SOLID weekly meeting, April 16, 2019

GlueX Experience with the Open Science Grid

past experience, present challenges, future prospects

Richard Jones, University of Connecticut

This work is supported by the U.S. National Science Foundation under grant 1812415

GlueX event, data rates

Experimental data rates

- 80 kHz trigger rate (~1.6GB/s)
- many exclusive final states
- < 50% are simple topologies

Simulated data estimates

- based on *GlueX-doc-3813* (2018)
- **36 Mcore-hr/yr** (2020 and beyond)
- *primarily targeted for OSG*

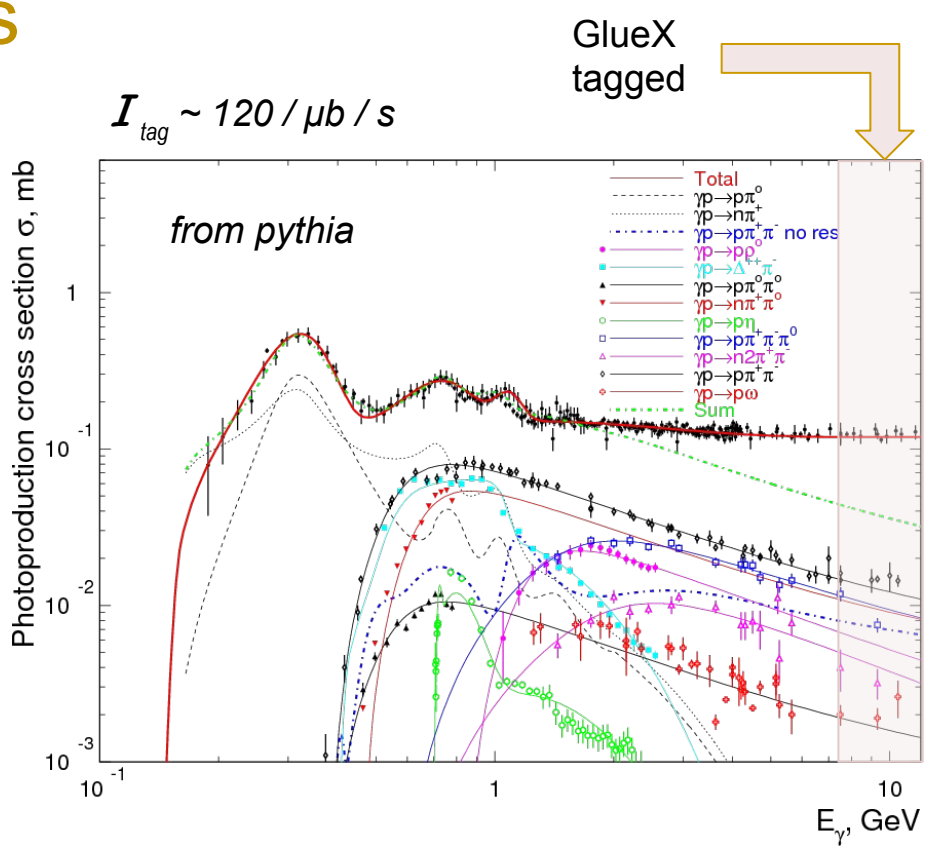

This work is supported by the National Science Foundation under grant 1508238

GlueX event, data rates

Experimental data rates

- 80 kHz trigger rate (~1.6GB/s)
- many exclusive final states
- < 50% are simple topologies

Real data calibration + monit.

- based on *GlueX-doc-3813* (2018)
- **35 Mcore-hr/yr** (2020 and beyond)
- *primarily targeted for JLab farm*

GlueX tagged

I_{tag} ~ 120 / μb / s

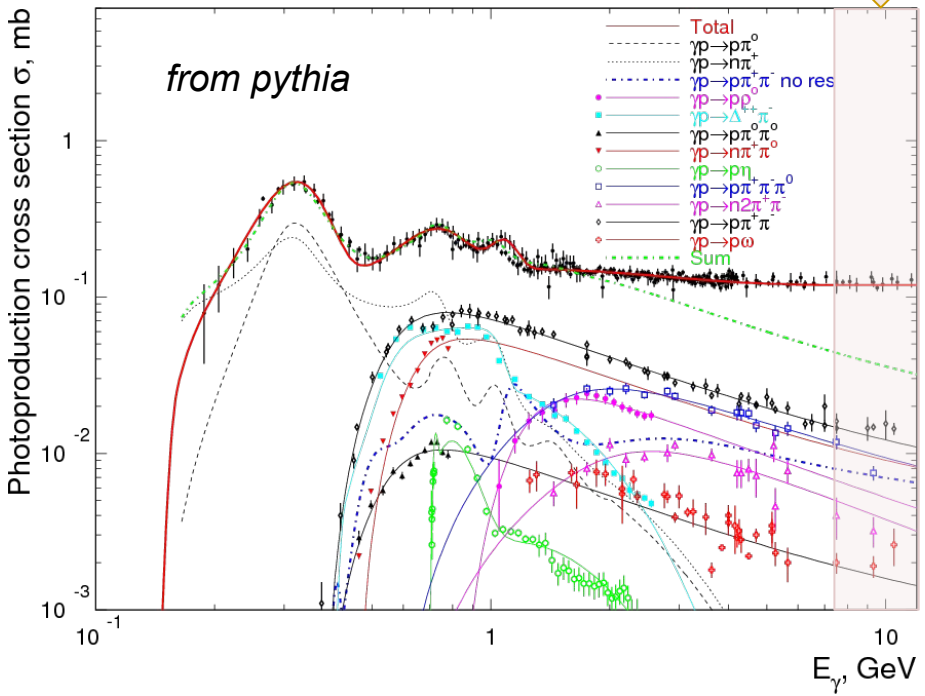

This work is supported by the National Science Foundation under grant 1508238

GlueX event, data rates

Experimental data rates

- 80 kHz trigger rate (~1.6GB/s)
- many exclusive final states
- < 50% are simple topologies

Real data reconstruction

- based on *GlueX-doc-3813* (2018)
- **130 Mcore-hr/yr** (2020 and beyond)
- *targeted for NERSC + ? + OSG??*

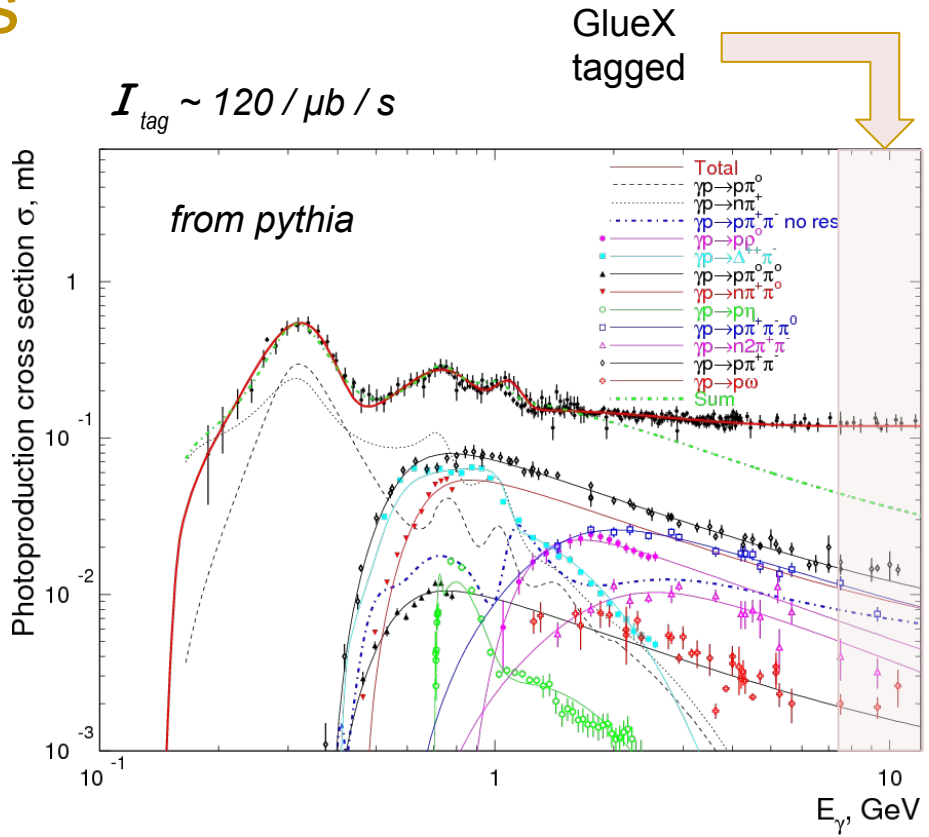

This work is supported by the National Science Foundation under grant 1508238

OSG: *the Gluex VO*

Open Science Grid

- Gluex vo created in 2009 (10 years!)
- *NSF Physics at the Information Frontier*
- underwent an early series of data challenges
  - *data challenge I* - 2 Mcore-hr, 10 days
  - *data challenge II* - 6 Mcore-hr, 30 days
  - ***campus cluster campaign*** - 1 Mcore-hr, 15 days
- reconstruction is maturing, analysis is ramping up...

This work is supported by the National Science Foundation under grant 1508238

Data Challenge 1: *results*

Open Science Grid

- cpu availability was very high (>10,000 cores peak)
- production efficiency was not great (40 – 60%)
- part of inefficiency is due to pre-emption (opportunistic)
- understanding sources of inefficiency is reason why we stopped @5B events

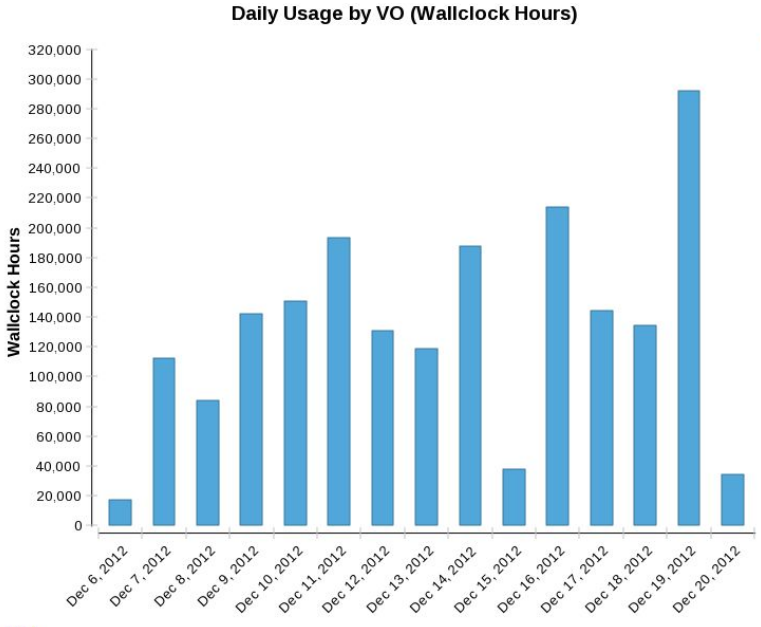

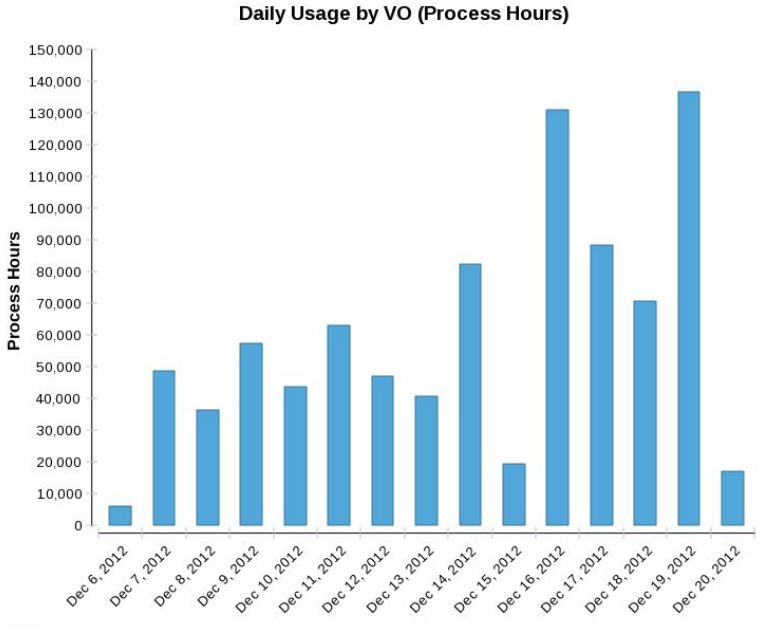

Data Challenge 2: *results*

Open Science Grid

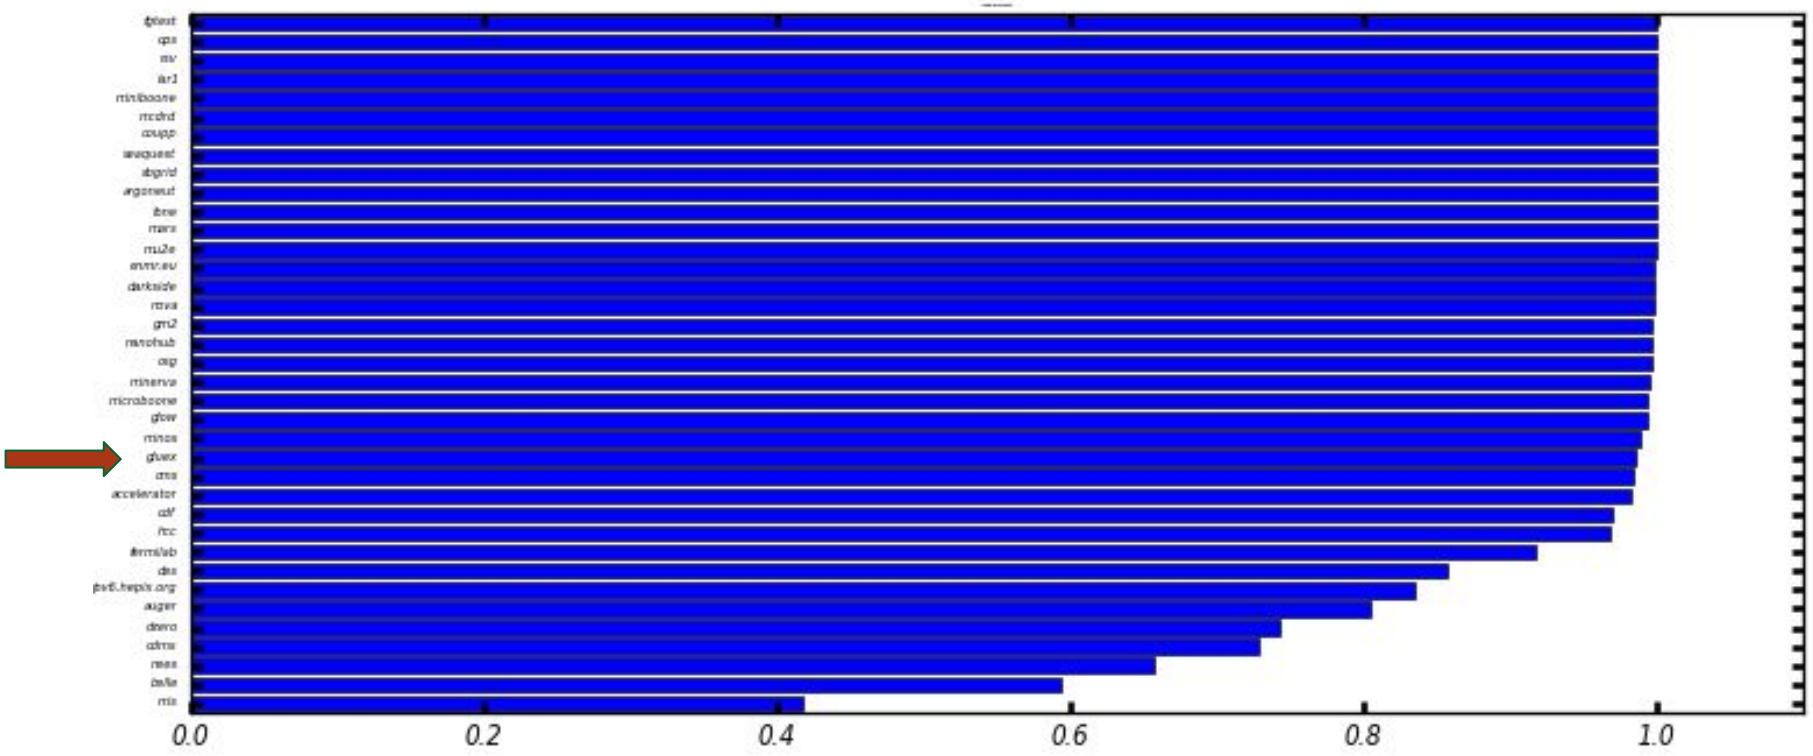

cpu time / wall clock time, daily average

This work is supported by the National Science Foundation under grant 1508238

Evolution in methodology

1. OSG_APPS, OSG_DATA ⟶ /cvmfs/oasis.opensciencegrid.org
2. singularity containers ⟶ /cvmfs/singularity.opensciencegrid.org

Big gains in opportunistic throughput seen by adapting software to run on the widest possible range of platforms.

For Gluex, this was a iterative, labor-intensive, experts-only process until ...

All Gluex jobs containerized, can even run on sites without singularity installed.

This work is supported by the National Science Foundation under grant 1508238

Evolution in methodology (2)

Open Science Grid

1. Nightly builds inside standard container oasis updates (as needed)
2. Software release management using github tags + *versions.xml*
3. Container hosted on cvmfs, rarely needs updating (once per year?)
4. Multiple binary releases maintained on oasis
   a. which one to use is selected at runtime
   b. current oasis footprint on the high side - 150 GB
   c. best practice to follow in this area? *-- got no push-back at OSG mtg*
   d. github-based updates are automatic, no human intervention needed

This work is supported by the National Science Foundation under grant 1508238

Evolution in *resources*

Open Science Grid

GlueX off-site computing resource needs *(GlueX-doc-3813)*

1. 130 Mcore-hr/yr - experimental data reconstruction
   - *Jefferson Lab compute facility (total 70 Mcore-hr/yr, all experiments)*
   - *NERSC (proven option, but competitive)*
   - *other (includes OSG, active study)?*
2. 36 Mcore-hr/yr - Monte Carlo simulation
   - *primarily targeted for OSG*
   - *cannot live on donations alone*

This work is supported by the National Science Foundation under grant 1508238

Evolution in *resources*

Open Science Grid

Existing OSG resources for GlueX:

1. **UConn_OSG** site: 600-core cluster
   - active on OSG since ca. 2010
   - contributed 3-4 Mhr/yr opportunistic OSG cycles over past decade
2. **GLUEX_US_FSU_HNPGRID** site: “entry-level” cluster
   - active on OSG since ca. 2017
   - contributed 100 khr/yr to OSG over the past 2 years
   - starting point for future growth in GlueX computing at FSU

*This amounts to **10%** of the projected need for GlueX simulations over past 2019.*

This work is supported by the National Science Foundation under grant 1508238

Evolution in *resources*

Open Science Grid

Where are the OSG resources for GlueX?

A number of *resourced GlueX institutions* have offered to contribute:

a. Carnegie Mellon University
b. Indiana University - stanley, karst, BigRed
c. Florida State University - RACF
d. George Washington University - colonialone
e. College of William and Mary - sciclone
f. University of Regina - computecanada
g. UConn Health Center HPC - xanadu
h. UConn Storrs HPC - storrs-hpc

This work is supported by the National Science Foundation under grant 1508238

Evolution in *resources*

Open Science Grid

Where are the OSG resources for GlueX?

A number of *resourced GlueX institutions* have offered to contribute, *but the barriers turned out to be higher than expected -- 2017 summer of setbacks.*

Summer 2018 -- what can be done to move forward?

- for the time being, *skip OSG site integration*
- **implement a separate stand-alone condor pool (at UConn)**
- get access to individual user accounts on every member's cluster
- **customize a glidein** for each individual cluster (bosco, 9 in total)
- install local copy of complete GlueX stack + container
- diagnose, debug, optimize...

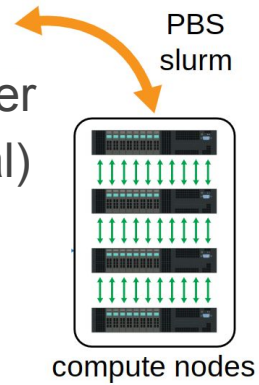

This work is supported by the National Science Foundation under grant 1508238

Evolution in *resources*

Open Science Grid

Summer 2018 -- lessons from "bosco" exercise

1. It was successful
   - 1 Mcore-hr of simulations completed in 15 days
   - average 5k cores active during periods when not debugging
   - spanned very different types: included BigRed Cray HPC @ IU
2. It was a management headache
   - like going back to opportunistic OSG ca. 2012 *(rsync, pilot jobs, etc.)*
   - no two clusters the same, each one like a green field
   - ***whack-a-mole ops*** -- not sustainable for GlueX production long-term.

This work is supported by the National Science Foundation under grant 1508238

Evolution in *resources*

Open Science Grid

Where are the OSG resources for GlueX?

A number of *resourced GlueX institutions* have offered to contribute, *and through the summer 2018 bosco exercise all have done so, except one.*

BUT

- There are good reasons to go the route of full OSG site integration.
- This requires significant effort from individual site admins, but...
- *... it also requires buy-in from management -- a very hard sell.*
- More than *technical expertise* is needed *(goodwill, friends in high places, ...)*
- Several early efforts were derailed at this point.

This work is supported by the National Science Foundation under grant 1508238

Evolution in *resources*

Open Science Grid

Where are the OSG resources for GlueX?

Broader lessons from the GlueX bosco exercise:

1. *Private* cluster resources owned by individual groups are not keeping pace with the needs of our science, trend is not likely to reverse.
2. Growth is happening in *shared computing resources* at universities.
3. Hurdles to executing grid jobs there are *primarily administrative, not technical.*
4. *In-advance discussions, agreements* with the central IT managers of these resources are needed -- they can be very helpful ***or not.***

This work is supported by the National Science Foundation under grant 1508238

Evolution in *resources*

Open Science Grid

Example framework for successful discussions:

1. GlueX researcher Prof Zisis Papandreou and his students would like to contribute resources on Compute Canada toward GlueX simulations. GlueX is a multi-national scientific collaboration based around the GlueX experiment at Jefferson Lab in Newport News, Virginia.

2. GlueX simulations are needed by and benefit the entire collaboration, not individual researchers or groups. As such, they are a shared responsibility of all groups. All groups are being asked to contribute a share toward the total anticipated load of 36 Mcore-hr per year. Currently 9 universities have expressed willingness to contribute, including Univ. of Regina and my own Univ. of Connecticut.

This work is supported by the National Science Foundation under grant 1508238

Evolution in *resources*

Open Science Grid

Example framework for successful discussions:

3. To do central coordination of all of these resources, GlueX uses a central simulations production system called "glidein workflow management system" (glideinWMS) that is supported by the Open Science Grid and provided to us as a service. GlueX is an authorized virtual organization in OSG, on the same footing as LIGO, USATLAS, and USCMS.

4. This glideinWMS job factory interacts seamlessly with slurm, and has been tested and shown to work on slurm clusters at the University of Connecticut.

This work is supported by the National Science Foundation under grant 1508238

Evolution in *resources*

Open Science Grid

Example framework for successful discussions:

5. We would like permission to run GlueX glideins on WestGrid resources. To do that, all we would need is that a single GlueX group account be created on the WestGrid submit host and regular ssh access from the OSG glidein factory to the slurm head node for starting and managing our GlueX simulation jobs.

6. Responsibility for security of the account would begin and end with Zisis, but it has the full OSG security apparatus behind it. I (RJ) can explain more about that when we speak.

This work is supported by the National Science Foundation under grant 1508238

Evolution in *resources*

Example framework for successful discussions:

5. Batch job environment: Westgrid slurm with queues, priorities, and job policies set by you -- that we will respect -- nothing more.
6. Access to /cvmfs on the workers is desired but not required. If no /cvmfs is present, we request ~150 GB of scratch space (can be read-only on worker nodes) for staging software, databases, etc.
7. Use of singularity on the workers is desired but not required.

This work is supported by the National Science Foundation under grant 1508238

Evolution in *resources*

Open Science Grid

Where are the OSG resources for GlueX?

So what is the plan going forward?

1. Discussions have begun with some of our central campus IT admins (UConn, Compute Canada, CMU, IU)
2. Idea is for project account on the campus cluster head node (separate from individual user accounts) configured for connection from the OSG glidein factory over ssh to start OSG jobs on the local batch system (eg. slurm, pbs, condor).
3. The jobs that start up are miniature condor workers -- connect to the OSG Connect condor pool and pull jobs according to preset rules.

This work is supported by the National Science Foundation under grant 1508238

Evolution in *resources*

Open Science Grid

Where are the OSG resources for GlueX?

So what is the plan going forward?

3. The jobs that start up are miniature condor workers -- connect to the OSG Connect condor pool and pull jobs according to preset rules.
4. Things like number of cores/job, memory/job, disk/job, and total lifetime of the glidein are all configured by the admin on per-cluster basis.
5. The jobs are matched to resources that meet their requirements.
6. *Jobs with the least restrictive requirements get the most resources.*

This work is supported by the National Science Foundation under grant 1508238

Evolution in *resources*

Open Science Grid

Where are the OSG resources for GlueX?

It really helps if we bring something to the table!

1. Offer access to OSG opportunistic cycles to any university researcher who is willing to learn how to use it.
2. Take advantage of NSF funding opportunities to grow university-based OSG resources.
3. Push the growth boundary along the data storage axis.

This work is supported by the National Science Foundation under grant 1508238

Evolution in *resources*

Where are the OSG resources for GlueX?

National Science Foundation: Campus Computing and the Computing Continuum

NSF 19-553 sollicitation: *"Local campus computing resources have emerged as an important aggregated and shared layer of scientific computing, as evidenced by the growth in* ***Open Science Grid (an NSF-funded distributed scientific computing fabric of shared computing clusters across more than 100 institutions)*** *productivity that will approach two billion CPU hours delivered in scientific computing for the calendar year 2018."*

This work is supported by the National Science Foundation under grant 1508238

Evolution in *resources*

Where are the OSG resources for GlueX?

University of Connecticut proposal 1925716

- submitted February 20, 2019
- $400,000 for compute nodes (1020 cores) + storage (1 PB)
- enables a broad range of science at UConn
  - *experimental nuclear physics*
  - *geophysics*
  - *astrophysics*
  - *public health*

This work is supported by the National Science Foundation under grant 1508238

Evolution in *resources*

Open Science Grid

Where are the OSG resources for GlueX?

Jefferson Lab computing review, November 2018:

- UConn - 10M core hours
- FSU - 5M core hours (so far, more on the horizon)
- Northwestern - 2M core hours
- Regina - 2M core hours (so far, asking for more)
- Indiana - 4M core hours
- Florida International - 2M core hours
- George Washington - 5M core hours (rough estimate)
- College of William and Mary - 2M core hours
- opportunistic cycles - 10M core hours

This work is supported by the National Science Foundation under grant 1508238

Evolution: *beyond simulation?*

Open Science Grid

GlueX offline computing resource needs *(GlueX-doc-3813)*

1. 130 Mcore-hr/yr - experimental data reconstruction
   - *Jefferson Lab compute facility (total 70 Mcore-hr/yr, all experiments)*
   - *NERSC (proven option, but competitive)*
   - *other ??*

In the future, maybe OSG can contribute to the greater need here

- This is intrinsically a **HTC problem**
- To solve it we are *looking primarily to HPC resources* (technical reasons)
- These problems should be readily solvable (UConn working with WCHTC)

This work is supported by the National Science Foundation under grant 1508238

Evolution *beyond GlueX*

Open Science Grid

Also from the November, 2018 Jefferson Lab computing review:

Plan to provision for simulation offsite ~63 M core hours per year

- OSG – pioneered by GLUEX, we will follow.
  - Submit jobs from JLab OSG submit node.
  - Collaborators contribute to OSG –in particular MIT.
- Other
  - Submit simulation in Docker container locally at the remote site using remote site's staff and batch system. Return results to Jlab.
  - INFN and others interested in this.
- NERSC
  - Submit via SWIF2 workflow tool (follow GLUEX).
  - Have requested NERSC allocation of 30M core hours.
  - Enough to cover 50% of the annual simulation workload.

This work is supported by the National Science Foundation under grant 1508238

Backup slides

This work is supported by the National Science Foundation under grant 1508238

History: *slide from Oct. 2012, rtj*

- Experiment is in *construction phase until 2014*
- **Usage increasing with demand for Monte Carlo**

run period	usage
9/2009 – 9/2010	26.4 khr
9/2010 – 9/2011	1.1 Mhr
9/2011 – present	**2.1 Mhr**

- **Growth has slowed as work turns to digesting the results**
- **Task:** simulation of background QCD photoproduction (Pythia)
- **Purpose:** develop cuts to suppress background, measure leakage from minimum-bias events into signal sample after cuts, requires very large statistics MC samples, shared between analysis tasks.
- **Plans**: saturate at the level 5-10M core-hr/yr until physics data collection begins ca. 2015.
- **Strategy**: glideinWMS – support from OSG admins *outstanding !*

This work is supported by the National Science Foundation under grant 1508238

Data Challenge 1: Dec. 2012

Open Science Grid

Purpose of the exercise:

1. **Test** the current simulation and reconstruction tools
   - **bggen** – pythia-based background Monte Carlo generator
   - **hdgeant** – geant3-based physics simulation, base detector
   - **mcsmear** – detector efficiency and resolution models
   - **hd-ana** – reconstruction of tracks, neutrals
   - **REST** plugin – summary of reconstruction results

2. **Develop** the ability to manage simulation production and data storage at rates approaching GlueX Phase I.

3. **Produce** a large sample of background simulation data.

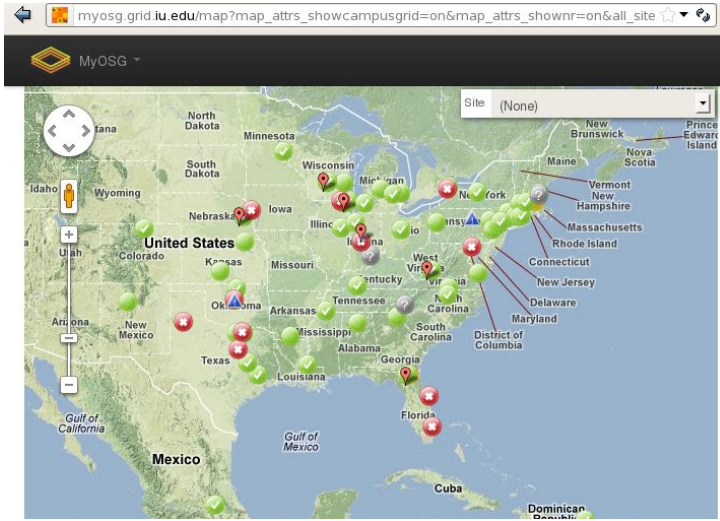

initial goal: 10 billion events, **60 days** at startup intensity

This work is supported by the National Science Foundation under grant 1508238

Data Challenge 1: *results*

Open Science Grid

- total of 5.56B events simulated
  - **4.24B** on the OSG
  - **0.96B** at Jefferson Lab
  - **0.36B** at CMU
- completed over a period of 14 days

Ran into several limiting factors:

1. security event
2. software staging
3. freeze-ups in hd-ana
4. memory hogging in hd-ana
5. segfaults in hdgeant
6. irreproducibility in mcsmear

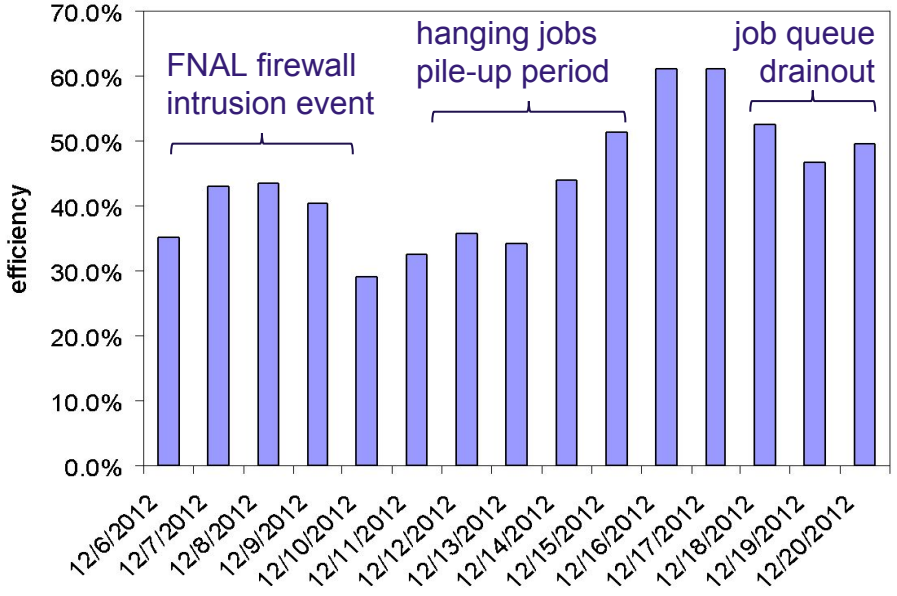

This work is supported by the National Science Foundation under grant 1508238

Data Challenge 2: Apr. 5-24, 2014

Similar in purpose to DC1:

1. **Test** the current simulation and reconstruction tools, see if we fixed problems from DC1, check for new ones.

2. **Develop** the ability to manage production and data storage at rates approaching GlueX Phase I.

3. **Produce** a large sample of background simulation data, sufficient statistics to address issues.

This work is supported by the National Science Foundation under grant 1508238

Data Challenge 2: Apr. 5-24, 2014

Similar in purpose to DC1:

1. **Test** the current simulation and reconstruction tools, see if we fixed problems from DC1, check for new ones.
    - more realistic simulation
    - include electromagnetic background
    - improved reconstruction
2. **Develop** the ability to manage production and data storage at rates approaching GlueX Phase I.
3. **Produce** a large sample of background simulation data, sufficient statistics to address issues.

This work is supported by the National Science Foundation under grant 1508238

Data Challenge 2: Apr. 5-24, 2014

Open Science Grid

Similar in purpose to DC1:

1. **Test** the current simulation and reconstruction tools, see if we fixed problems from DC1, check for new ones.
   - more realistic simulation
   - include electromagnetic background
   - improved reconstruction
2. **Develop** the ability to manage production and data storage at rates approaching GlueX Phase I.
   - software distribution using cervnvm / oasis
   - particular focus on job efficiency
3. **Produce** a large sample of background simulation data, sufficient statistics to address issues.

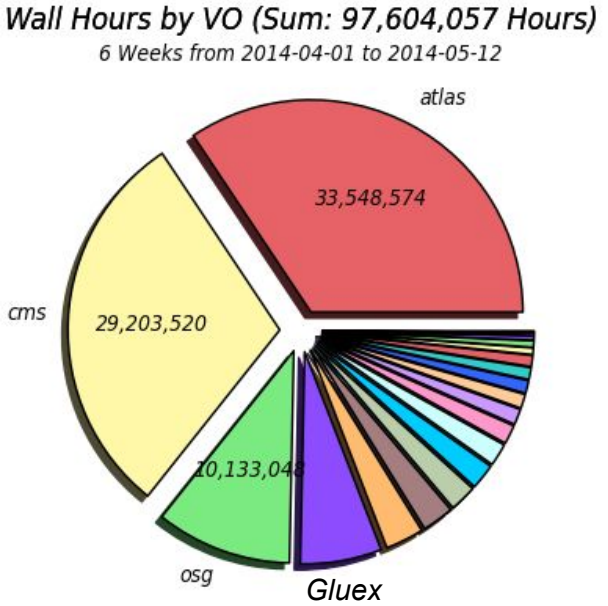

6M core-hr = DC1 x 2

This work is supported by the National Science Foundation under grant 1508238

Data Challenge 2: *results*

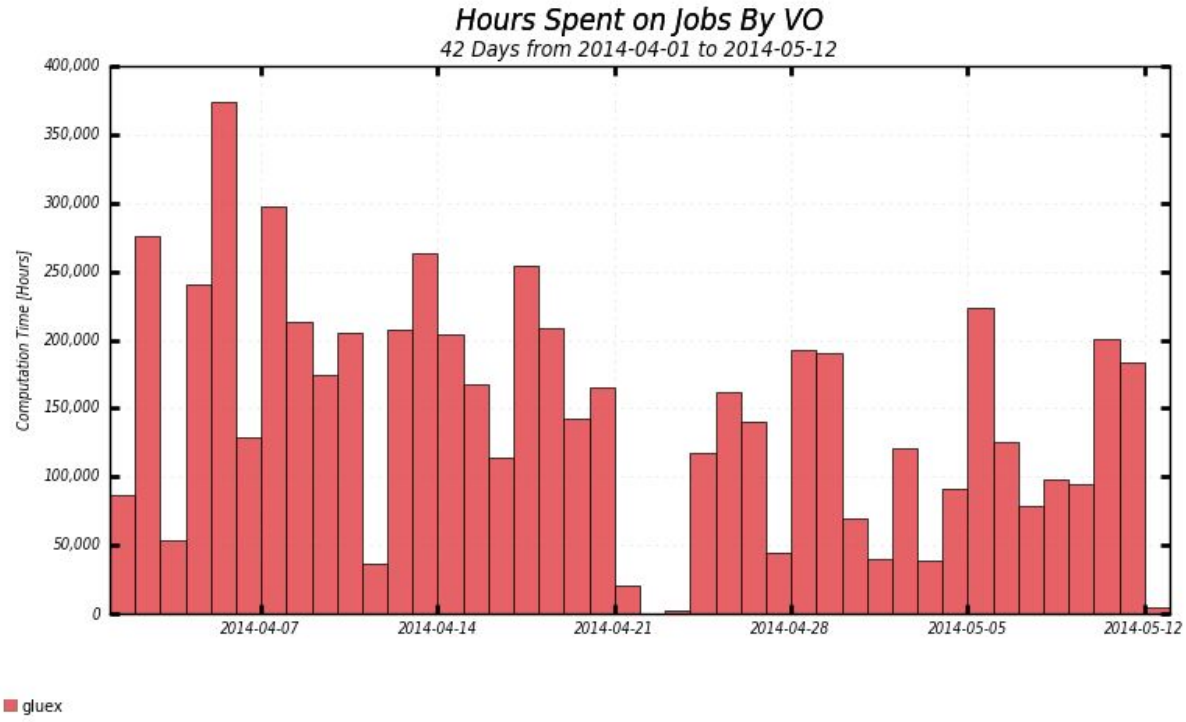

Final event tally

CMU	170M	2%
MIT	760M	9%
JLAB	2000M	25%
OSG	5200M	64%
total	**8100M**	**100%**

This work is supported by the National Science Foundation under grant 1508238

Data Challenge 2: *results*

Open Science Grid

Gluex usage on the Fermilab site

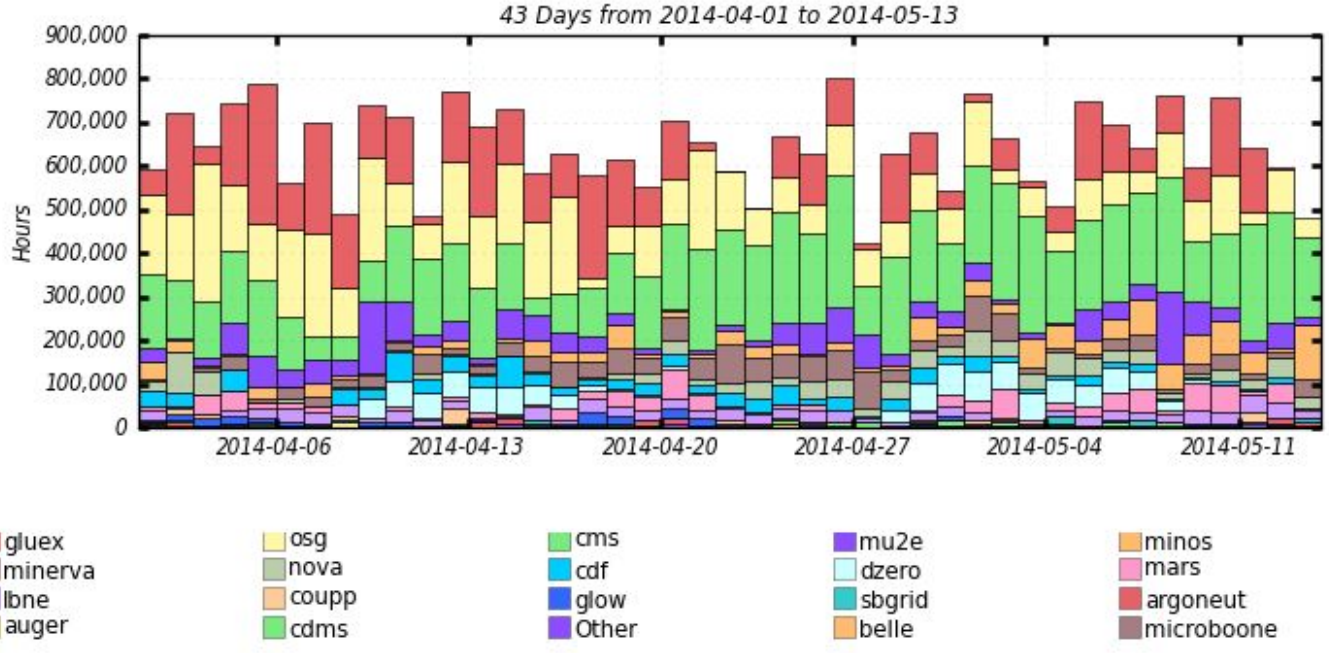

Data Challenge 2: *results*

Open Science Grid

Gluex usage on the Purdue site

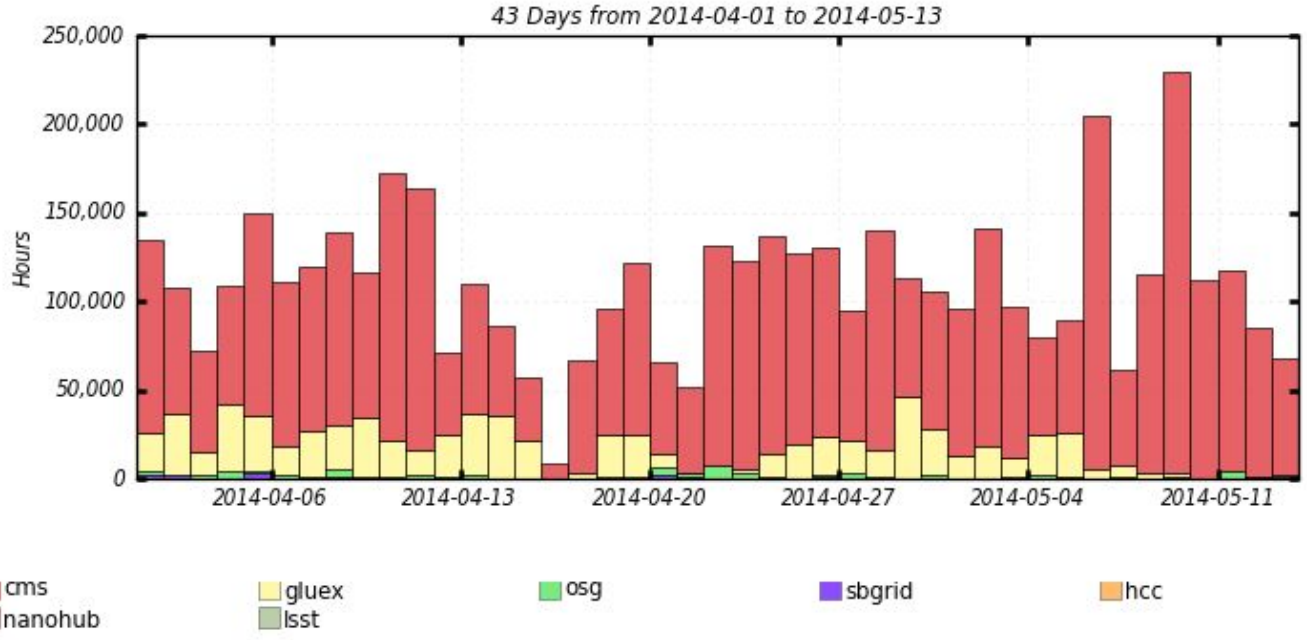

Data Challenge 2: *results*

Open Science Grid

Gluex usage on the Northwestern site

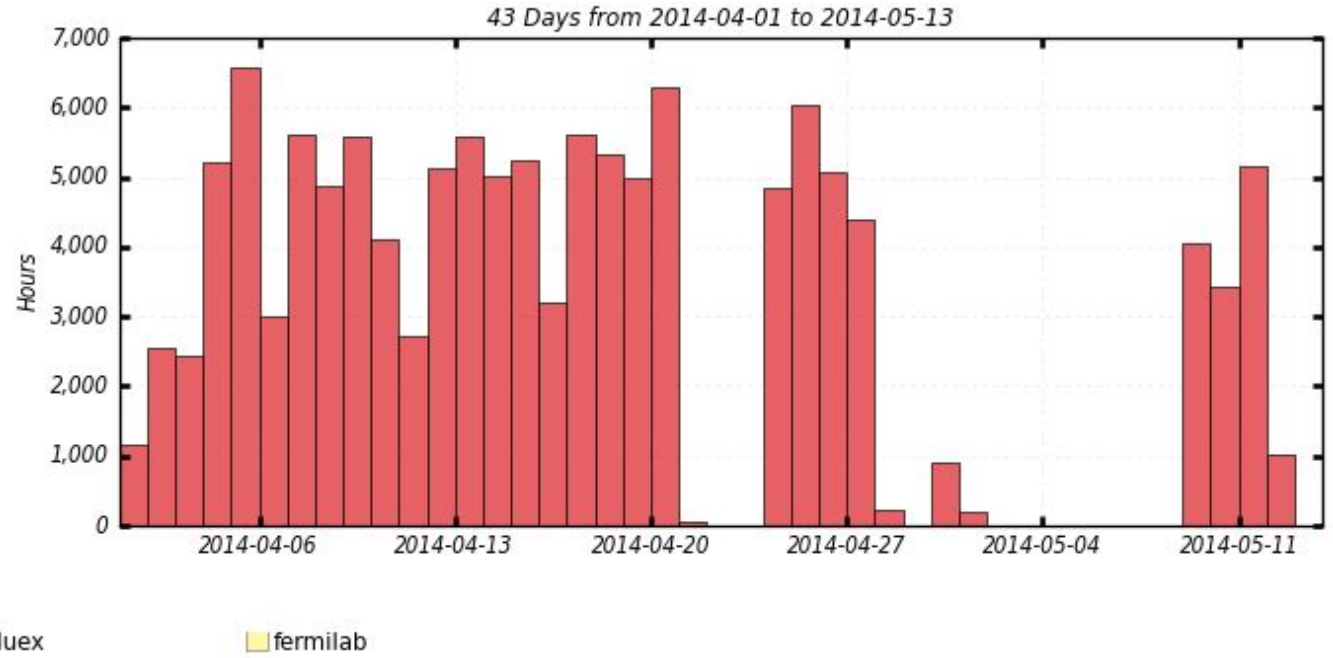

Data Challenge 2: *results*

Open Science Grid

Gluex usage on the UConn site

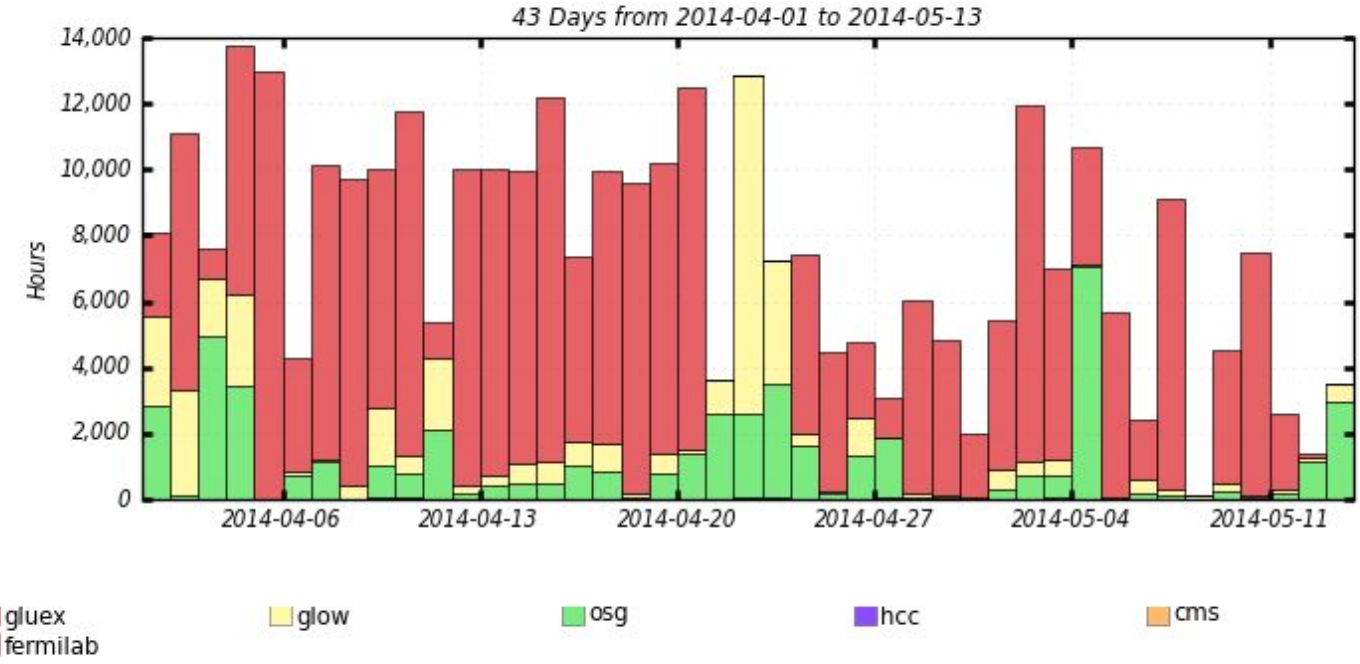

Data Challenge 2: *results*

Open Science Grid

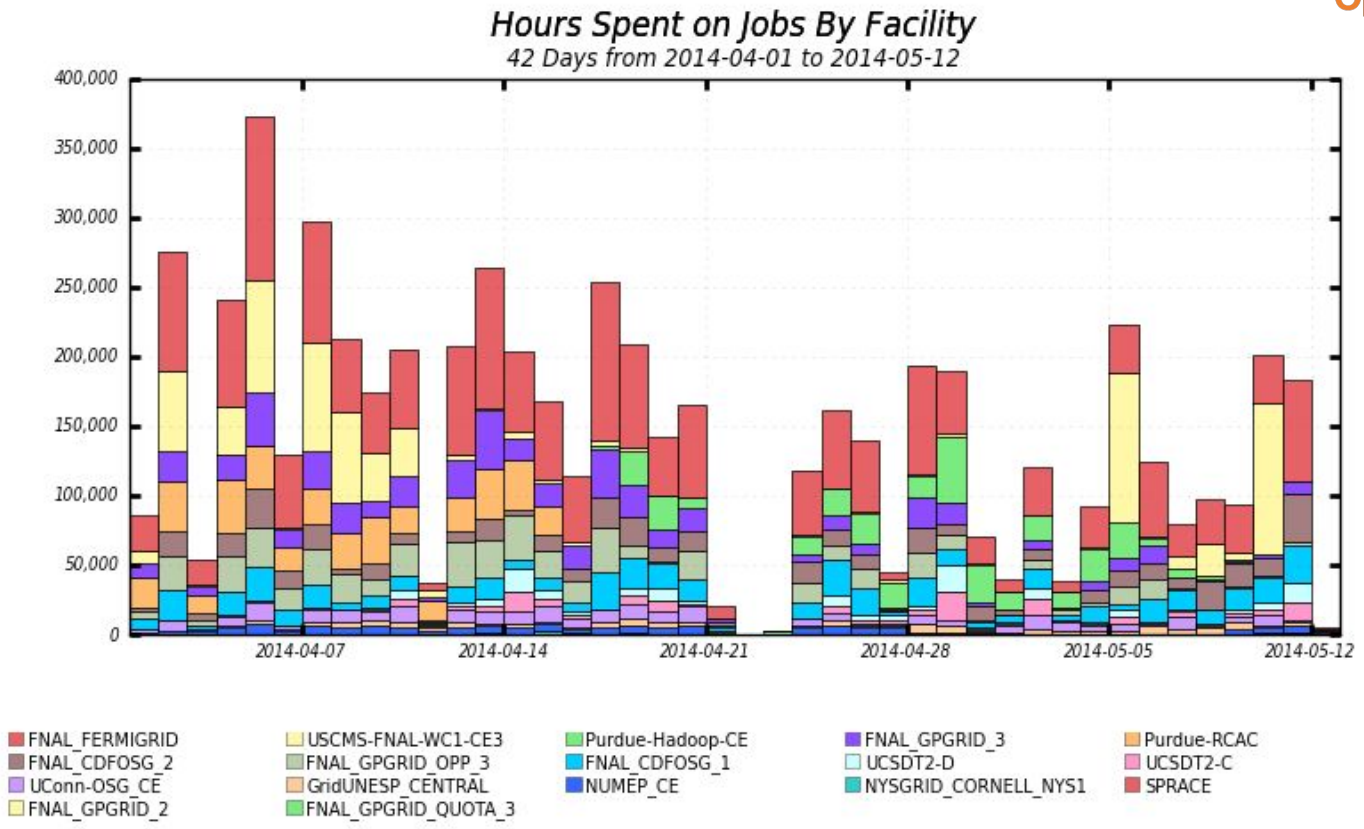

Gluex activity on osg 2014-2016

Open Science Grid

This work is supported by the National Science Foundation under grant 1508238

Gluex @ – the reboot

- **OSG Executive Director, Frank Wuerthwein** speaks at *NP Computing Workshop,* Newport News, VA in March, 2016.
- **JLab CIO, Amber Boehnlein** initiates a pilot project for JLab users.

scosg16: a GWMS submit host for JLab users

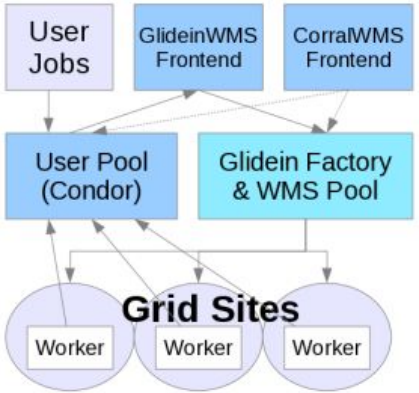

- ➢ located at JLab
- ➢ supported by JLab IT staff
- ➢ GlueX to be among the first users
- ➢ only out-flow of work is currently envisioned
- ➢ server configuration recommended, tested by OSG expert
- ➢ server installed, configured in 2Q 2017, testing by GlueX is now underway.

This work is supported by the National Science Foundation under grant 1508238

Gluex @ – the reboot

New infrastructure for osg @ jlab:

1. scosg16: GWMS submit host for JLab users
2. GWMS Frontend service provided by OSG ops

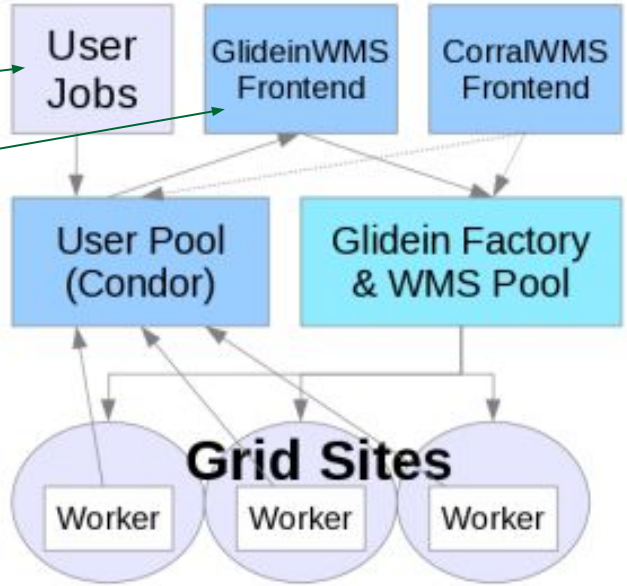

This work is supported by the National Science Foundation under grant 1508238

Gluex @ – the reboot

New infrastructure for osg @ jlab:

1. scosg16: GWMS submit host for JLab users
2. GWMS Frontend service provided by OSG ops
3. Opportunistic cycles on OSG continue to grow
4. Two new member universities in Gluex moving this summer to stand up local resources on osg
5. Software distribution is now greatly simplified by the use of the new ***Gluex singularity container***:
   - **singularity.opensciencegrid.org**
   - **oasis.opensciencegrid.org**

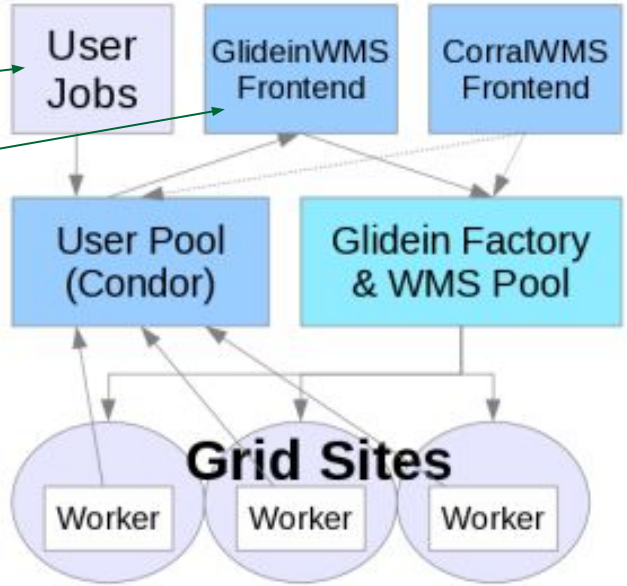

This work is supported by the National Science Foundation under grant 1508238

GlueX @  - opportunity cost

- osg represents a new way of working for JLab users
- lab IT management conscious of *user support issues*
- JLab collaborations are small, developing new expertise can be expensive

This work is supported by the National Science Foundation under grant 1508238

GlueX @  - opportunity cost

- osg represents a new way of working for JLab users
- lab IT management conscious of *user support issues*
- JLab collaborations are small, developing new expertise can be expensive

BUT

- grid production is a good match to GlueX needs for simulations
- recent work by ***OSG + JLab*** staff has been ***a real boost***
- new effort is underway to enable us to exploit OSG for Gluex

This work is supported by the National Science Foundation under grant 1508238

Backup slides

This work is supported by the National Science Foundation under grant 1508238

Support for Gluex users

- Support for resource consumers (15 users registered)
  - howto get a grid certificate
  - howto access data from DC
  - howto test your code on osg
  - howto run your skims on osg

Quickstart users guide for Gluex
https://halldweb.jlab.org/wiki/index.php/Using_the_Grid

Gluex OSG HOWTO series (R.Jones)
https://halldweb.jlab.org/wiki/index.php/HOWTO_get_your_jobs_to_run_on_the_Grid

- Support for resource providers (UConn, NWU, FIU, FSU, CMU, IU, MIT?)
  - NOT a commitment to 100% allocation to OSG jobs
  - OSG site framework assumes that the local admin retains full control over resource utilization (eg. supports priority of local users)
  - UConn Gluex site running for 8 years
  - Northwestern Gluex site running for 3 years

This work is supported by the National Science Foundation under grant 1508238

GlueX Data Challenge #1

- total of 5,561,650000 events *successfully* generated
  - 4G events produced on the OSG (~2M core-hours)
  - 0.9G events at Jefferson Lab
  - 0.3G events at CMU
- completed over a period of 14 days in Dec., 2012
- output data saved in REST format
  - Reconstructed Event Summary Type (no hits information)
  - approx. 2.2 kB/event, including MC generator event info
  - hadronic interaction in every event (pythia 8.4 – 9.0 GeV)
  - no em beam background or hadronic pile-up included
  - 111236 files stored, 50k events each
  - typical run time 8 hours / job on Intel i7

This work is supported by the National Science Foundation under grant 1508238

Problems encountered in OSG production

1. GlueX software environment staging
   - 20 packages to install (counting all of sim-recon as 1)
   - production spread over 8 sites (fnal.gov, cornell.edu, purdue.edu, ucllnl.org, ucsd.edu, unesp.br, org.br, uconn.edu)
2. freeze-ups in hd-ana
   - occurred any time an event took >30s to process
   - dependent on other things happening at the site
   - tended to occur in clusters, many jobs at once
3. memory hogging in hd-ana (feeds into 2)
4. segfaults in hdgeant
   - artifact from one node at UConn – bad SDRAM chip
5. irreproducibility in mcsmear

This work is supported by the National Science Foundation under grant 1508238

Production inefficiency

- 10% jobs would hang in hd_ana, up to 24hr.
- 24hr is 300% inflation of normal job time
- Ejected jobs would get requeued for later execution.
- Some fraction of these would hang 2nd, 3rd time around…
- Ad-hoc scripts were written to prune jobs that were stuck looping.
- Other known factors (store output to SRM, thrashing on memory hogs…) not quantified.

FNAL firewall intrusion event

hung job script development

job queue drainout

This work is supported by the National Science Foundation under grant 1508238