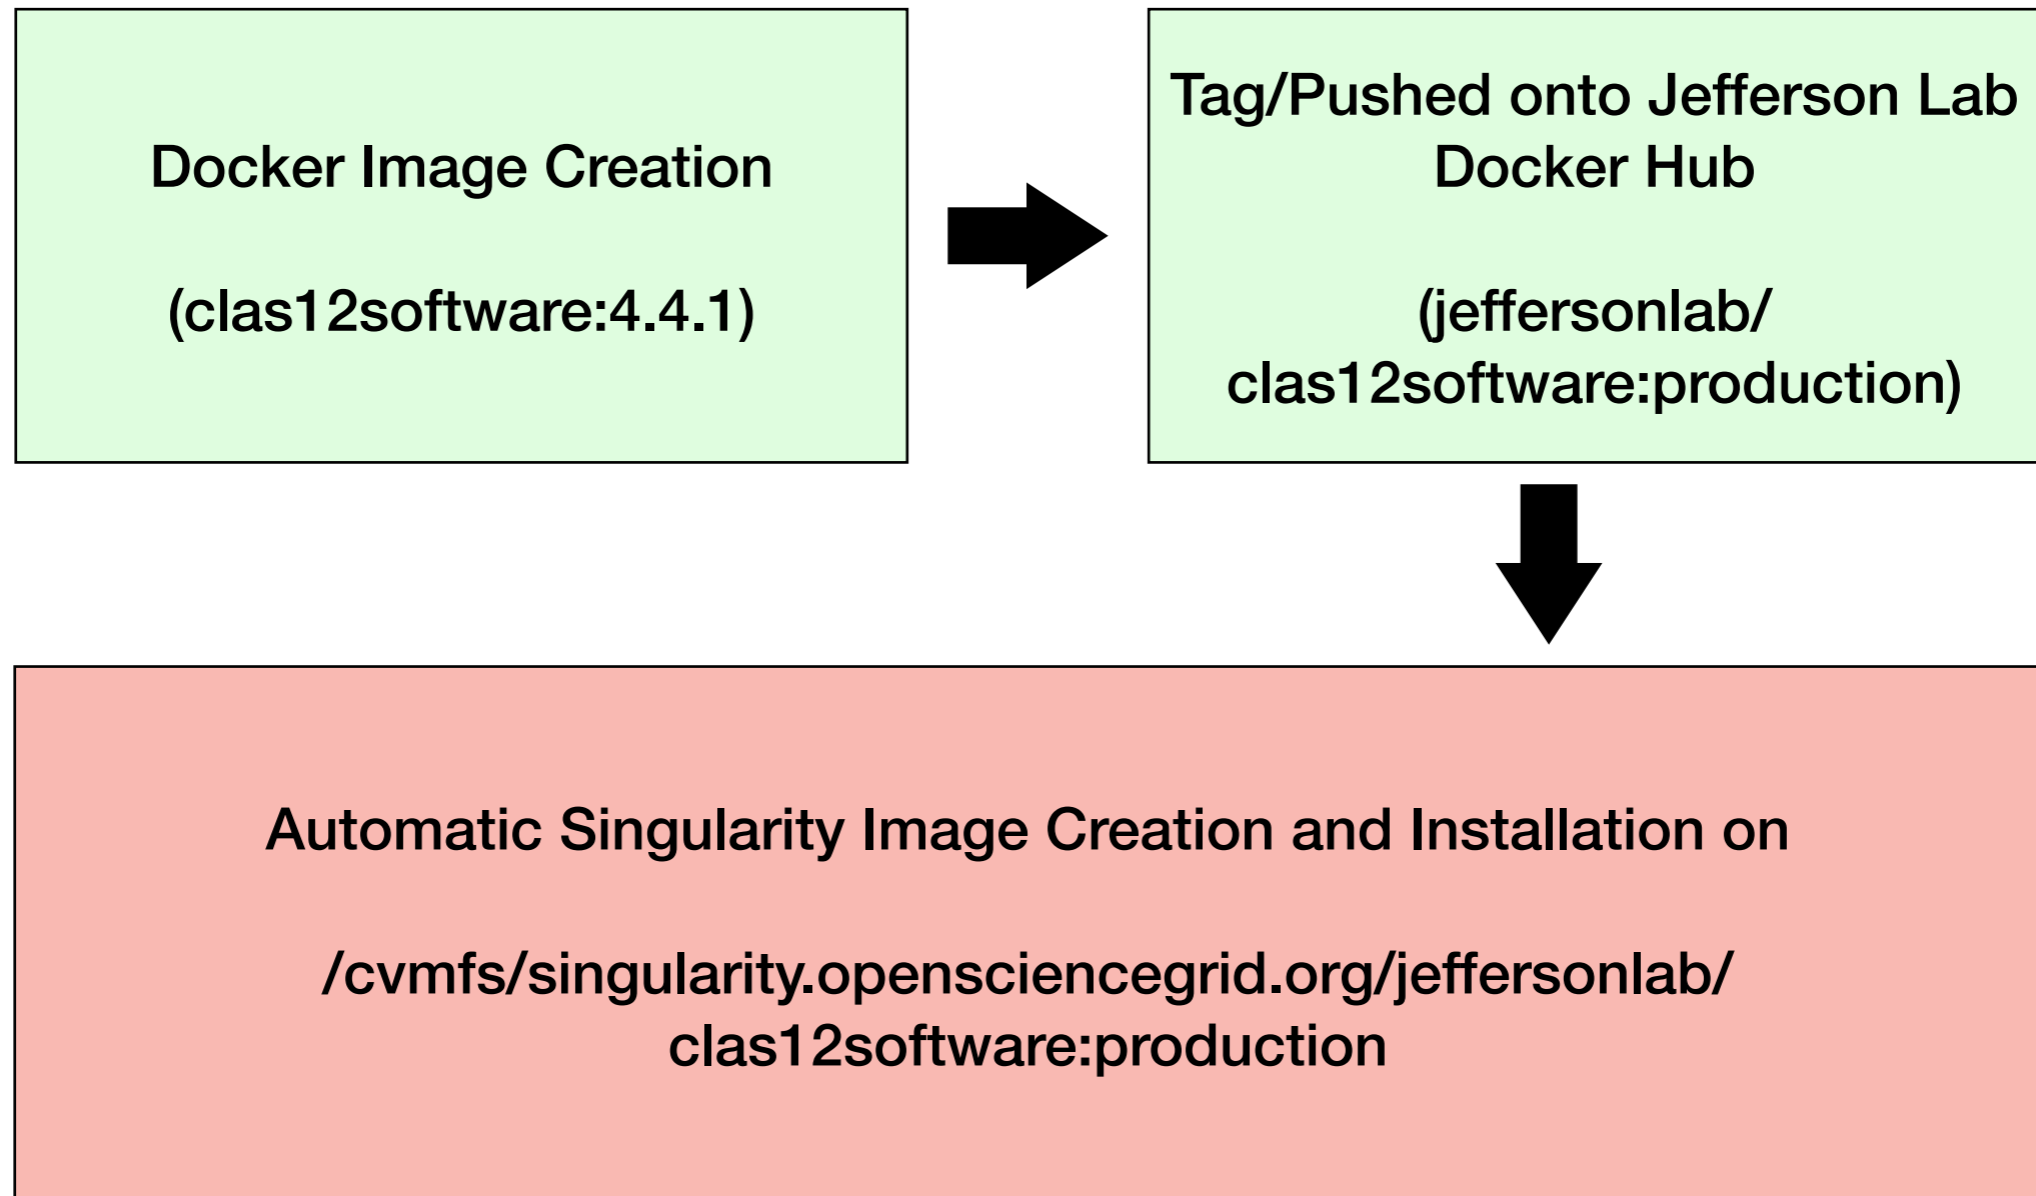


OSG Status Report

M. Ungaro, N. Baltzell
CLAS Collaboration Meeting
June 1, 2021

CLAS12 OSG Images



Both docker and singularity images rely on /cvmfs being mounted to access environment, fields, software

CLAS12 OSG Submissions Workflow

Web Portal on gemc.jlab.org
(php, javascript, html)

| | | |
|---|--|------------------------|
| Configuration | <input type="text" value="rgb_fall2019"/> | |
| Magnetic Fields | <input type="text" value="tor+1.00_sol-1.00"/> | |
| Generator | <input type="text" value="clasdis"/> | |
| Generator Options | <input type="text" value="--t 15 25"/> | clasdis options |
| <p>After selecting the generator, check the documentation and paste the needed options above. Notice: do no use the following options as they are automatically passed for you: --docker, output file name --trig options.</p> | | |
| Number of Events / Job | <input type="text" value="10000"/> | |
| Number of Jobs | <input type="text" value="2000"/> | |
| Total Number of Events | <input type="text" value="20"/> M | |
| Background Merging | <input type="text" value="40nA_10410MeV"/> | |
| <input type="button" value="Submit"/> | | |



Define and upload
submission parameters to
JLAB MYSQL server

submit.jlab.org



**Cronjob on JLAB Submit Node
(scosg16)**

- Reads submissions pars
- Create Condor scripts
- Create Run scripts
- Fetch LUND files if applicable
- Submit Jobs to OSG
- Modify Submission Status on submit.jlab.org

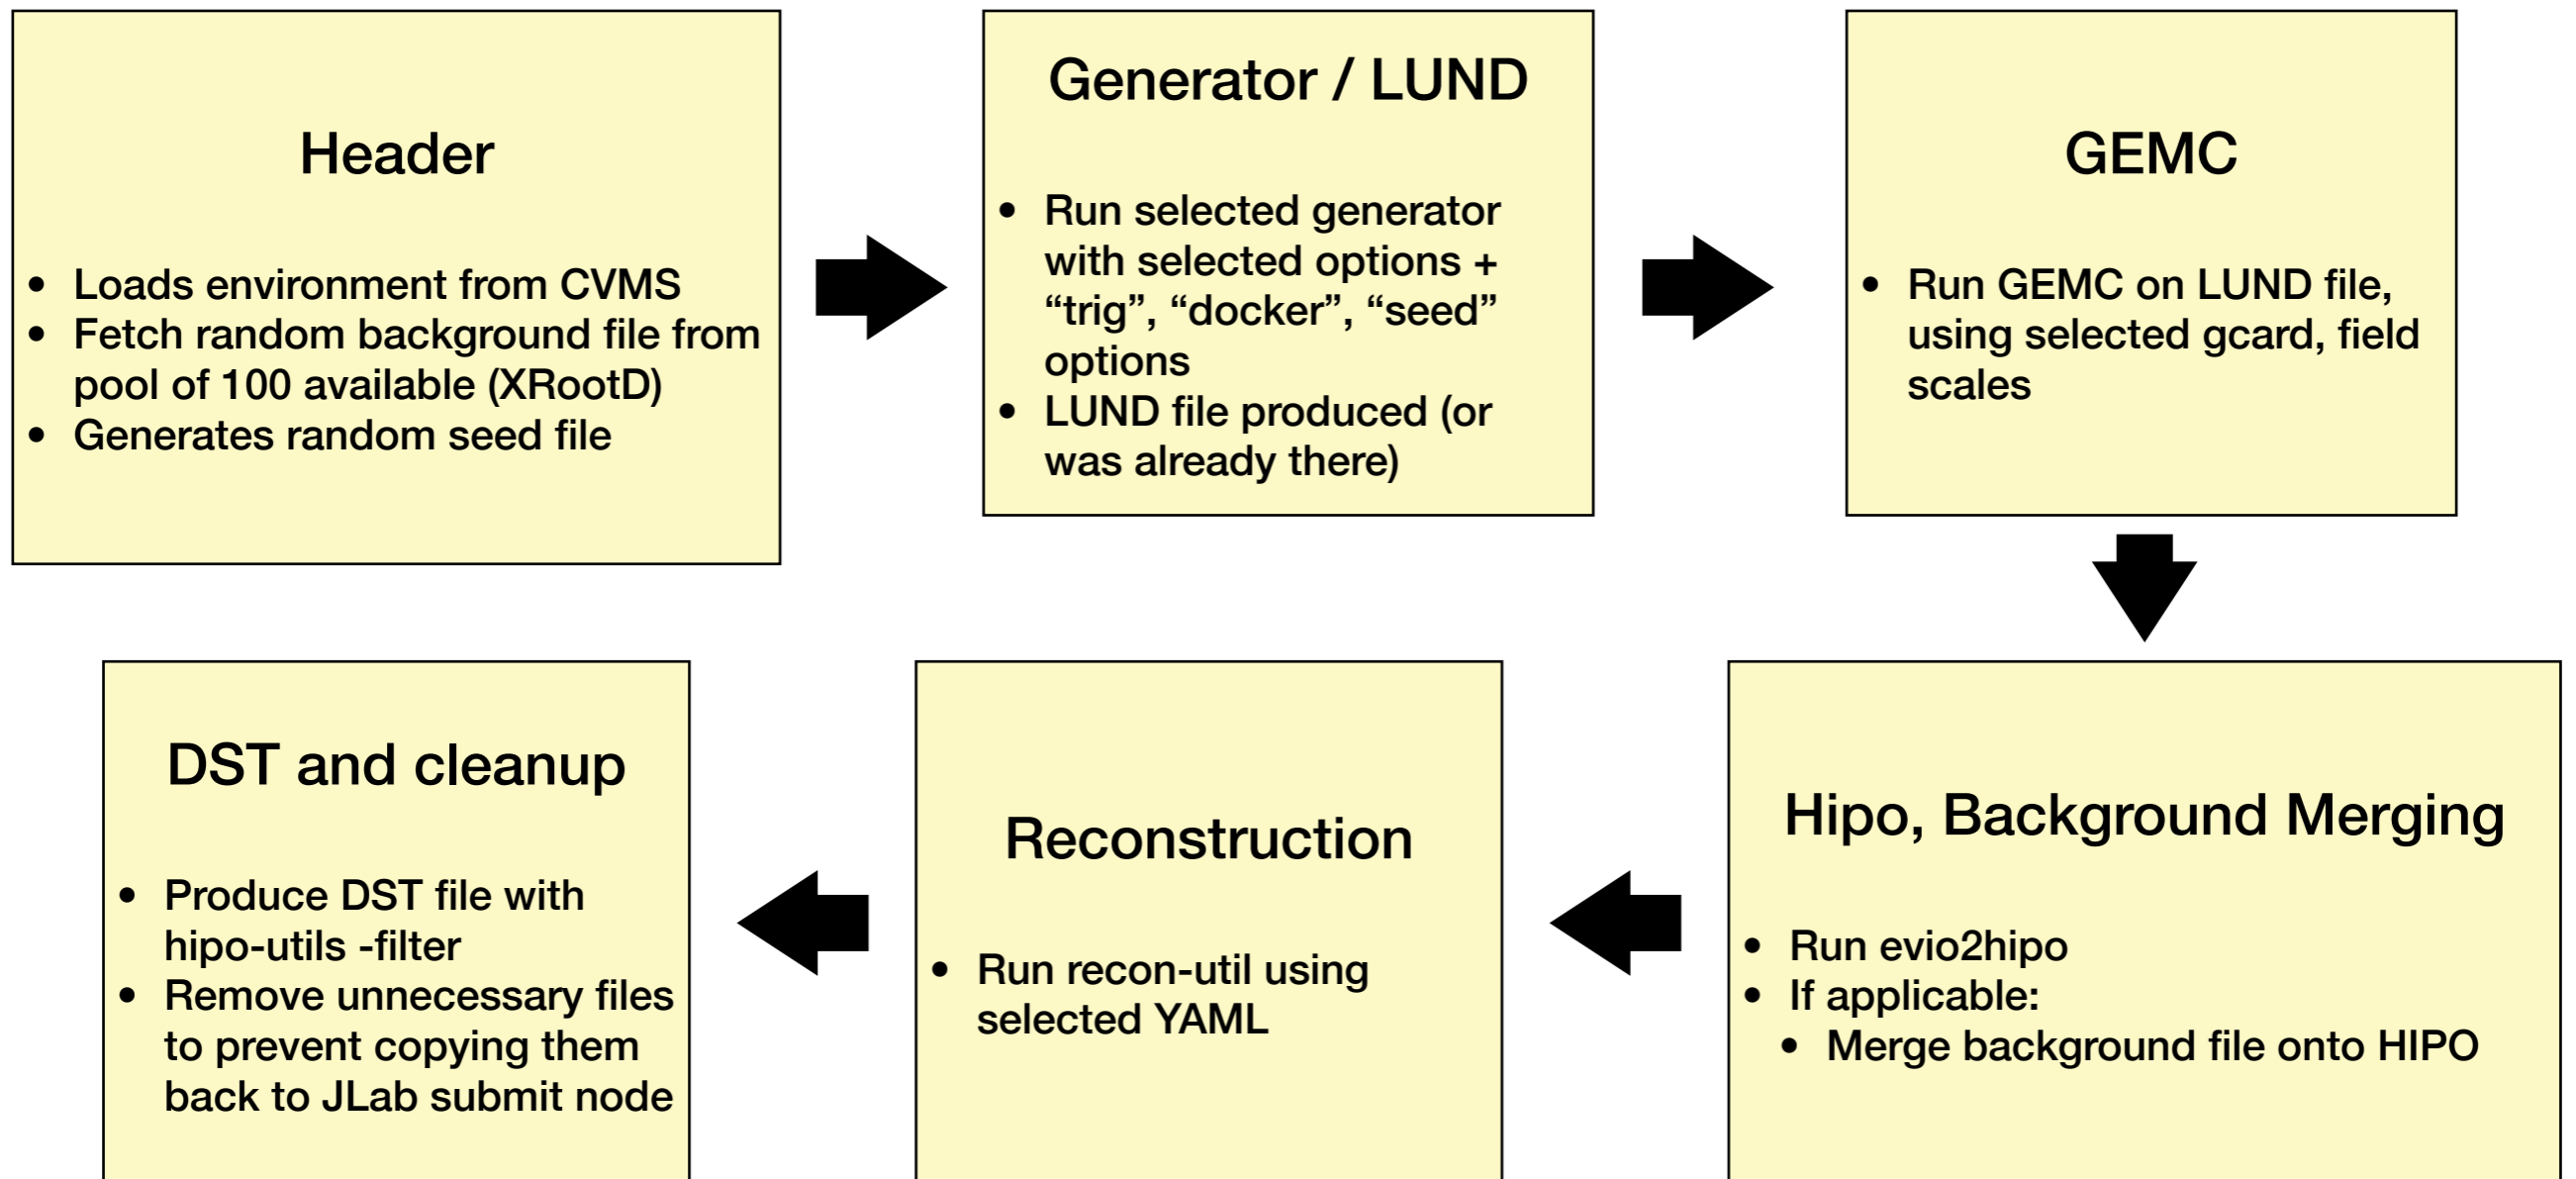
Additional condor parameters:

Offsite Resources "Rank"

Offsite Resources minimum requirements

CLAS12 OSG Node Workflow

nodescript.sh is run on the offsite node



At any of these stages, if an error occurs:

- a file cleanup is done
- the job exits with dedicated exit codes

CLAS12 OSG Additional Workflows

Cronjobs on scog16, python backends

Monitoring

- Reads Submission Status from JLAB MYSQL Server jsubmit
- Reads gemc jobs from Condor
- Writes JSON file with summary and detailed status report

Priority

- Reads User priority from JLAB MYSQL Server jsubmit
- Adjusts job priority, based on running jobs

Summary of current jobs

| user | submission | total | done | run | idle |
|----------|------------|-------|-------|-------|------|
| kneupane | 4 | 31000 | 9999 | 12944 | 8028 |
| osoto | 4 | 40000 | 38915 | 1083 | 2 |
| tyson | 1 | 501 | 0 | 496 | 5 |
| total | 9 | 71501 | 48914 | 14523 | 8035 |

Details of current OSG Jobs

| user | job id | submitted | total | done | run | idle | hold | osg id |
|----------|--------|-------------|-------|------|------|------|------|---------|
| kneupane | 2952 | 05/07 08:32 | 10000 | 1 | 9967 | 29 | 3 | 3610109 |
| osoto | 2951 | 05/06 12:04 | 10000 | 9044 | 954 | 2 | 0 | 3604617 |
| osoto | 2949 | 05/06 12:04 | 10000 | 9990 | 10 | 0 | 0 | 3604613 |
| osoto | 2950 | 05/06 12:04 | 10000 | 9883 | 117 | 0 | 0 | 3604616 |
| osoto | 2935 | 05/03 21:58 | 10000 | 9998 | 2 | 0 | 0 | 3592222 |
| kneupane | 2940 | 05/04 10:06 | 10000 | 9998 | 2 | 0 | 0 | 3595470 |
| kneupane | 2953 | 05/07 10:20 | 10000 | 0 | 2975 | 6999 | 26 | 3610164 |
| kneupane | 2954 | 05/07 10:20 | 1000 | 0 | 0 | 1000 | 0 | 3610165 |
| tyson | 2955 | 05/07 10:55 | 501 | 0 | 496 | 5 | 0 | 3610176 |

Offsite Dedicated Resources

*1 core = 24*365 = 8760 hours / year*

| OSG Name | Country/ Institution | Number of CPUs | MHours / Year |
|---------------------------------------|--|-------------------|------------------|
| INFN-T1 | Istituto Nazionale di Fisica Nucleare, Italy | 400 | 3.5 |
| UKI- SCOTGRID- GLASGOW | Glasgow University, Scotland | 700 | 6.1 |
| MIT* | Massachusetts Institute of Technology, USA | 320 | 4.4 |
| Upcoming: LAMAR | Lamar University, USA | 1-200 | 1.7 |

Total Dedicated:

1,420 CPUs

(1-200 upcoming)

12.4 MH / Year

* accepts other projects if cores are idles. May expand # cores.

Offsite High Priority and Opportunistic Resources

| OSG Name | Country/Institution | Number of CPUs | MHours / Year |
|--------------|--|----------------|---------------|
| UCONN | University of Connecticut, USA | 0-2,000 | 0-17.5 |
| SU | Syracuse University | 0-12,000 | 0-105 |
| GRIF | Grille au service de la Recherche en Ile-de-France, France | 0-500? | 0-4.4? |
| OSG | OSG Opportunistic | 0-3000 | 26.2 |

Not dedicated but CLAS12 has high priority on these

Total High Priority:

0 - 14,500 CPUs

0 - 127 MHours / year

CLAS12 Project Summary of Resources

| OSG Name | Number of CPUs | MHours / Year |
|----------------------|----------------|------------------------------|
| Dedicated | 1,420 | 12.4 |
| Priority | 0 to 14,500 | 127 |
| Opportunistic | 0 to 3,000 | 26 |
| Total | 1420-18920 | ⁸ 14 - 166 |

Note: event time may vary from site to site, depending on the hardware used.


For example, SU CPUs are a factor of 2 slower than others. Need to confirm this, or if it's a software inefficiency.

CLAS12 @OSG: Last 6 months

usage (includes dedicated experiments)

Total Wall Hours for Top 10 Projects


| | total |
|------------------------|---------|
| cms | 605 Mil |
| GLOW | 42 Mil |
| cms.org.cern | 35 Mil |
| IceCube | 32 Mil |
| COVID19_FoldingAtHome | 28 Mil |
| CLAS12 | 24 Mil |
| dune | 21 Mil |
| nova | 16 Mil |
| microboone | 15 Mil |
| Hawaii_Doetinchem | 12 Mil |
| cms.org.ucsb | 12 Mil |
| CpDarkMatterSimulation | 11 Mil |
| gm2 | 11 Mil |
| WSU_3DHydro | 10 Mil |
| chemml | 10 Mil |
| CompBinFormMod | 9 Mil |
| LIGO | 7 Mil |
| cms.org.baylor | 6 Mil |
| electrolytes | 6 Mil |
| REDTOP | 6 Mil |



OSG opportunistic

| | total |
|------------------------|--------|
| COVID19_FoldingAtHome | 28 Mil |
| CLAS12 | 24 Mil |
| Hawaii_Doetinchem | 12 Mil |
| CpDarkMatterSimulation | 10 Mil |
| WSU_3DHydro | 10 Mil |
| chemml | 10 Mil |
| CompBinFormMod | 9 Mil |
| electrolytes | 6 Mil |
| REDTOP | 6 Mil |
| SourceCoding | 6 Mil |

Core Hours by Site per 1d



Output Data Staging

- **Switched to staging mechanism on the JLab submit/collector node 2.5 months ago**
 - Fast enough to keep up with ~50k simultaneous OSG jobs (*DSTs only*) with single serial rsync ~100 MB/s
 - We have peaked around 20k, average usually much less
- **Jobs no longer write directly to nor write lots of unnecessary small files to Lustre**
 - Was a significant contributor to performance issues with both Lustre and the submit node, limiting job throughput
 - “Unfortunately” that means just the top level job script (until its info is embedded in HIPO), one per submission, and HIPO files. Everything else, e.g. logs, remain only on the submit node and only for a couple weeks, can be retrieved upon request.
 - /volatile/clas12/osg2 (the old “osg” directory will eventually disappear)
 - Currently a ~1 hour lag between job success on the portal and appearance of outputs on Lustre
 - This can be addressed if there’s high demand, currently decoupled for easier reliability
- **Transfer software recovers automatically after hardware issues are fixed (e.g. Lustre/stage/network), preserves logs on all transfers, automatically cleans up old stuff (staging filesystem is finite and shared!)**



Improved Diagnostics

- **New utilities to analyze condor OSG job data** (e.g. by user/generator/site/etc, log file analysis)
- **This helped improve job workflow and submission requirements, e.g.**
 - Understanding efficiency, preemption, failure modes (true software crashes are <0.1%, usually GEANT)
 - Error detection and reporting in the jobs for immediate exit and [automatic retries](#)
 - *But have to be careful here, jobs need to be well configured, failure modes understood and not due to our software or submission mistakes, testing before submission, not retrying jobs with terminal issues*
 - Detecting and sometimes avoiding sites with huge problems until they're addressed via feedback to OSG and site admins, usually related to XRootD or CVMFS, sometimes unavailable memory
 - Preferring dedicated resources to keep them full without reducing overall throughput
- **Background processes run on the JLab submit node to automate everything:**

Frequent
Submit and Prioritize
Jobs, Update Portal
Status

Hourly
Transfer, Cleanup,
Collect Metrics,
Problem Alerts

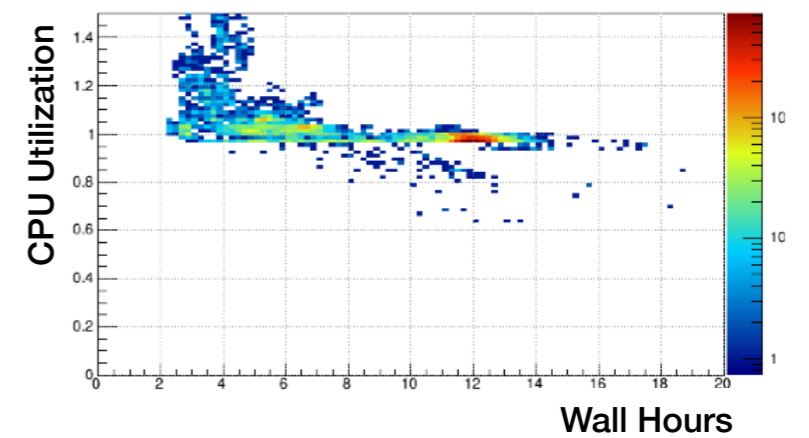
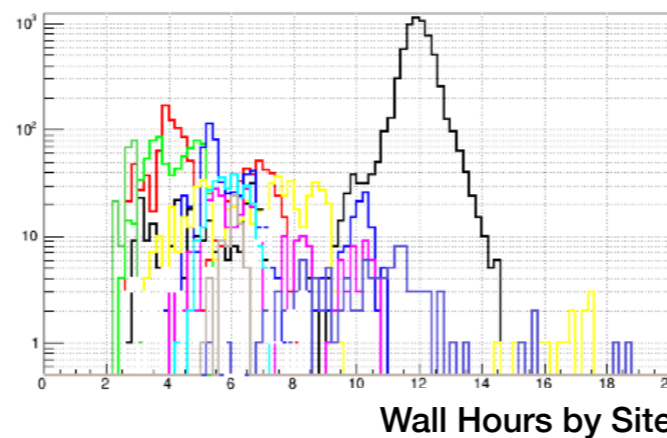



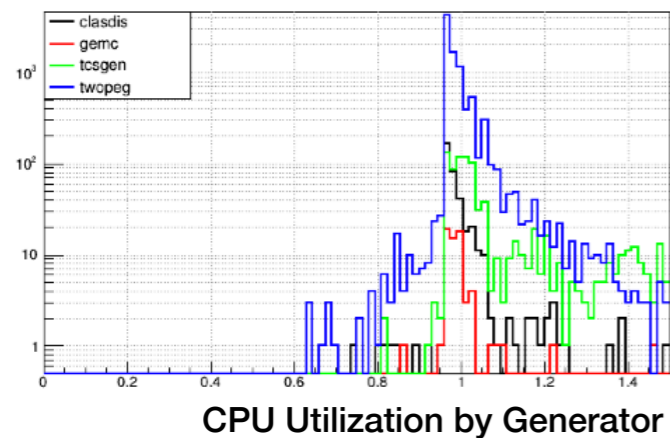


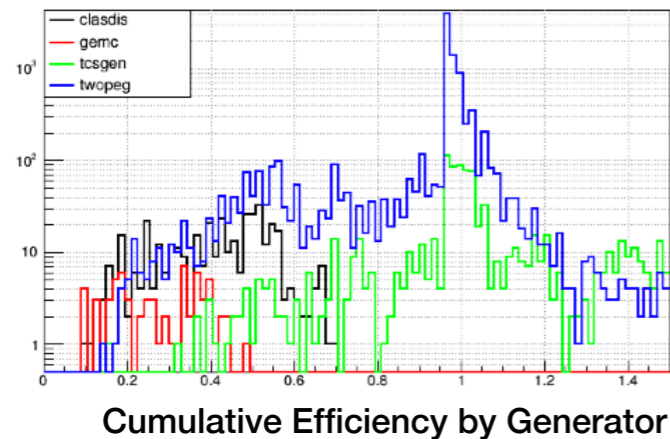
Daily
Send Diagnostic
Report

- **We now ultimately get 100% job success, hands free**
 - Average number of attempts is usually ~1.2 (preemption and node issues), so still have the occasional straggler (<<1%) that take a few tries to succeed. So while the majority of a submission may finish, it can take another day or two for that remaining job (not really a problem with that particular job, it just got unlucky).

Improved Diagnostics

We've had a case where efficiency dropped, attempts increase, due to misconfiguration which caused losing dedicated and priority sites, burning cycles on opportunistic resources by jobs that failed or got preempted. That's a red flag for OSG admins and we need to keep our jobs efficient and not waste resources.

A "normal" snapshot of jobs completed in the past 24 hours ... (note, all are log scale)



CLAS12 Project Resources Summary

- CLAS12 Project Goal for Simulations: **55 Million CPU Hours Per Year**

- That's **6k simultaneous jobs**

- In the past months, when we keep our OSG queue full:

- We routinely hit 10-15k simultaneous running jobs on OSG
 - and peaked at 20k a few times (which is the size of the JLab farm)
 - largest contributor is Syracuse, a priority site

- But this is when there's low competition for the priority and opportunistic sites, e.g. when GlueX runs they get ~half*

- Maybe we should anticipate "normal" being closer to the rightmost column

```

-----
site           run      idle
-----
-              0       22423
SU-ITS        11185   87
UChicago     4810    105
UConn-HPC    1033     0
SGridGLA     510      2
UConn        372      3
CNAF         365      1
SLATE ILS NMS 227      0
tally        19885   22653
-----
    
```

Yearly CLAS12 Computing Projections on OSG

| OSG Site Type | Number of CPU Cores | MHours | MHours*, Normalized | MHours* @ 30% share of Non-Dedicated |
|---------------|---------------------|------------|---------------------|--------------------------------------|
| Dedicated | 1420 | 12 | 12 | 12 |
| Priority | 14500 | 127 | 76 | 23 |
| Opportunistic | 3000 | 26 | 26 | 8 |
| Total | 18920 | 166 | 115 | 43 |

* = Normalized for CPU speeds at different sites