azimuthal angles resolutions, suggest to push the outermost two layers at large radii and close to each other (see https://indico.bnl.gov/event/12052/contributions/50502/attachments/34967/56877/AI_Supported_Detector_Design.pdf). This provided a substantial improvement in momentum resolution as compared to the baseline design.

- We received a comment to leave like 5 cm between these two layers. We checked this “manually” and showed (left figure) that the effect on resolution is almost negligible.

- The obtained resolution is also close to what one would obtain with an extra GEM layer outside the DIRC

Backups

Utilizing BNL and JLab storage for remote jobs

- JLab
  - Use built-in GridFTP mechanism
    - will require 1PB storage be mounted on submit node. Need to check if this is issue.

- BNL
  - Write to S3 storage from remote node
  - Requires minio client. Trivial to add to singularity container:
    - */cvmfs/oasis.opensciencegrid.org/jlab/epsci/singularity/images/rhic_sl7_ext_S3.simg*
  - Authentication:
    - transfer script from secure area with job
    - run script using modified HOME to install secrets in cwd on remote node
    - delete *".s3"* (aka *".mc"*) directory when job finishes

```
1 #!/bin/bash
2
3 # "s3" is a symbolic link to /usr/local/bin/mc
4 export HOME=$PWD
5 s3 config host add eic https://dtn01.sdcc.bnl.gov:9000/ [redacted] [redacted]
```

OSG

Potential issue:

- Generated events files are large (5GB for 2M events for SIDIS)
  - Standard job size of 2k events uses 0.1% of this
  - Significant bandwidth overhead to send whole file to every job that needs it
    - S3 access requires whole file transfers
    - stashcache lets OSG distribute copies so nearest one can be copied to job node
    - xrootd allows random access limiting required bandwidth
      - channels all requests to single server (possible bandwidth issue, but probably not)
      - JLab internal jobs require different path than external ones
  - Alternative: create EIC-smear files with only 2k events

***recommendation**: For OSG jobs use xrootd hosted at JLab for input files.*

1. Allows us to use YP generated events files
2. Cameron's production scripts already set up for splitting files into finer jobs
3. *rhic_sl7_ext.simg* singularity image already contains xrootd

Strawman schedule (I've not seen Cameron's timeline so this is based on what I heard at the Bi-weekly meeting)

week of	activity
June 14	10M particle gun events + validation
June 21	**Physics generators campaign I**
June 28	**Physics generators campaign I** + Analyze campaign I data
July 5	Analyze campaign I data
July 12	Implement changes and validate in preparation for campaign II
July 19	**Physics generators campaign II**
July 26	**Physics generators campaign II** + ...

ECCE Compute Estimate

Disclaimer:
These numbers are preliminary and based on a very small sample test using events generated for sPHENIX and not necessarily ECCE physics

2 jobs run on OSG nodes

Default events from tutorial

https://www.phenix.bnl.gov/WWW/publish/phnxbld/sPHENIX/files/sPHENIX_G4Hits_sHijing_9-11fm_00000_00010.root

Total committed by [(BNL + JLab)*75% + OSG] / 2 = 4k cores x 4 months (x 75%) = **10 Mcore-hrs**

expect ~ ½ for ECCE

102 Events - avg. time : 2,644 sec (single thread)
2 Events - avg. time : 215 sec (single thread)

time/event: 24.3 sec
overhead: 166 sec

assume 2k events/job

for **2.25B** events:
15.2 Mcore-hrs for event simulation
51.9 kcore-hrs for overhead *(program startup and shutdown)*

for **1.0B** events:
6.8 Mcore-hrs for event simulation
23.1 kcore-hrs for overhead *(program startup and shutdown)*