If you give a GlueX researcher a computer....

Thomas Britton
HallD Postdoc

What is GlueX ?

- GlueX is an experiment based in **HallD**
  - One of Jefferson Lab's four experimental halls.
  - Uses a polarized photon beam on liquid hydrogen to study the spectrum of mesons
  - Comprised of an international collaboration of 116 members at 27 institutions

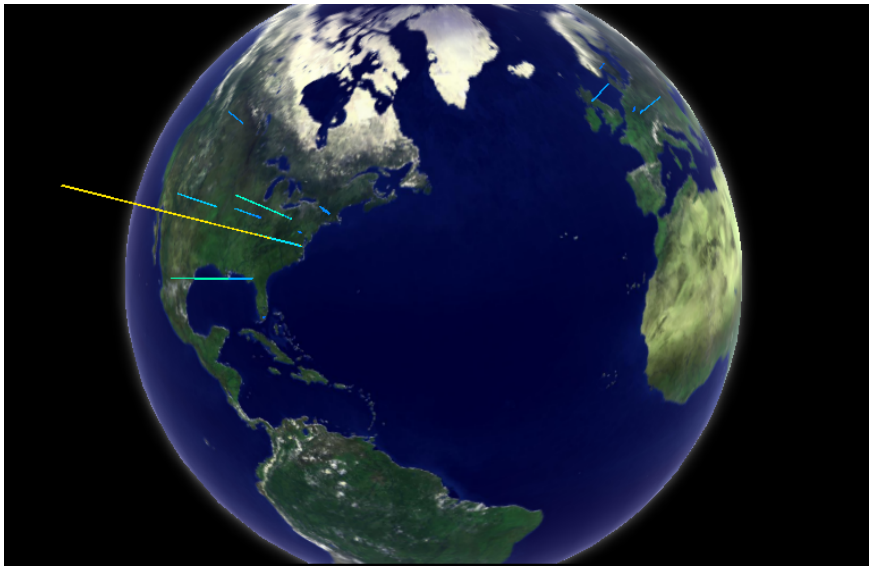

What is GlueX?

- Take 1 GB every 2 seconds.
  - This volume of data requires heavy computational lifting to:
    - Reconstruct events
    - Sift through events
    - Analyze events
    - **Simulate events**

- Most analyses are statistical in nature
  - Many require large data sets
  - No one would complain about having “too much data”
    - Unless it becomes unfeasible to analyze

If you give a GlueX researcher a computer...

- **He's going to need to run some Monte Carlo.**
  - Monte Carlo (MC) is how we simulate things, from physical processes to detector response
- MC is very parallel
  - Each event is done completely isolated
- Meaningful production of MC usually involves millions of events
- Given the lack of data to ship and multitude of variants for every study this production is usually left to individual members
- The concept for MCwrapper started when it came time for me to run Monte Carlo
  - The person I shared an office with showed me a configuration file and pointed me to all of the various packages used

Short Rant

- Lots of duplication of code
  - Easy to go wrong and never know it
- Many comparisons of MC were apples to oranges
  - Little to no provenance
    - Settings, software etc.
- Critical options spread out
  - “you mean you didn’t see that flag at line 500 in that file?”
    - Some duplicate options present in up to 3 configuration files and don’t get me started on the command line arguments...

MC Production Line

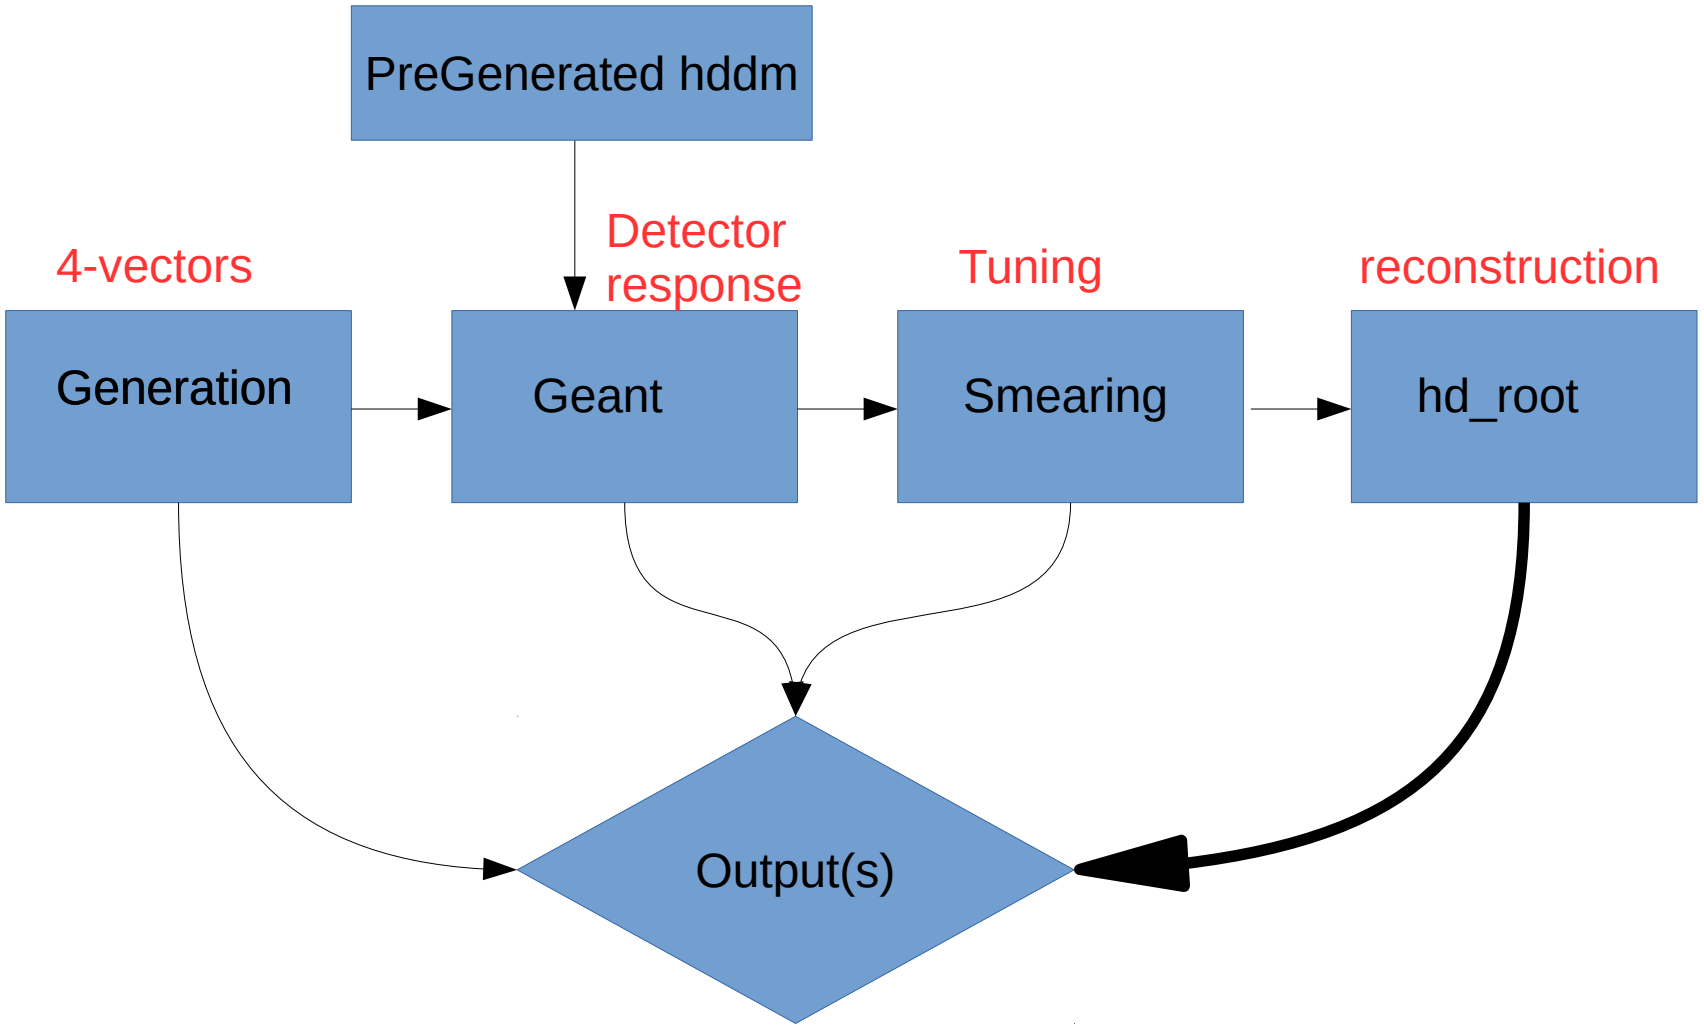

Philosophy

- MCwrapper seeks to be a one-stop-shop for simulation for GlueX/HallD. It needs to be able to:
  - Complete the production chain; from generation through (hd_root) plugins
  - Run both geant3 and geant4 easily
  - **Provide basic standards of simulation**
  - Be customizable for individual studies
  - **Utilize various batch systems**
  - Provide support for new users

And when he runs his Monte Carlo he'll probably ask for a cluster...

- Development began interactively and moved to the local farm
  - Jlab essentially has an interactive farm for users and batch nodes. Submission is handled by a program called "SWIF".
  - SWIF was quickly wrapped and MCwrapper was "born" in late 2016

- **And when he's done he'll want other clusters**

Many systems

- The GlueX collaboration is international
  - Graduate students et al have to learn how to configure their jobs properly for running both at Jlab and at their home institution
  - Often re-configuring each time something new is needed
    - Hand-me-downs are awful
- Wrapped qsub and condor and a few others thus covering most collaborators

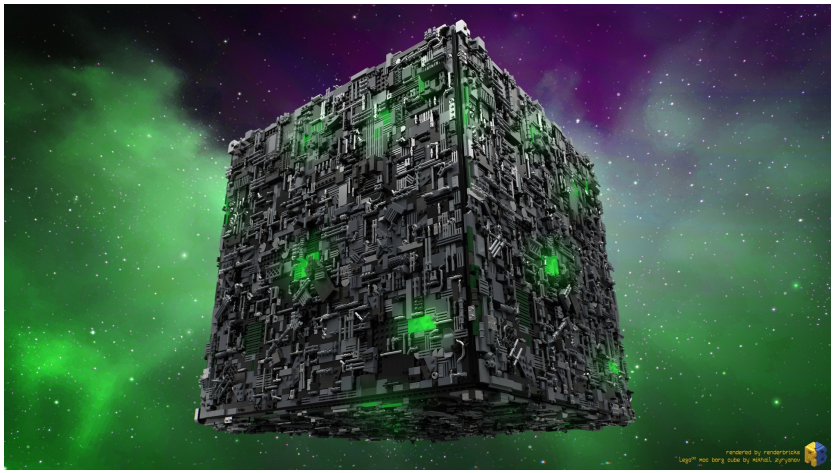

A Problem

- Jlab's farm is used by 4 experimental halls and is involved with data taking, reconstruction, and analysis
  - Moving the raw data is costly so those jobs typically run at Jlab.
- If only there was a place MC could run without interference...
  - MC has no real data to move

Murmurs of an ancient relic

- I was hired fall of 2016. While getting settled I was unaware that an OSG submit node had been set up just months prior.
  - It essentially lay dormant for two years
- Additionally, there was ongoing work with containers
  - This combined with CVMFS promised to make running elsewhere much easier

The OSG

- The OSG seemed to be the greener pasture. A place were MC can run freely.
- MC jobs are very parallelizable (many small jobs), which is exactly what the OSG likes.
- Even better. Everything was ready to be used…

Assimilation of the OSG

- Submission to the OSG was first incorporated into MCwrapper Mar. 8, 2018
- I was using the OSG almost exclusively for MC and encouraged collaborators to do the same with MCwrapper
  - Even with the tools I was often asked to run MC for others. My system of choice was the OSG

Growing pains

- It wasn't all smooth sailing….
- Some flavors of MC require taking a relatively small amount of data with the job
  - Submits of moderately sized requests (a few thousand jobs) would take *forever* (6 to 8 hours)
  - MCwrapper utilizes a database backend for book keeping. And when submitting all users were blocked from the database server from the submit node
    - Turns out, after quite a bit of investigating, the node only had a 1 gigabit connection which was saturating

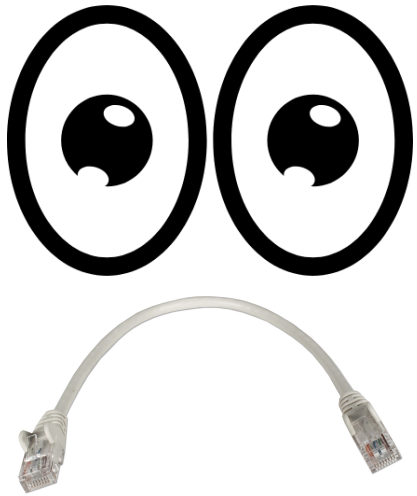

Growing pains

- Easy fix: Upgrade the NIC card to 10 Gb
- Throughput was way up. No more Denial of Service. Those same projects that took hours took ~30min to submit
- Life was good until….

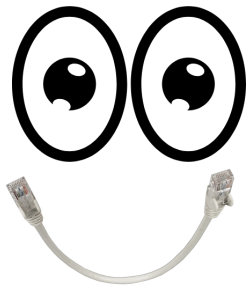

```
scosg16.jlab.org> condor_q 477165.0


-- Schedd: scosg16.jlab.org : <129.57.199.132:9615?... @ 11/20/18 14:37:55
OWNER    BATCH_NAME            SUBMITTED   DONE   RUN    IDLE  TOTAL JOB_IDS
tbritton CMD: osg-container.sh  11/20 10:43     _      1      _      1 477165.0

1 jobs; 0 completed, 0 removed, 0 idle, 1 running, 0 held, 0 suspended
scosg16.jlab.org> emacs -nw Utilities/MCOverlord.py
scosg16.jlab.org> condor_q 477165.0 -json
scosg16.jlab.org> condor_history 477165.0 -json
[
]
```

Where did the job go?!

This is only one of many condor_q inconsistencies I began to notice

Growing pains

- condor_q seemed buggy. I couldn’t get a consistent picture of the [running | idle | held] jobs
  - Speculated lag due to higher throughput?
- Started seeing jobs marked removed (in my database) that I could have sworn I never removed
  - Maybe I did and forgot….
- Why does admin see a different condor_q than I do?
  - condor_q … -json repeated

    gave different results
- One morning I asked condor_q

  about a job twice. I got different

  results; different jobs. Pretty spooky

Growing pains

- Looked into my database and saw multiple job entries with the same condor id number
  - Kurt+Thomas: “How can the same ID number be used multiple times?”
  - Edgar: “It can’t”
  - Kurt+Thomas: “It is though…”
  - Edgar: “This should never happen”

This is the abridged version of around one hundred emails sent

- Traced it back to between 10:13:01 and 10:13:44 on 11-06-2018

  coinciding with upgrade

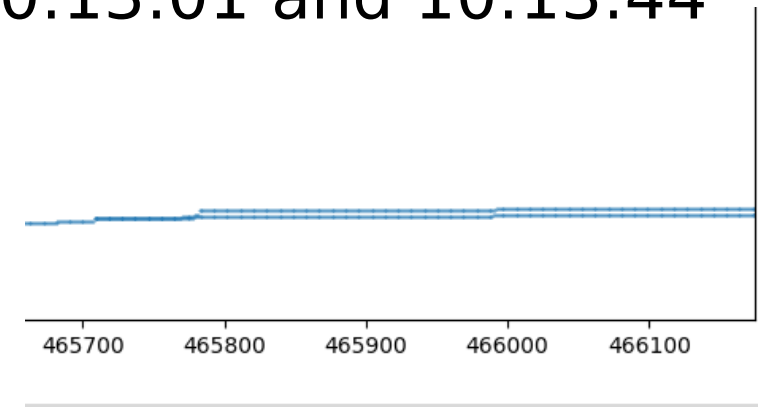

BIG thanks to Kurt and Edgar.
Twas a puzzling few weeks

The solution

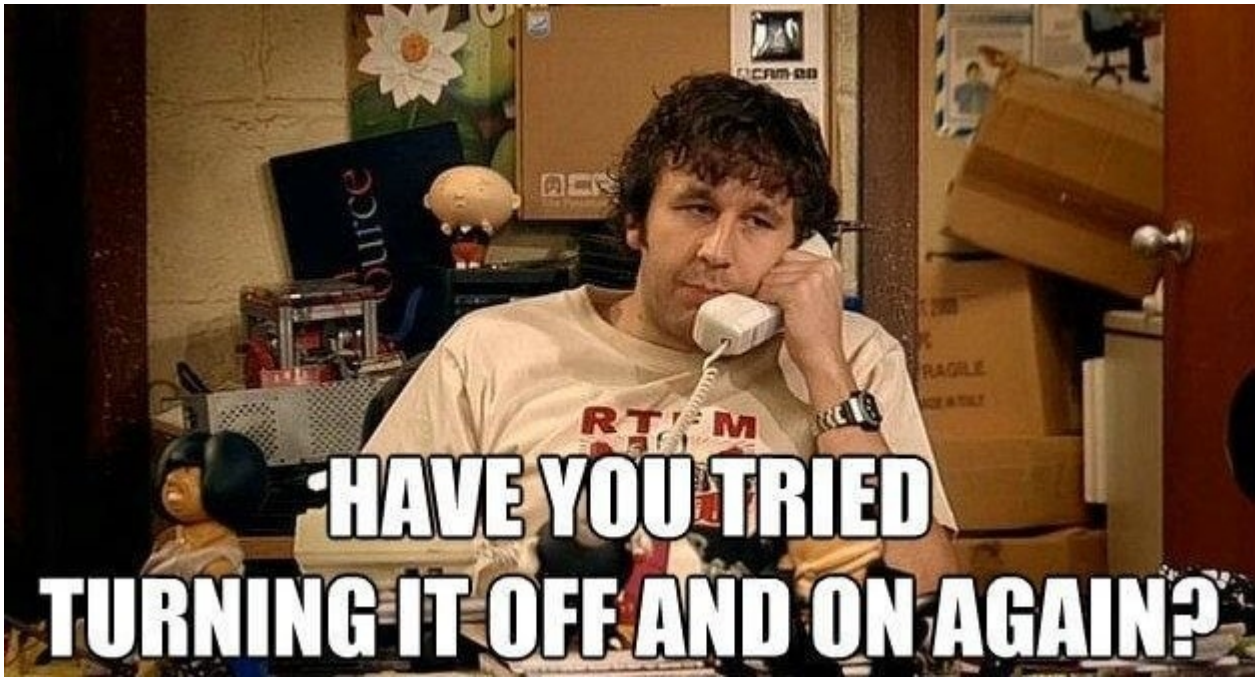

- Speculation. When the NIC card was installed something glitched and produced a second schedd. So when I would ask condor anything I would get a response from one or the other
- Mcwrapper-bot now checks for duplicate IDs and will issue an “AllStop”, notify, and shut everything down. Minimizing damage

Full operation

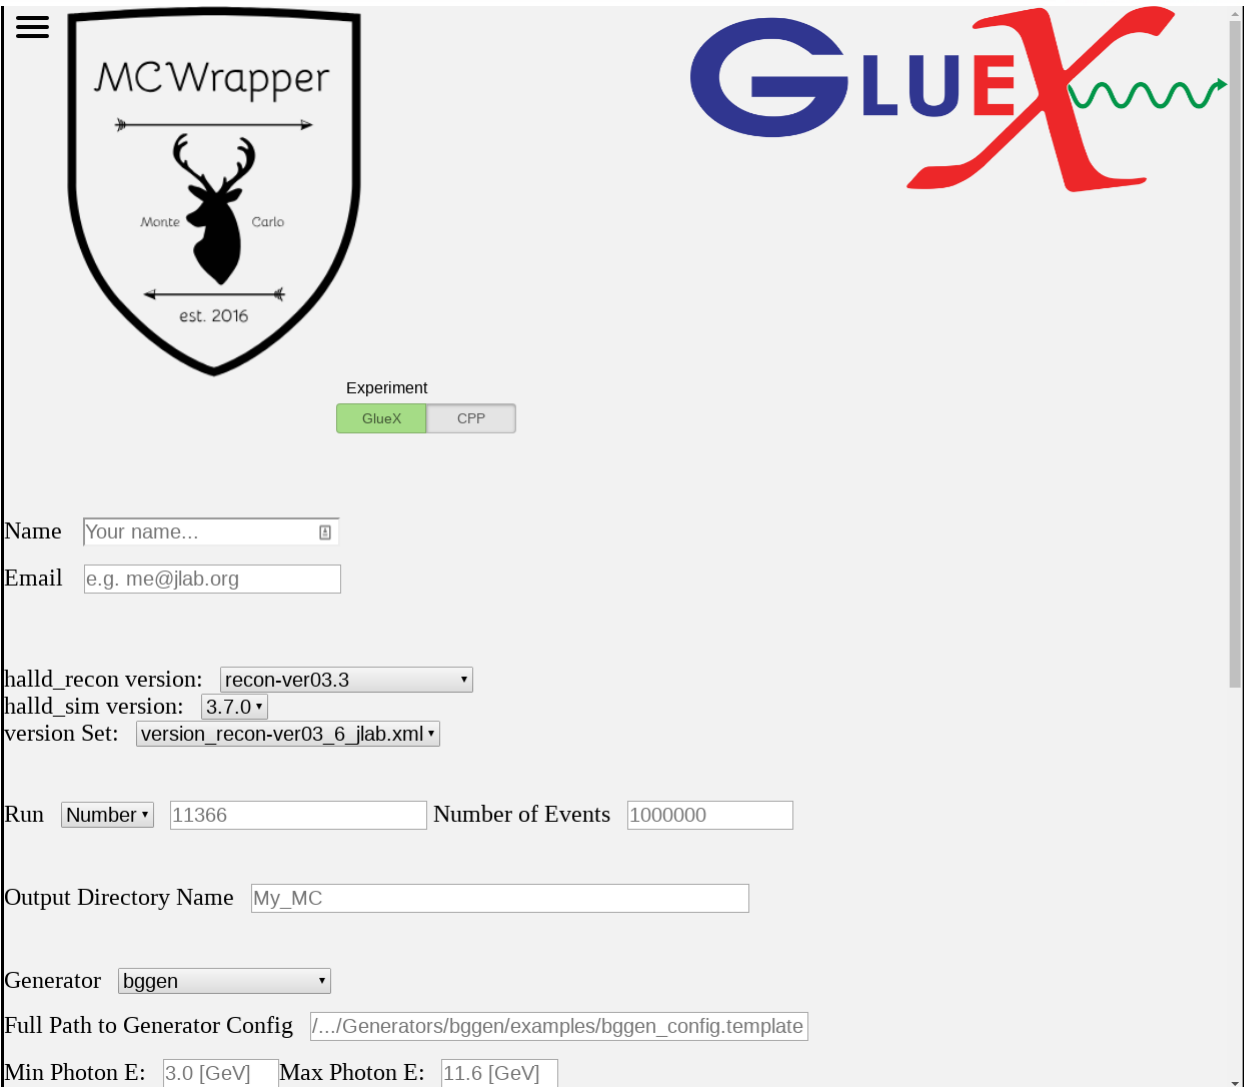

version_recon-ver03_6_jlab.xml

- **created:** 2018-12-17
- **description:** Updates of rcdb, halld_sim, hdgeant4, gluex_root_analysis, and MCwrapper built against recon-ver03.3 tag of halld_recon

Package	Version	Directory Tag
jana	0.7.9p1	ccdb166
halld_recon	recon-ver03.3	
halld_sim	3.7.0	ver03.3
hdds	recon-2017_01-ver03	
lapack	3.6.0	
cernlib	2005	
xerces-c	3.1.4	
root	6.08.06	
ccdb	1.06.06	
evio	4.4.6	
rcdb	0.03.01	
geant4	10.02.p02	
hdgeant4	1.10.0	ver03.3
hd_utilities	1.14	
gluex_MCwrapper	v2.0.4	
gluex_root_analysis	0.6	ver03.3
amptools	0.9.4	
sqlitecpp	2.2.0	bs130
sqlite	3.13.0	bs130

Full operation

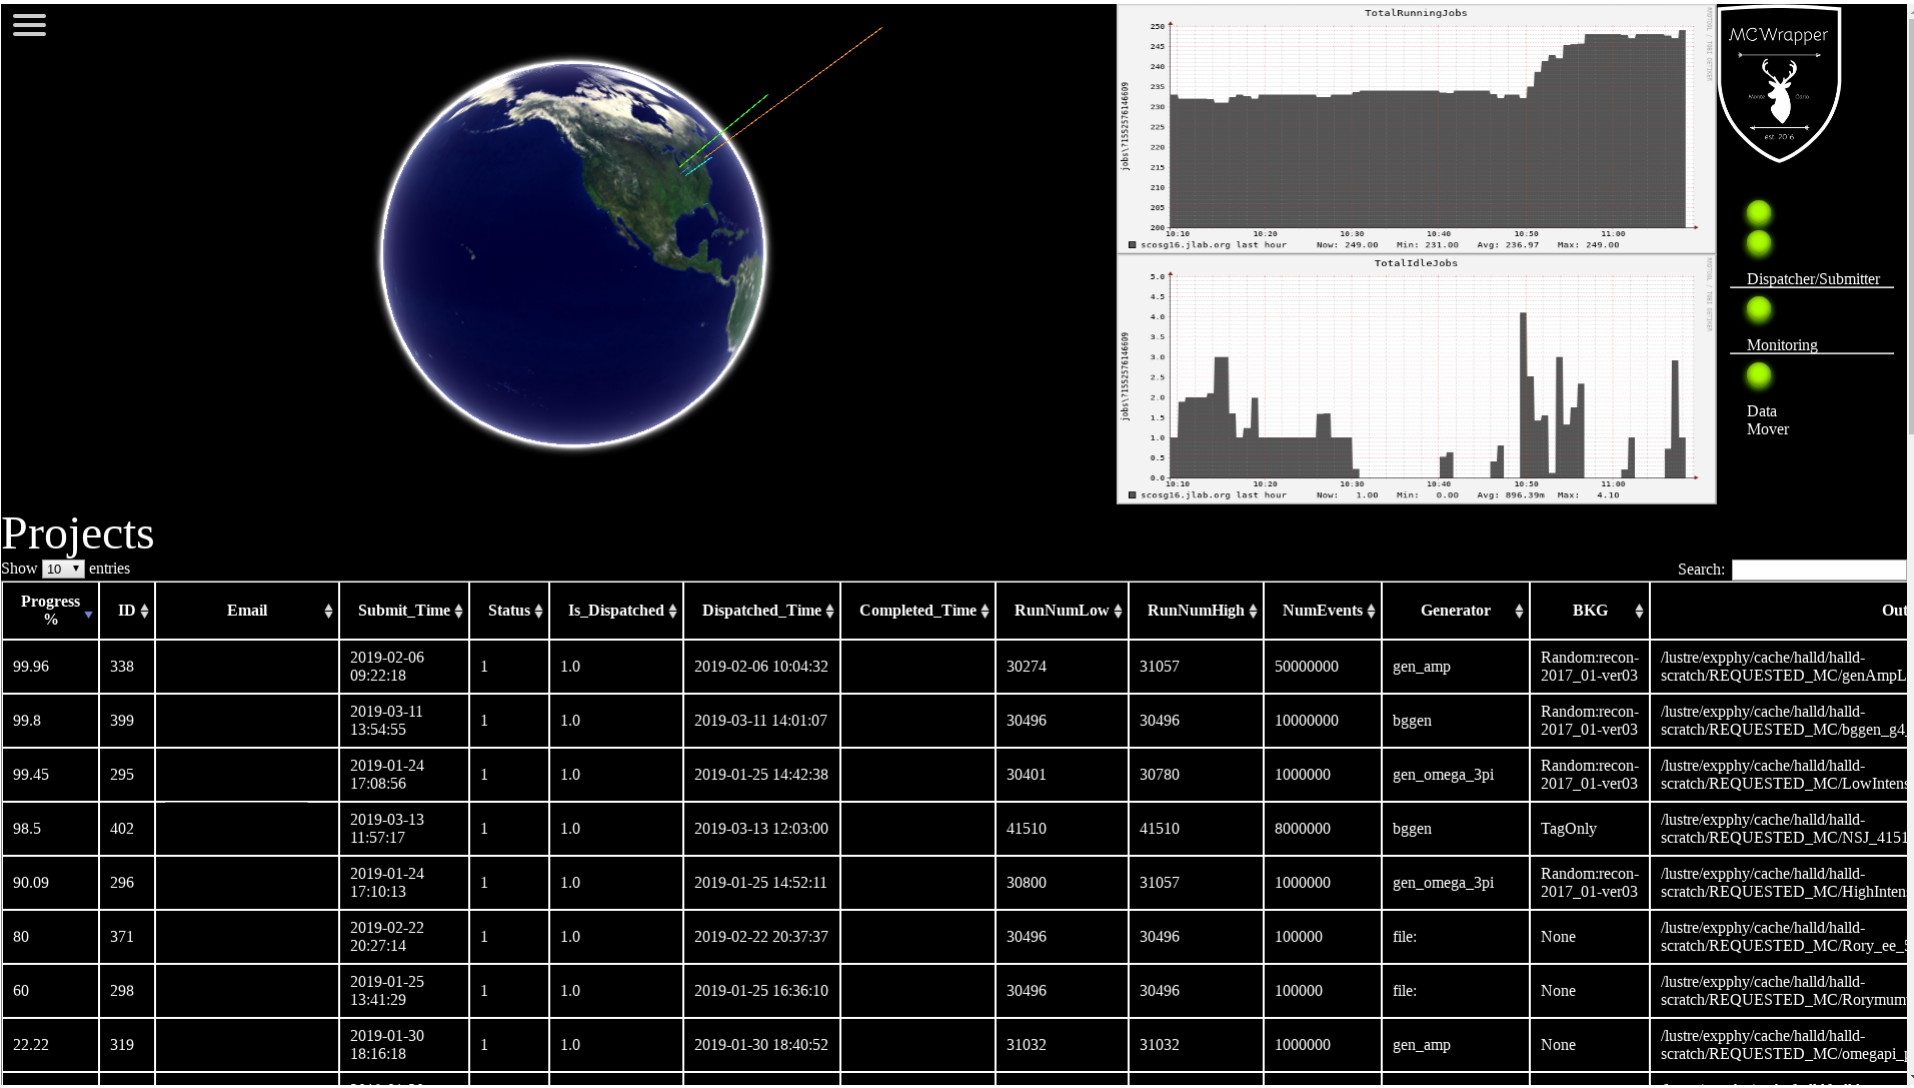

Projects

Show 10 entries Search:

Progress %	ID	Email	Submit_Time	Status	Is_Dispatched	Dispatched_Time	Completed_Time	RunNumLow	RunNumHigh	NumEvents	Generator	BKG	Out
99.96	338		2019-02-06 09:22:18	1	1.0	2019-02-06 10:04:32		30274	31057	50000000	gen_amp	Random:recon-2017_01-ver03	/lustre/expphy/cache/halld/halld-scratch/REQUESTED_MC/genAmpL
99.8	399		2019-03-11 13:54:55	1	1.0	2019-03-11 14:01:07		30496	30496	10000000	bggen	Random:recon-2017_01-ver03	/lustre/expphy/cache/halld/halld-scratch/REQUESTED_MC/bggen_g4
99.45	295		2019-01-24 17:08:56	1	1.0	2019-01-25 14:42:38		30401	30780	1000000	gen_omega_3pi	Random:recon-2017_01-ver03	/lustre/expphy/cache/halld/halld-scratch/REQUESTED_MC/LowInten
98.5	402		2019-03-13 11:57:17	1	1.0	2019-03-13 12:03:00		41510	41510	8000000	bggen	TagOnly	/lustre/expphy/cache/halld/halld-scratch/REQUESTED_MC/NSJ_4151
90.09	296		2019-01-24 17:10:13	1	1.0	2019-01-25 14:52:11		30800	31057	1000000	gen_omega_3pi	Random:recon-2017_01-ver03	/lustre/expphy/cache/halld/halld-scratch/REQUESTED_MC/HighInten
80	371		2019-02-22 20:27:14	1	1.0	2019-02-22 20:37:37		30496	30496	100000	file:	None	/lustre/expphy/cache/halld/halld-scratch/REQUESTED_MC/Rory_ee_
60	298		2019-01-25 13:41:29	1	1.0	2019-01-25 16:36:10		30496	30496	100000	file:	None	/lustre/expphy/cache/halld/halld-scratch/REQUESTED_MC/Rorymum
22.22	319		2019-01-30 18:16:18	1	1.0	2019-01-30 18:40:52		31032	31032	1000000	gen_amp	None	/lustre/expphy/cache/halld/halld-scratch/REQUESTED_MC/omegapi_

Future

- Increase throughput
  - Htcondor
  - Data streaming
    - Xrootd? Stashcache?
  - More submit nodes working in tandem
- Assimilate more resources
  - Transparent to the user

Conclusion

- If you give a GlueX researcher a computer
- He'll need to run Monte Carlo
- And when he runs Monte Carlo
- He'll ask for a cluster
- When you give him a cluster
- He'll want more clusters
- And when he has more clusters
- He will be asked to run everyone's Monte Carlo
- And when he runs everyone's Monte Carlo
- He will be forced to fight for resources
- And when he is forced to fight for resources
- He will ask for the OSG