CLAS Collaboration Meeting, November 12, 2019

GlueX Experience with Off-Site Simulation

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 Offsite Computing Plan

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)*
   - *PSC (XSEDE, also competitve), other ??*
2. 36 Mcore-hr/yr - Monte Carlo simulation
   - *primarily targeted for OSG*
   - *opportunistic usage alone is not adequate*

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

Existing OSG resources for GlueX

1. **UConn_OSG** site: 600-core cluster
   - active on OSG since ca. 2010
   - contributed 2-3 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 post-2019.

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

GlueX Opportunistic Usage on OSG

Open Science Grid

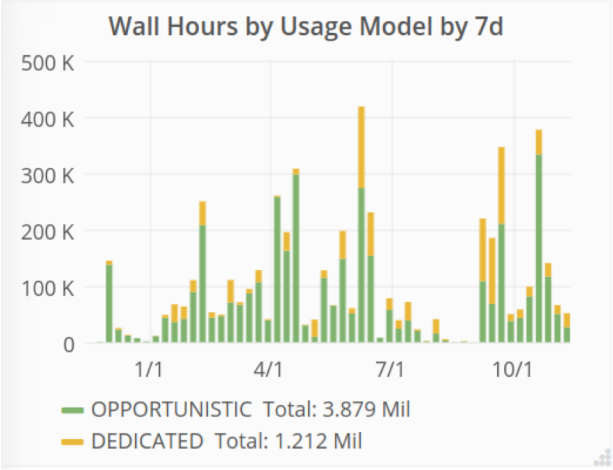

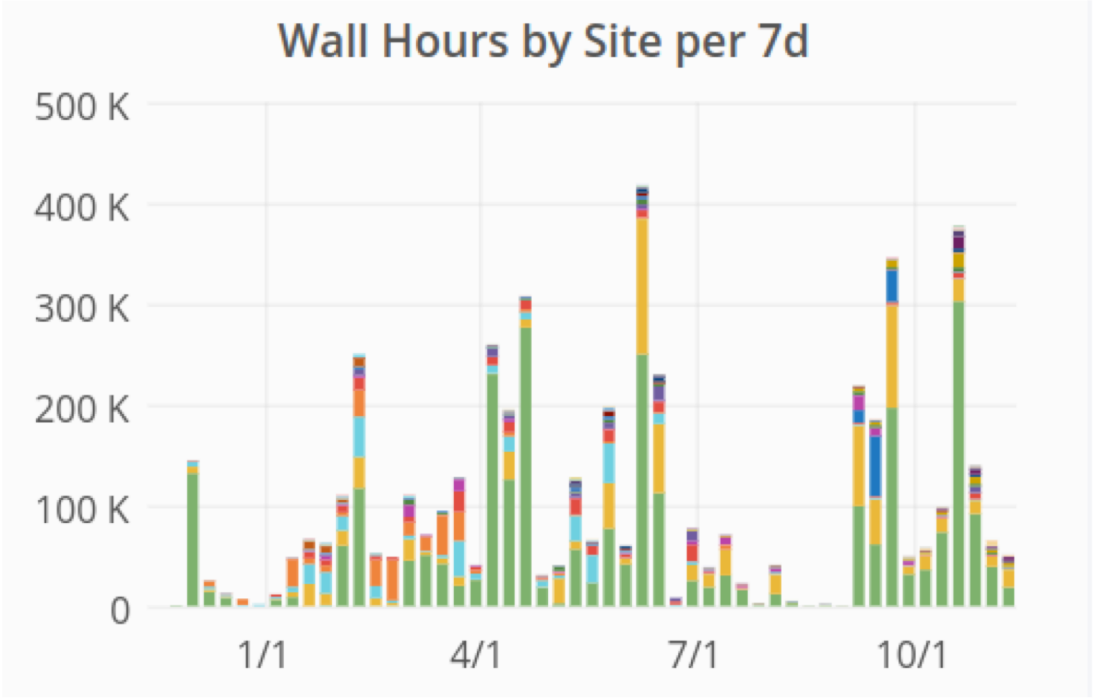

This work is supported by the National Science Foundation under grant 1508238

GlueX Opportunistic Usage on OSG

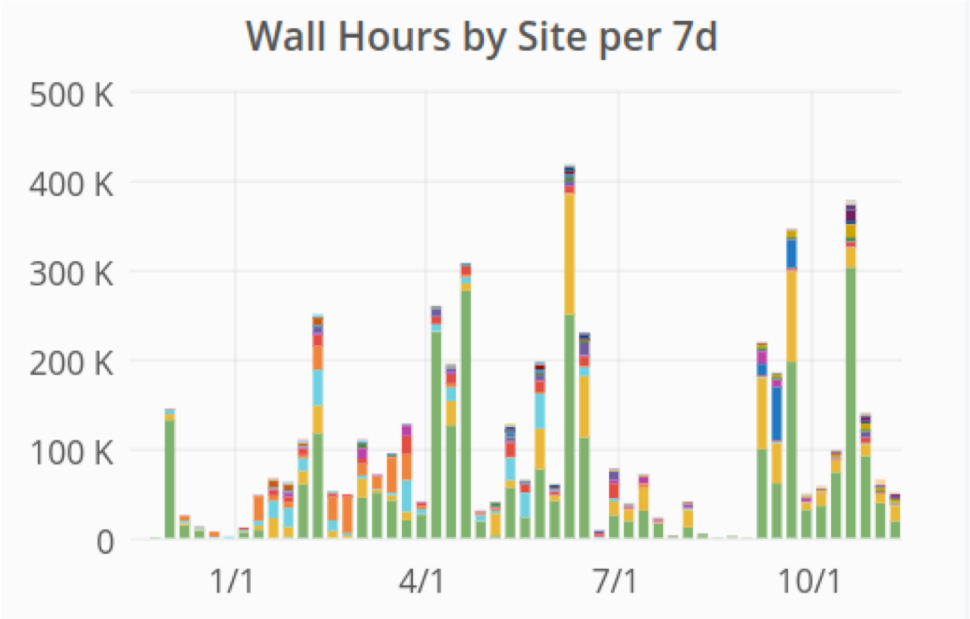

SU ITS
UConn-OSG
MWT2 ATLAS UC
SPRACE
UColorado_HEP
UKI-SCOTGRID-GLASGOW
UCSD CMS Tier2
NWICG_NDCMS
GLUEX_US_FSU_HNPGRID
ASU Research Computing
Nebraska-CMS
AGLT2
Information Sciences Institute
BNL ATLAS Tier1
NEWJERSEY_ELSA
New Mexico State Discovery
cinvestav
GLOW
Caltech CMS Tier2
UPRM_HEP
Nebraska-Omaha
Clemson IT
IIT - Illinois Institute of Technology
Stampede2
OSG_US_WSU_GRID_ce2
UMD_SIAB
JetStreamIU
GLUEX_US_UConn_cedar
Nebraska-Lincoln
Bellarmine University
FIU HPC
University of South Florida
UF HPC
OSG_US_GSU_ACORE
N/A

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

GlueX Opportunistic Usage on OSG

1. There are sizable opportunistic cycles available on OSG
   - *This is what grid computing is about!*
   - Probably not enough to accommodate the full GlueX need for offsite simulations.
2. Opportunity for growth: ***shared local resources***
   - Universities are developing local shared research IT
   - Intended to leverage local IT expertise, infrastructure to boost the productivity *(grant funding)* of local researchers.

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

Potential local GlueX resources

Open Science Grid

Survey of interested institutions taken in spring 2018:

a. Carnegie Mellon University - PSC, local cluster
b. Indiana University - stanley, karst, BigRed
c. Florida State University - rcc
d. George Washington University - colonialone
e. College of William and Mary - vortex
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

Potential local GlueX resources

Open Science Grid

Two options were offered:

1. Regular OSG site integration
   - significant initial effort by admins
   - entails buy-in to grid computing concept
   - minimal cost on the side of GlueX
2. Campus cluster site configuration
   - minimal effort by admins, uses a local user account
   - communication with admins is important, so they are on-board
   - non-trivial cost on the side of GlueX production manager

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

Potential local GlueX resources

Two options were offered: ***in 2018 this is what happened***

1. Regular OSG site integration -- ***nobody took this route***
   - significant initial effort by admins
   - entails buy-in to grid computing concept
   - minimal cost on the side of GlueX
2. Campus cluster site configuration -- ***6 universities opted-in***
   - minimal effort by admins, uses a local user account
   - communication with admins is important, so they are on-board
   - non-trivial cost on the side of GlueX production manager

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

GlueX experience: offsite university resource integration

Open Science Grid

Summer 2018

- 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, 8 in total)
- install local copy of complete GlueX stack + container
- diagnose, debug, optimize...

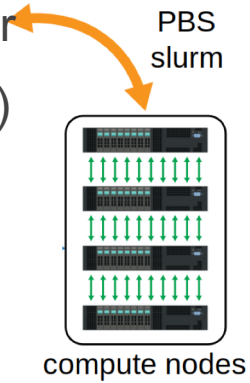

This work is supported by the National Science Foundation under grant 1508238

GlueX experience: clarification

Open Science Grid

What we never considered doing:

- Setting up custom workflows on each separate cluster using the local dialects of the campus cluster, custom scripts for each site, etc...
- This is what JLab users have been doing since forever, with local users managing the complexity of translating collaboration-wide scripts to the local dialect.
- This generally has worked for local analyses, limited scale, but...
- This does not scale up to a distributed production across many sites.

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

GlueX experience: clarification

Open Science Grid

What OSG workflows do well:

- Hide the complexity of a distributed environment
- Allow a single production to run across a diverse set of sites
- *Duplicates offsite what the JLab farm provides onsite*

What the challenge was:

- How to integrate campus clusters into the OSG production ecosystem without requiring the contributing clusters to become OSG grid sites?

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

GlueX experience: offsite university resource integration

1. Lessons from the summer 2018 integration test
   - 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. Operations required considerable effort
   - jobs flowed from one submit node at UConn to diverse remote sites
   - connections to individual clusters over ssh managed by condor
   - (mis)communication with cluster admins ***-- the unexpected hurdle!***

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

GlueX experience: offsite university resource integration

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.
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

GlueX experience: offsite university resource integration

What progress has been made over the past year?

1. OSG Central Ops have agreed to take over management of integrated GlueX campus cluster resources.
    - decision taken at the All-Hands Meeting (here) last March
    - implies some delay: additional layers of communication, knowledge transfer from GlueX to Campus Clusters Team at Wisconsin
    - *critical if this success is to be transferrable to other collaborations!*

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

GlueX experience: offsite university resource integration

Open Science Grid

What progress has been made over the past year?

2. Integration with computecanada is now complete.
3. Integration with UConn's xanadu and storrs.hpc clusters is underway.
4. More member university groups are queued up.
5. Major upgrade to UConn shared cluster with OSG integration for GLUEX + CLAS funded by NSF this past summer!

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

Other lessons learned: *negotiating resource integration*

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

Other lessons learned: *negotiating resource integration*

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.

5. Only outgoing connectivity to the internet is required by these jobs. No specialized software libraries and no license agreements are needed.

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

Other lessons learned:
negotiating resource integration

Example framework for successful discussions:

6. 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.

7. 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. We will provide a list of contacts with phone numbers and emails in case anything suspicious arises and you need to investigate.

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

Other lessons learned: *negotiating resource integration*

Example framework for successful discussions:

8. Batch job environment: Westgrid slurm with queues, priorities, and job policies set by you -- that we will respect -- nothing more.
9. 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.
10. Use of singularity on the workers is desired but not required.

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

GlueX experience: offsite university resource *expansion*

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

GlueX experience: offsite university resource *expansion*

University of Connecticut proposal 1925716

- submitted February 20, 2019
- $400,000 for compute nodes (2300 cores) + storage (1 PB)
- **notification of award in July, 2019**
- 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

GlueX experience: offsite university resource *expansion*

Just last week: notice from Allena Opper:

RFI on Data-Focused Cyberinfrastructure Needed to Support Future Data-Intensive S&E Research

I encourage you to respond to NSF 20-015 Dear Colleague Letter, *Request for Information (RFI) on Data-Focused Cyberinfrastructure Needed to Support Future Data-Intensive Science and Engineering Research,* https://www.nsf.gov/publications/pub_summ.jsp?ods_key=nsf20015.

The challenges of growing volumes of scientific data – their availability, transmission, accessibility, management, and utilization – have become urgent and ubiquitous across NSF-supported science, engineering, and education disciplines. To inform the formulation of a strategic NSF response to these imperatives, the RFI asks the research community to update NSF on their data-intensive scientific questions and challenges and associated needs specifically related to data-focused cyberinfrastructure.

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

OSG integration: *beyond simulation?*

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 ??*

Can OSG 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

Summary and outlook

- Plans for doing most of GlueX simulation offsite on the OSG are being implemented.
- So far, most of our needs have been met using opportunistic cycles, but this will not continue indefinitely.
- Within the last 2 years we have successfully demonstrated the extension of OSG-capable sites to include campus clusters.
- Integration of existing campus clusters within the OSG ecosystem has begun, with oversight by OSG Operations.
- Opportunity found for expanded campus resources for Jlab experiments!

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

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 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 rarely updated (once per year?)
4. Multiple binary releases maintained on oasis
    a. selected by demand
    b. size dominated by symbol-rich shared libraries
    c. currently on the high side - 270 GB oasis footprint
    d. may be excessive, but no complaints so far...

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.

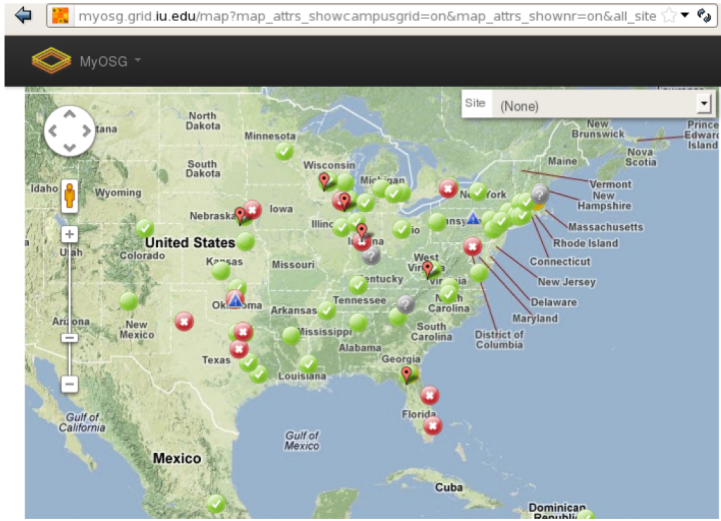

initial goal: 10 billion events, **60 days** at startup intensity

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

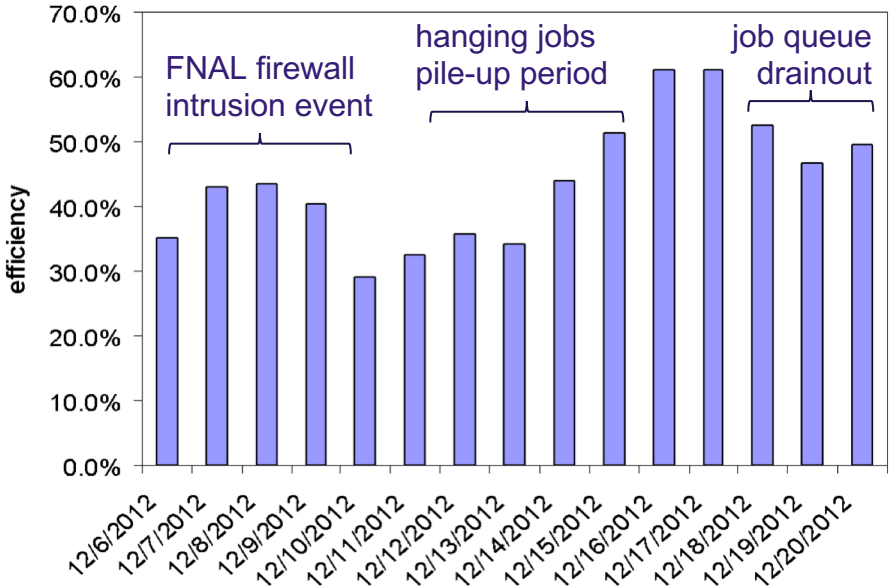

This work is supported by the National Science Foundation under grant 1508238

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.

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

1. **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.

1. **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.

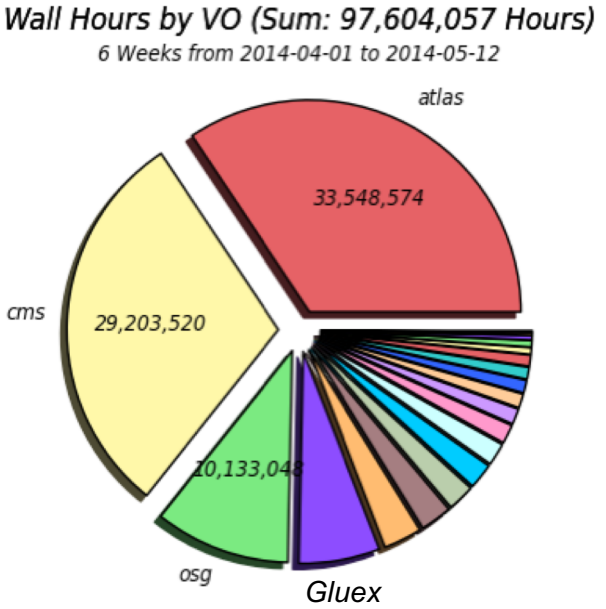

6M core-hr = DC1 x 2

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

Data Challenge 2: *results*

Open Science Grid

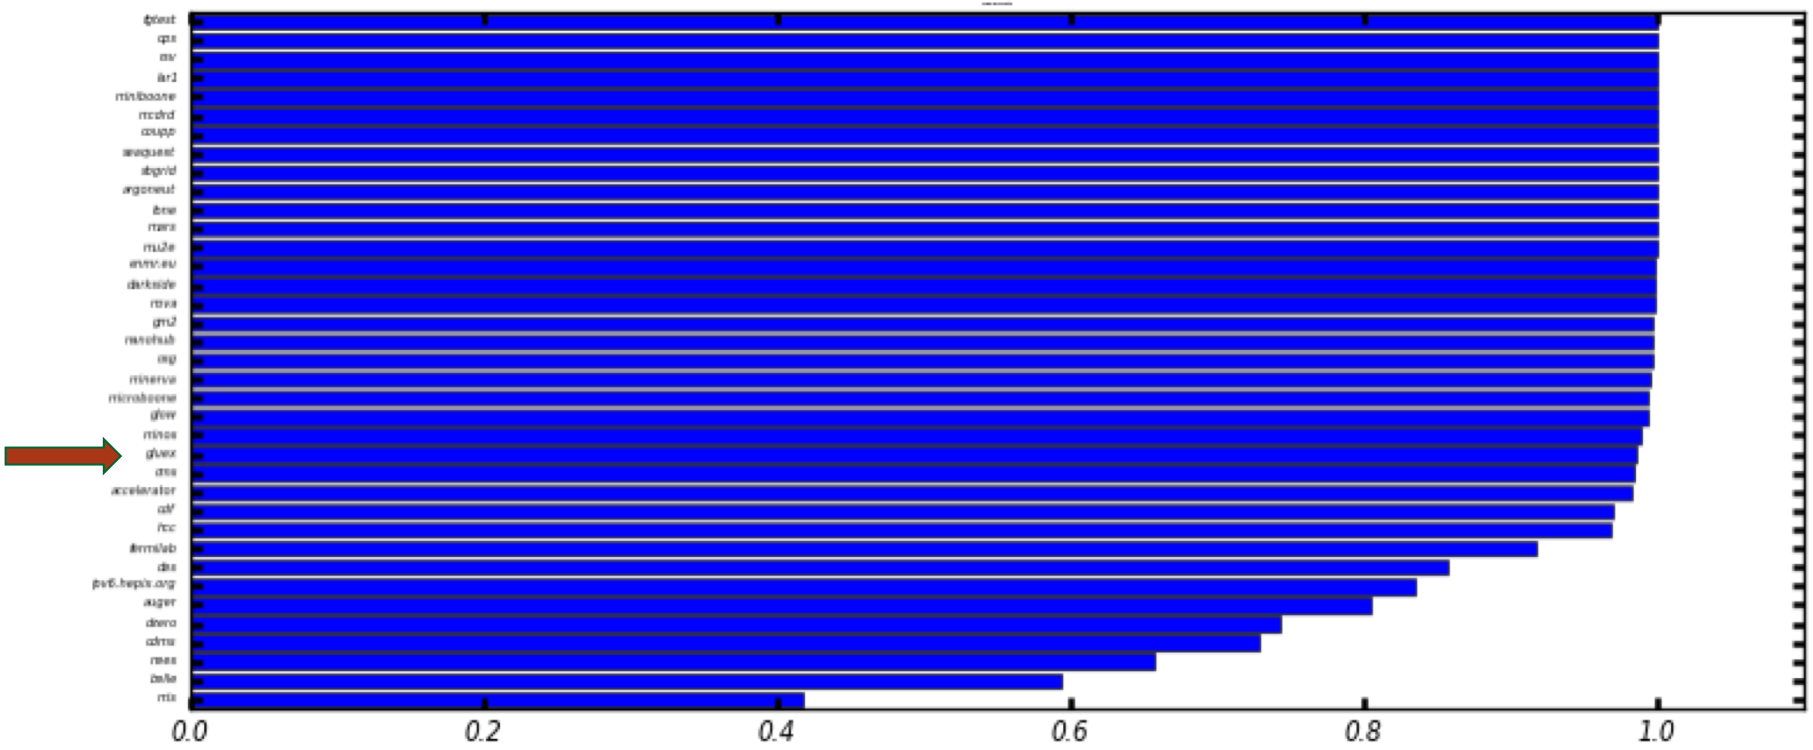

cpu time / wall clock time, daily average

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

Data Challenge 2: *results*

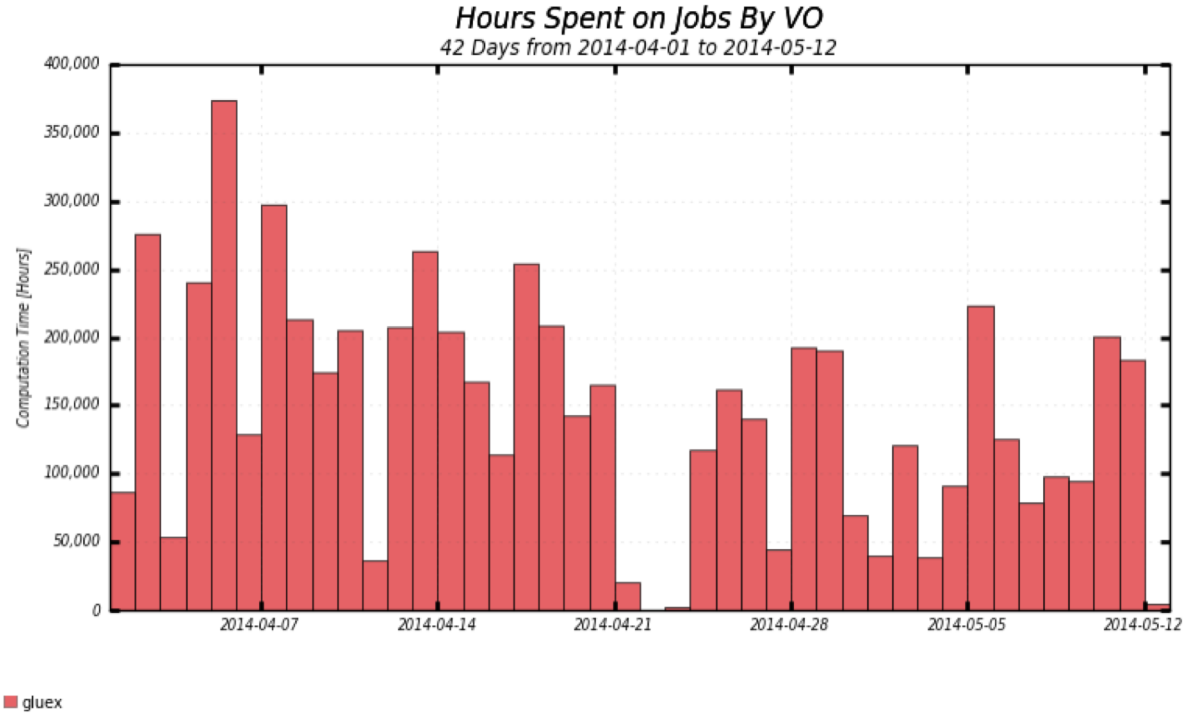

Final event tally

	CMU	
70M	2%	
	MIT	
60M	9%	
	JLAB	2000M
5%		
	OSG	
	5200M	64%
	total	

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

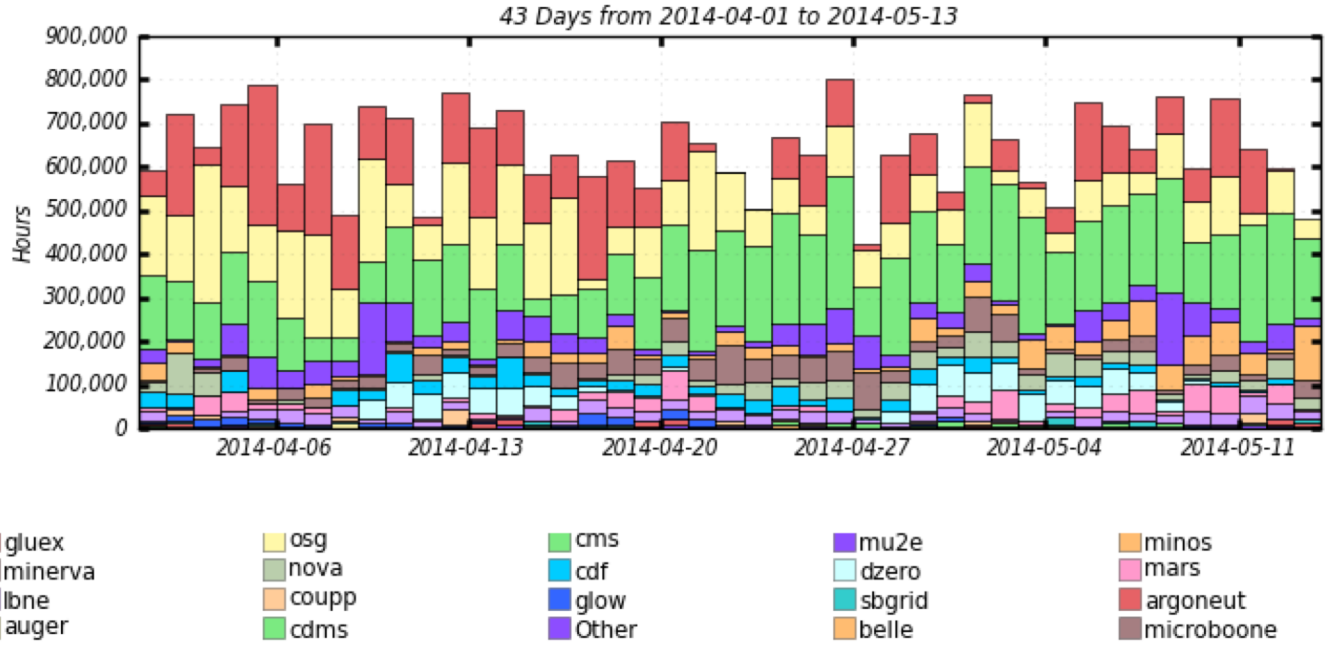

Data Challenge 2: *results*

Open Science Grid

Gluex usage on the Purdue site

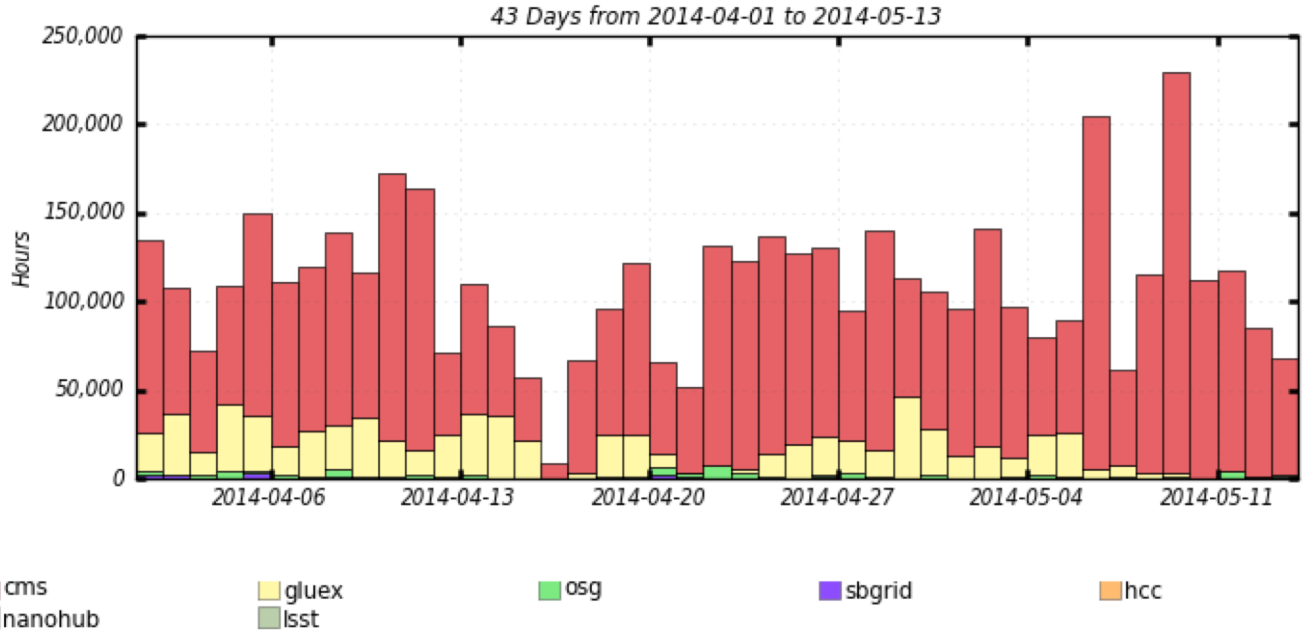

Data Challenge 2: *results*

Open Science Grid

Gluex usage on the Northwestern site

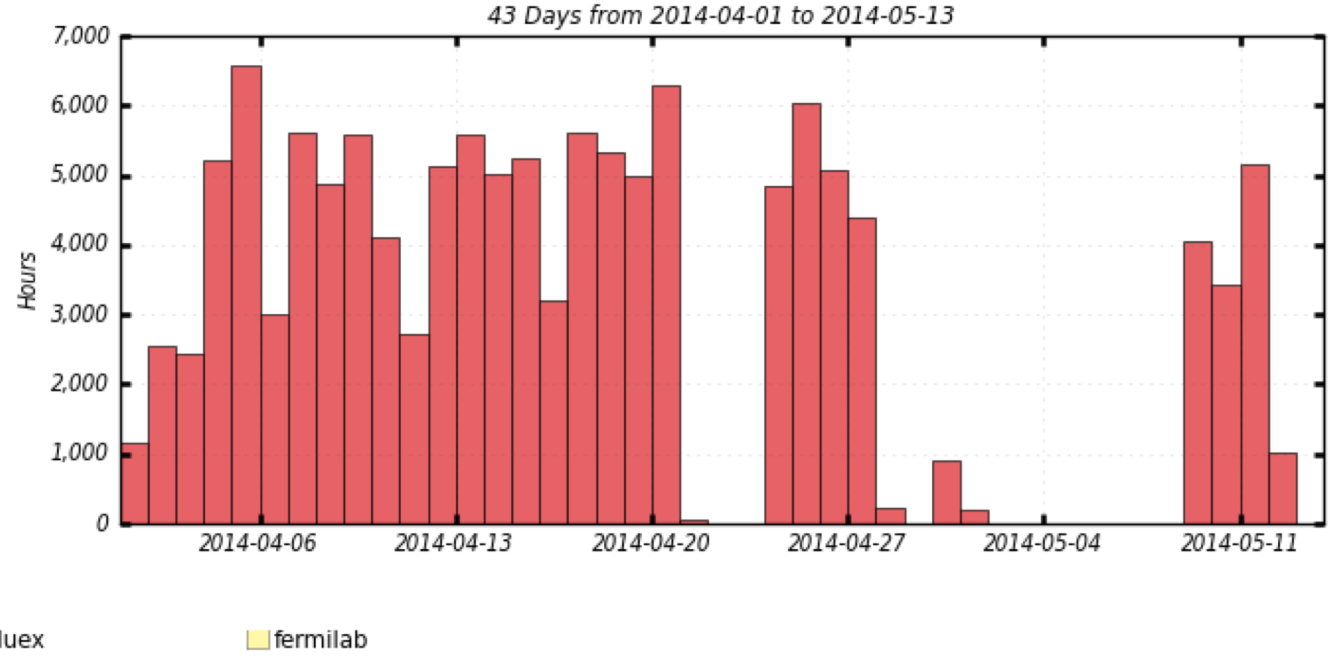

Data Challenge 2: *results*

Open Science Grid

Gluex usage on the UConn site

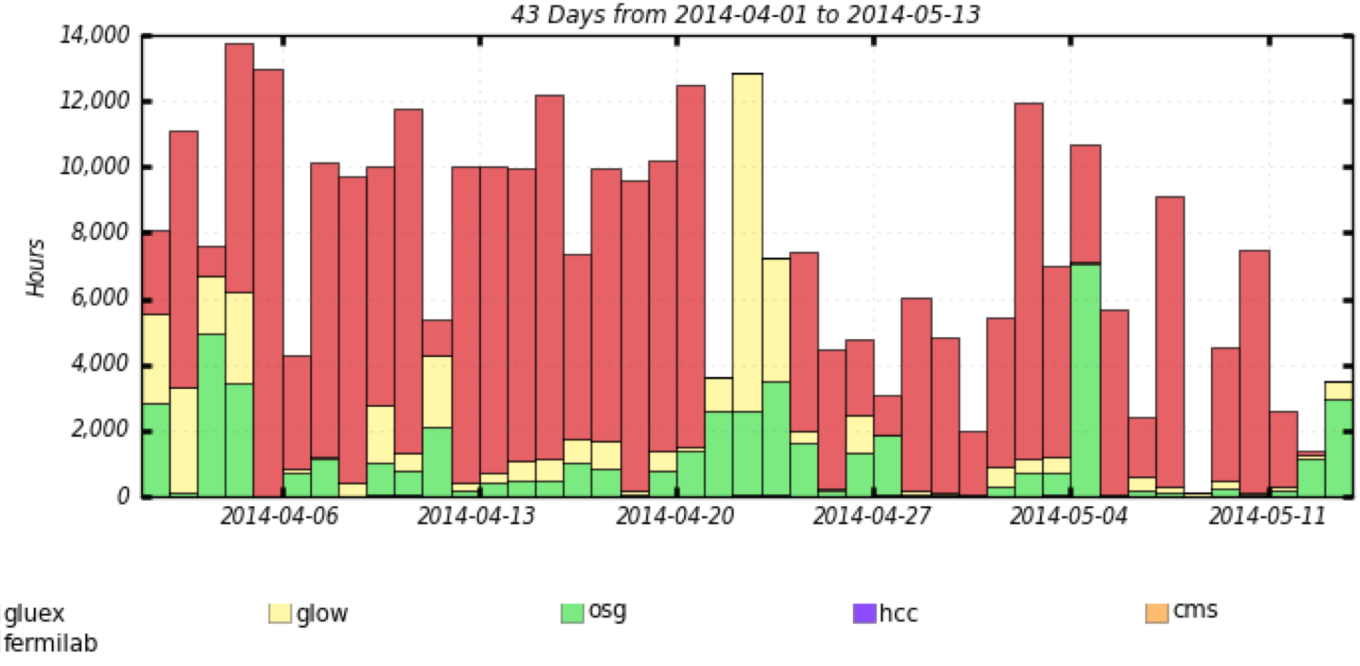

Data Challenge 2: *results*

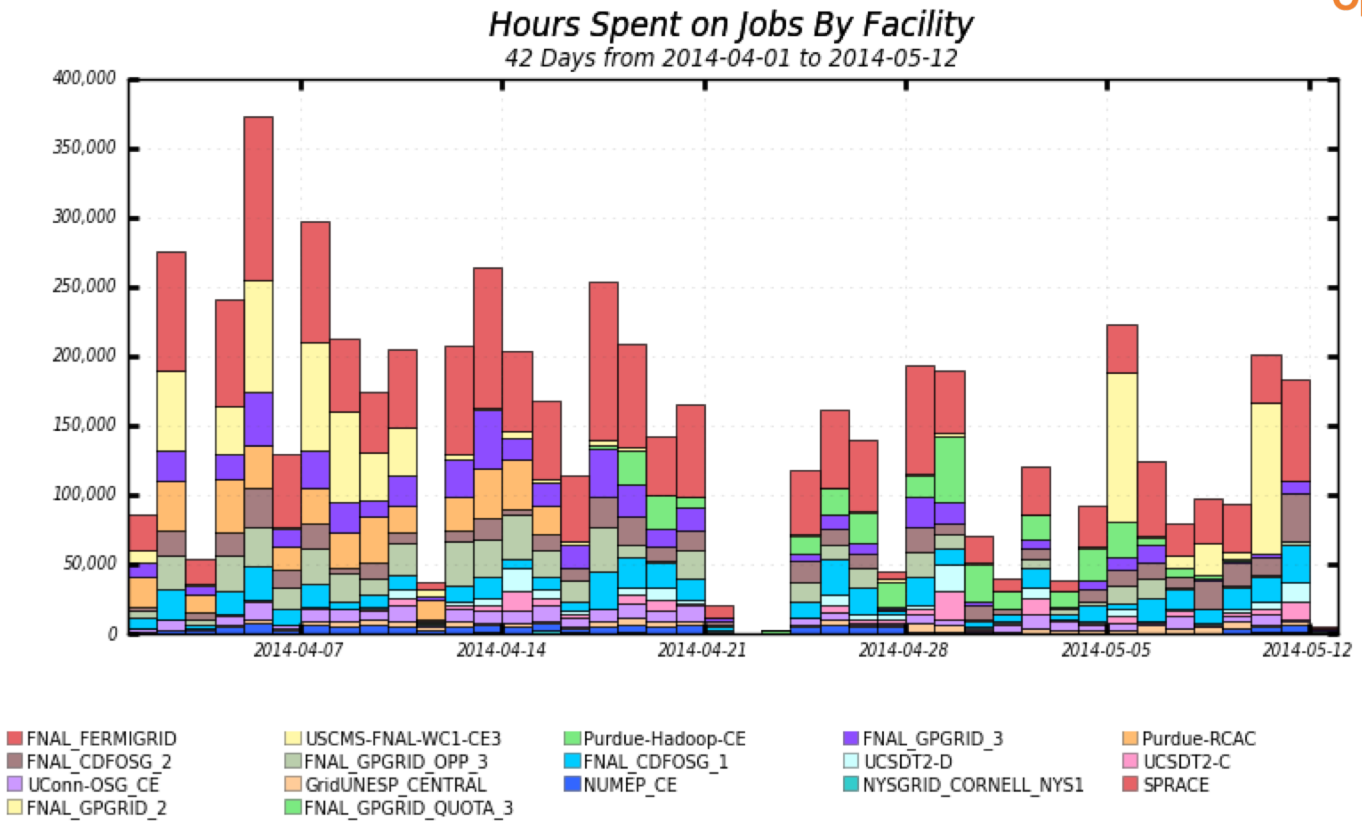

Gluex activity on osg 2014-2016

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

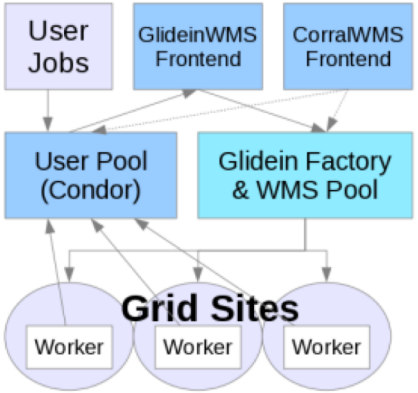

- located at JLab
- supported by JLab IT staff
- GlueX to be among the first users
- only out-flow of work is currently envisioned
- server configuration recommended, tested by OSG expert
- server installed, configured in 2Q 2017, testing by GlueX is now underway.

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

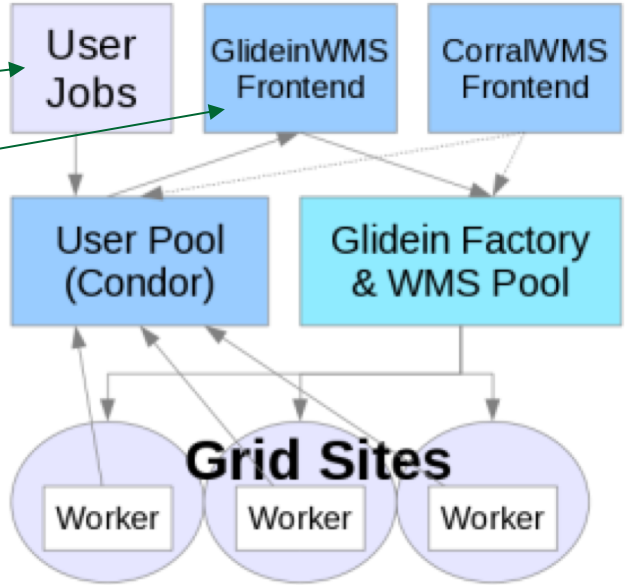

This work is supported by the National Science Foundation under grant 1508238

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**

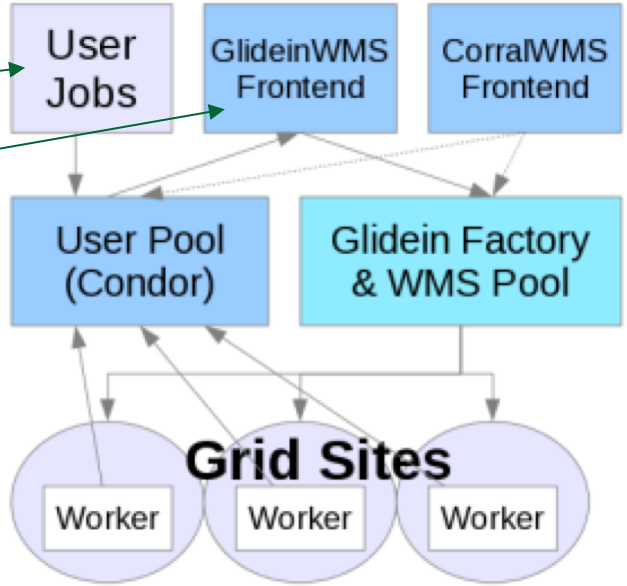

This work is supported by the National Science Foundation under grant 1508238

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

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