OSG: GlueX Experience

past experience and future prospects

Richard Jones, University of Connecticut

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

Background and Motivation

GlueX in Hall D

- search for hybrid mesons
- precision QCD measurements
- dark sector searches

installation: 2012-2014

commissioning: 2014-2016

physics data: 2017-2022, beyond...

Background and Motivation

GlueX data rates

- 50 kHz trigger rate (~1GB/s)
- many exclusive final states
- < 50% are simple topologies

Simulated data estimates

- based on *GlueX-doc-2350* (2013)
- **50 Mcore-hr/yr** (2019 and beyond)
- *large fraction targeted for OSG*

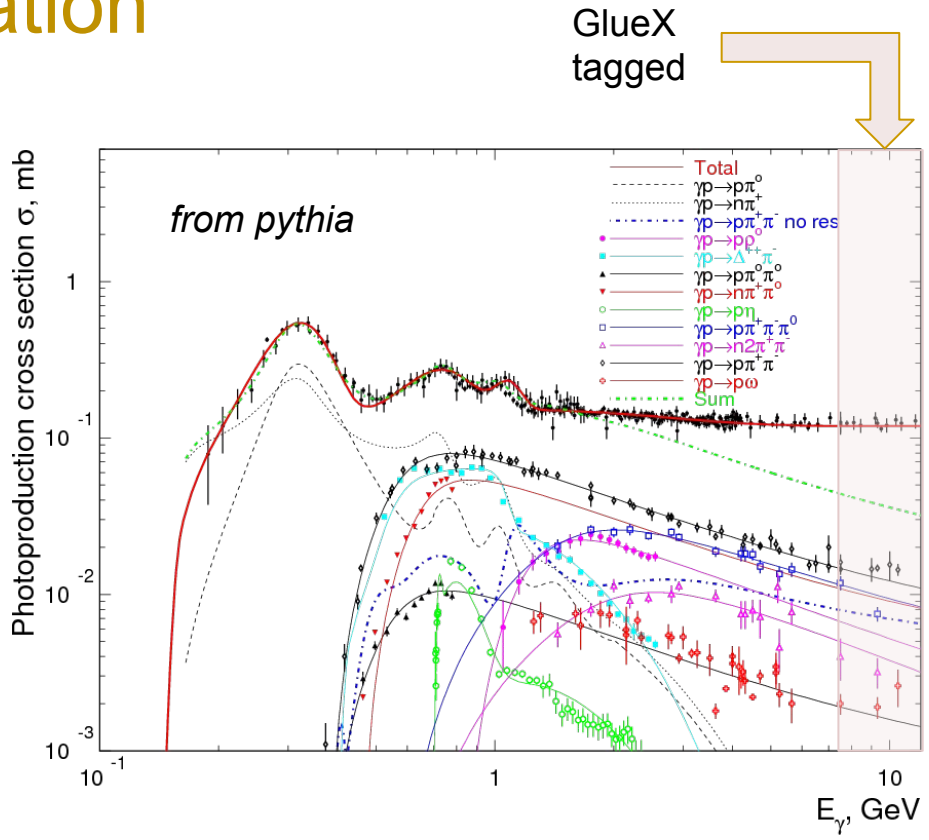

History of the Gluex VO

Open Science Grid

- Gluex vo created in 2009
- UConn initiative, funded by *NSF PIF* program
- underwent an early series of data challenges
  - data challenge I - December 2012, first major osg exercise
  - data challenge II - April 2014, second major osg exercise
  - demonstrated promise of osg for Gluex
  - revealed some issues, bottlenecks
- since then osg activity has been limited by several things
  - effort focused on commissioning the detectors
  - computing effort focused on **online data challenges**
- ***renewed interest arose in 2016***

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 !*

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.

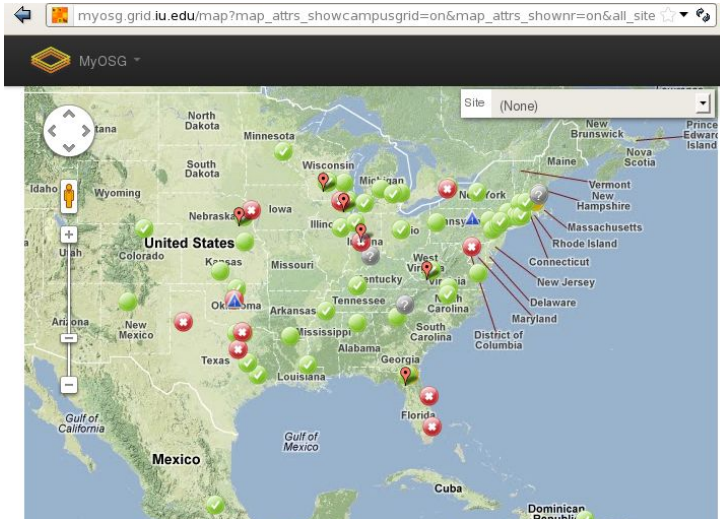

initial goal: 10 billion events, **60 days** at startup intensity

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

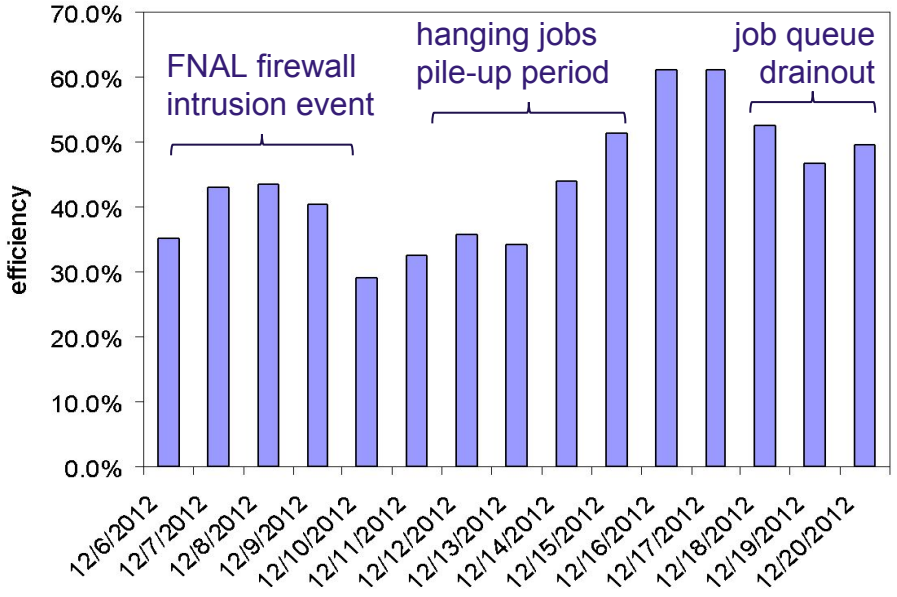

Data Challenge 1: *results*

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

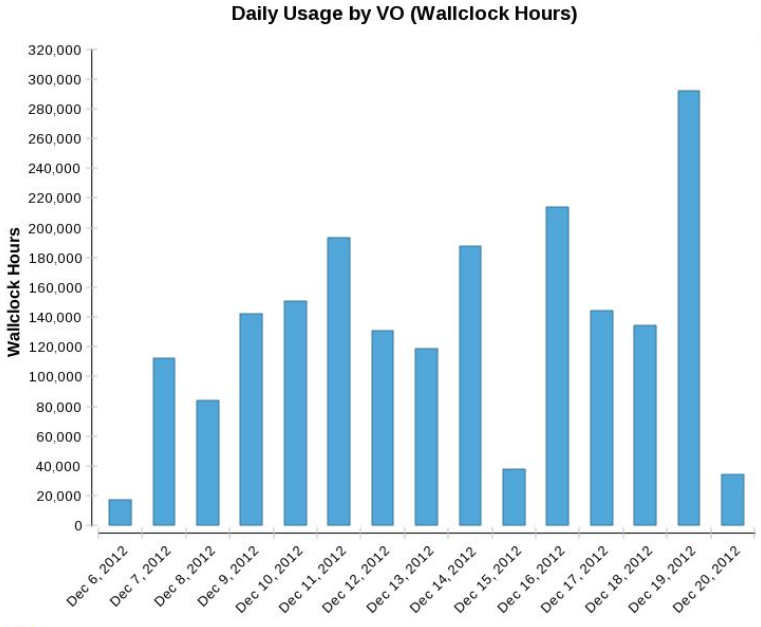

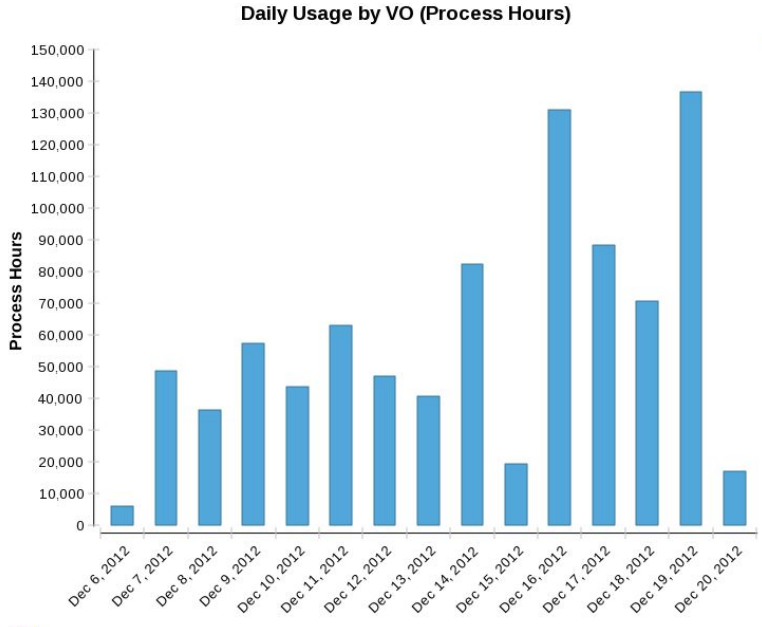

Data Challenge 2: Apr. 5-24, 2014

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.

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.

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.

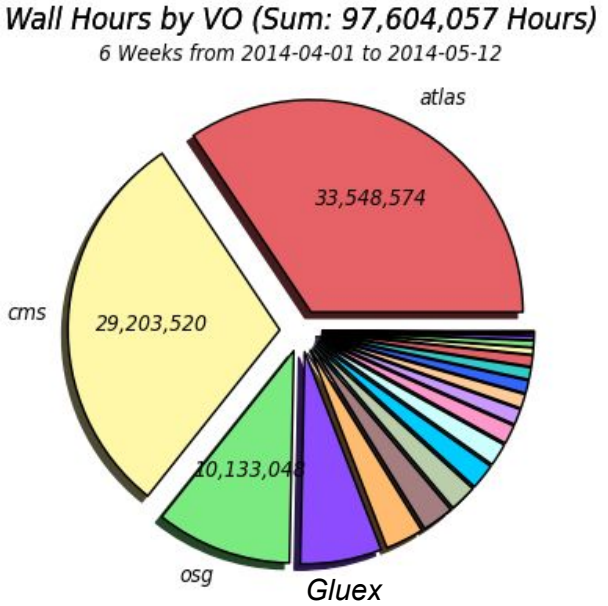

6M core-hr = DC1 x 2

Data Challenge 2: *results*

Open Science Grid

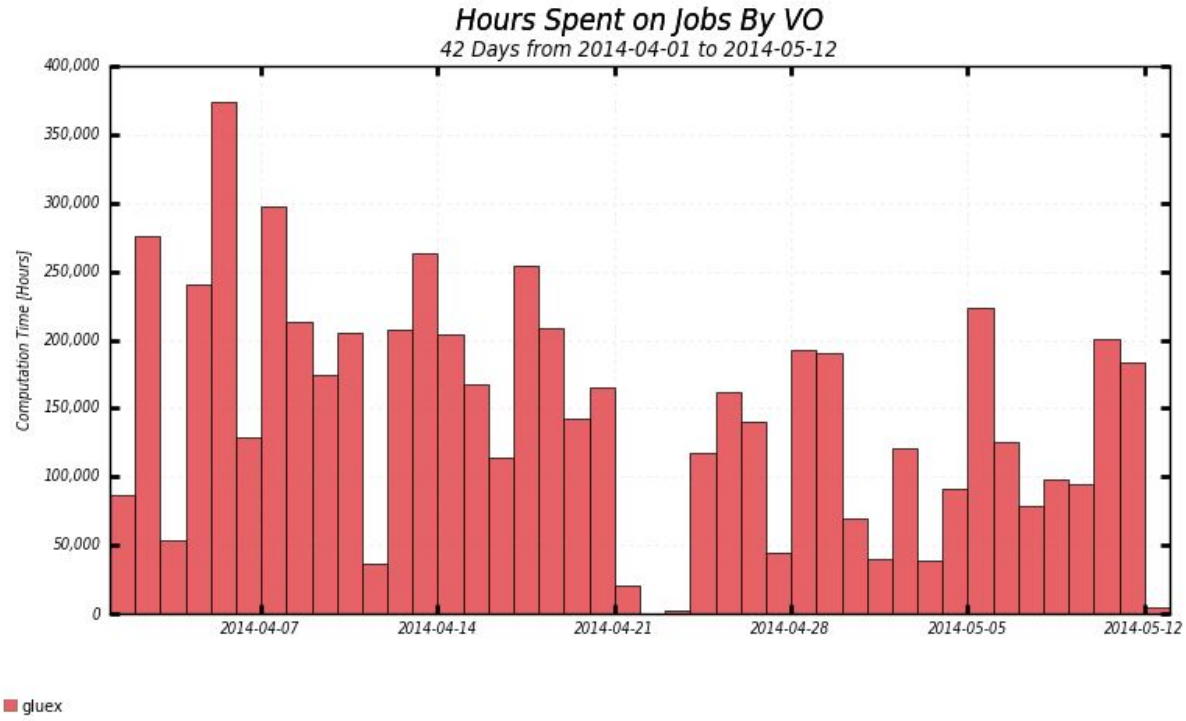

Final event tally

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

Data Challenge 2: *results*

Open Science Grid

Gluex usage on the Fermilab site

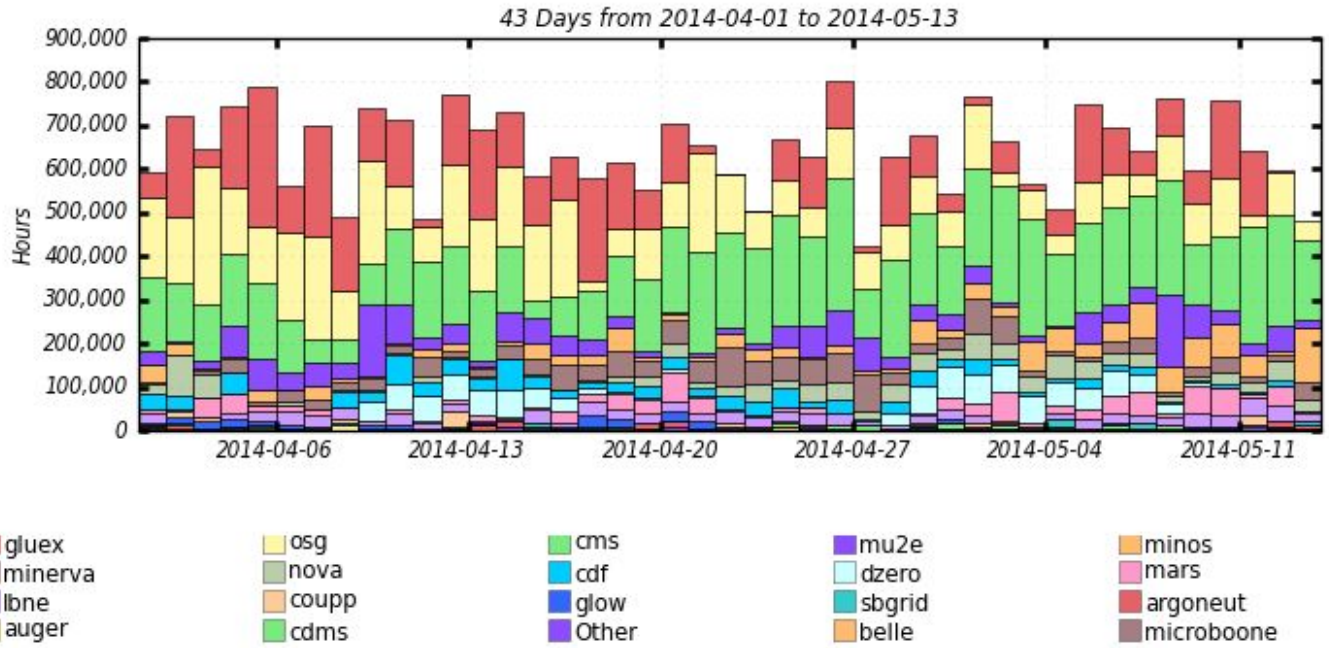

Data Challenge 2: *results*

Open Science Grid

Gluex usage on the Purdue site

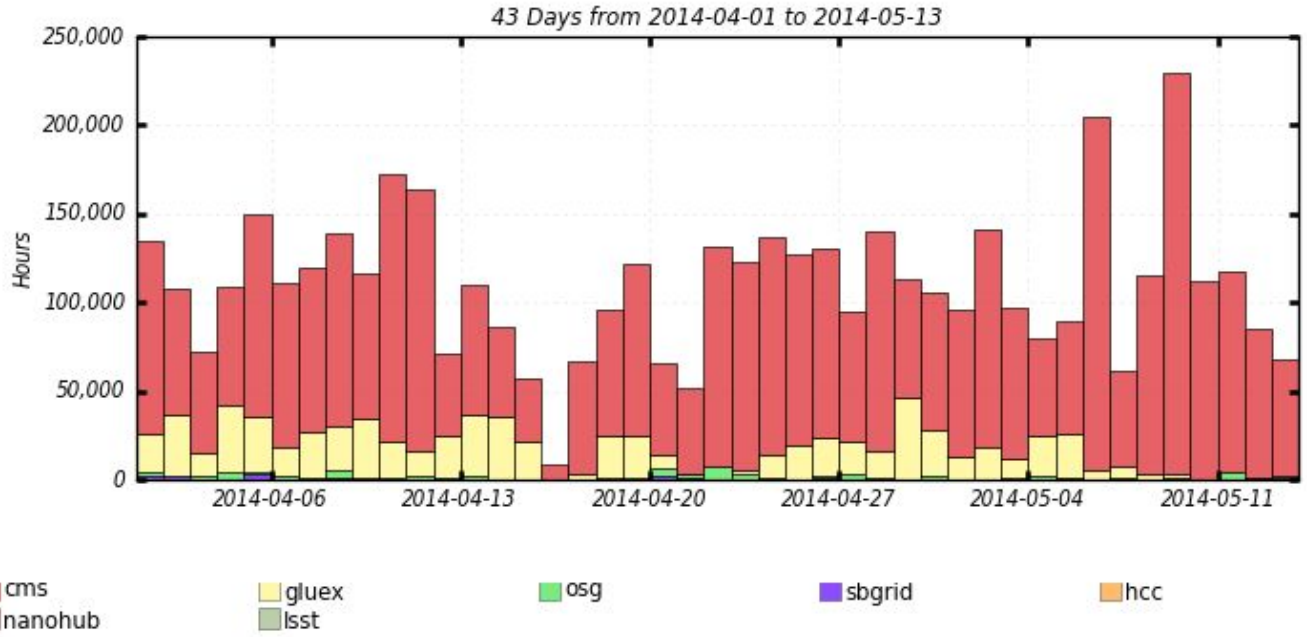

Data Challenge 2: *results*

Open Science Grid

Gluex usage on the Northwestern site

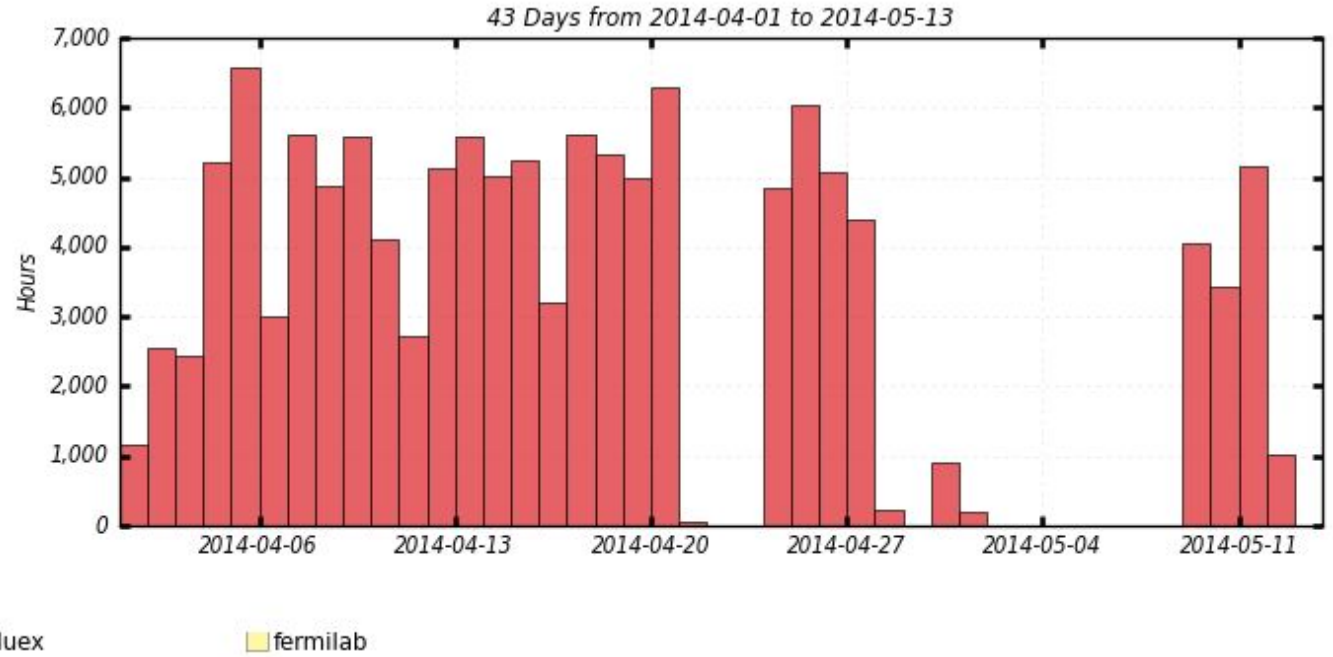

Data Challenge 2: *results*

Open Science Grid

Gluex usage on the UConn site

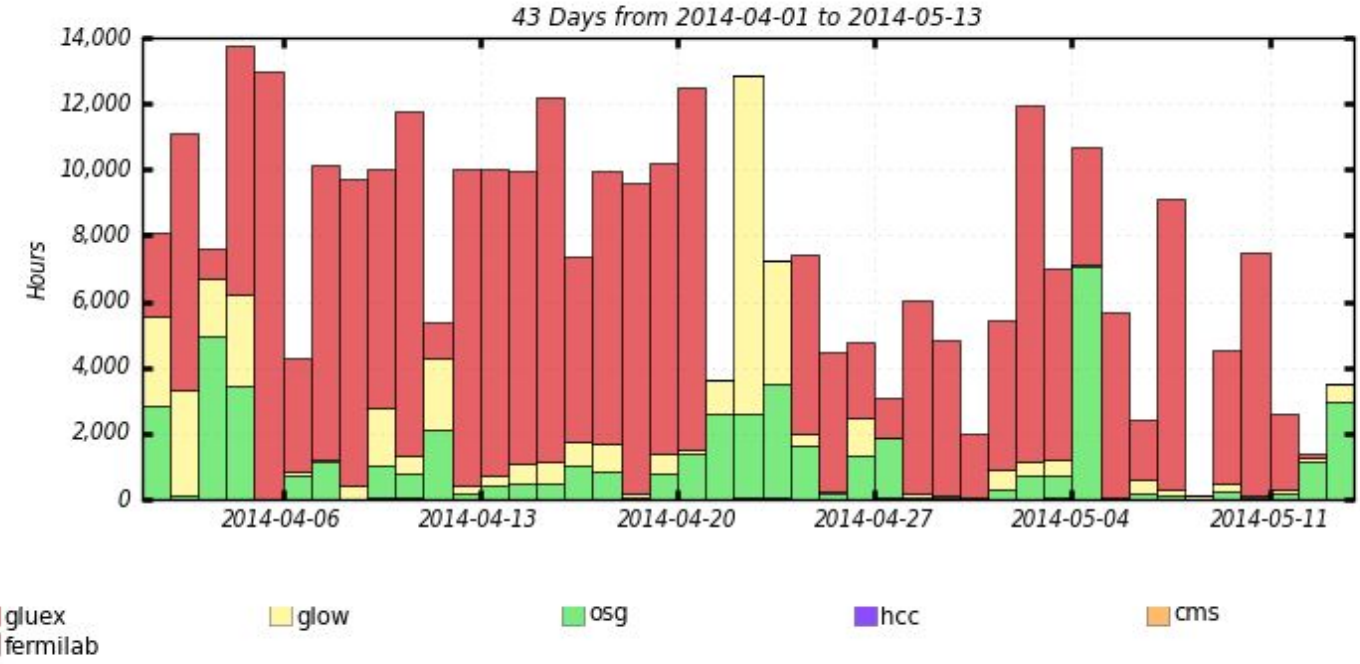

Data Challenge 2: *results*

Open Science Grid

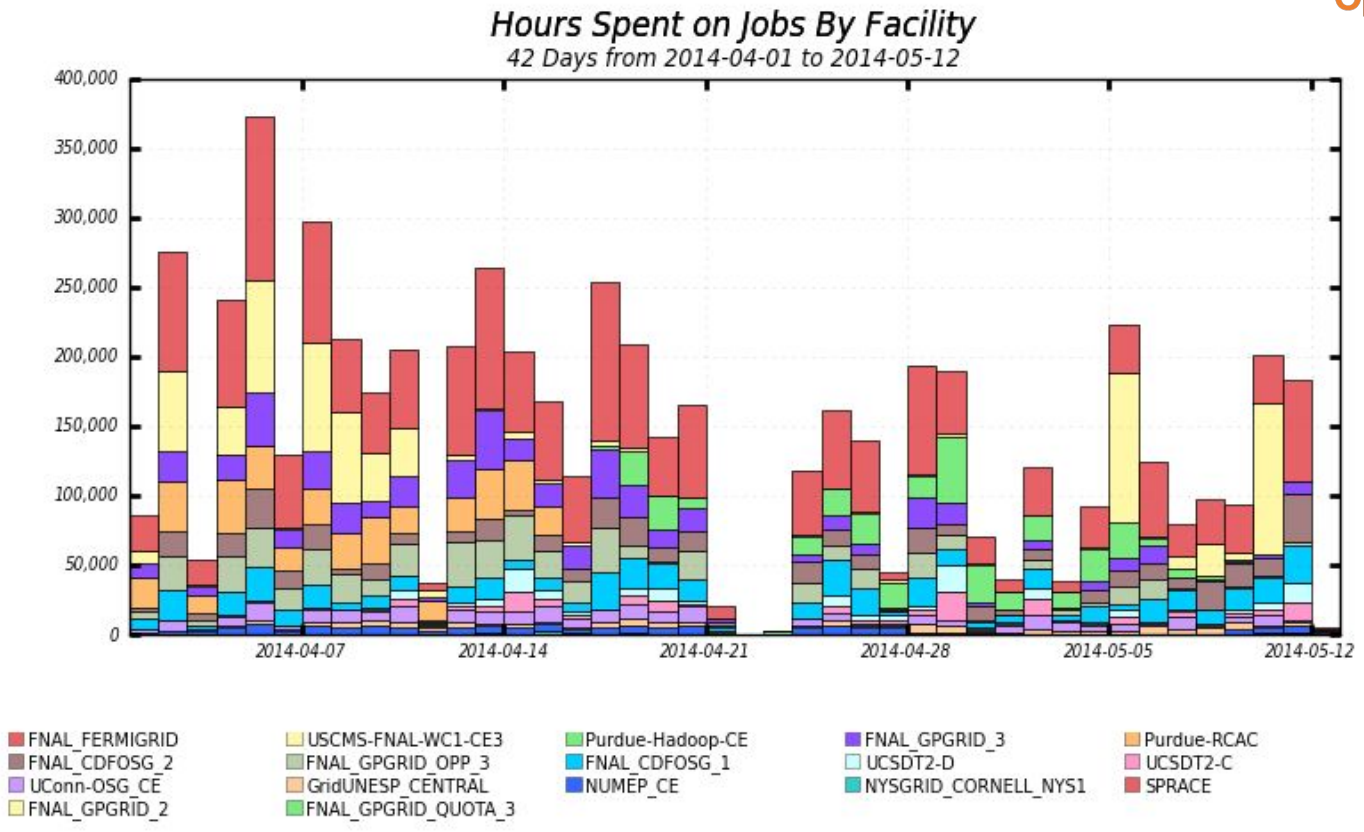

Data Challenge 2: *results*

Open Science Grid

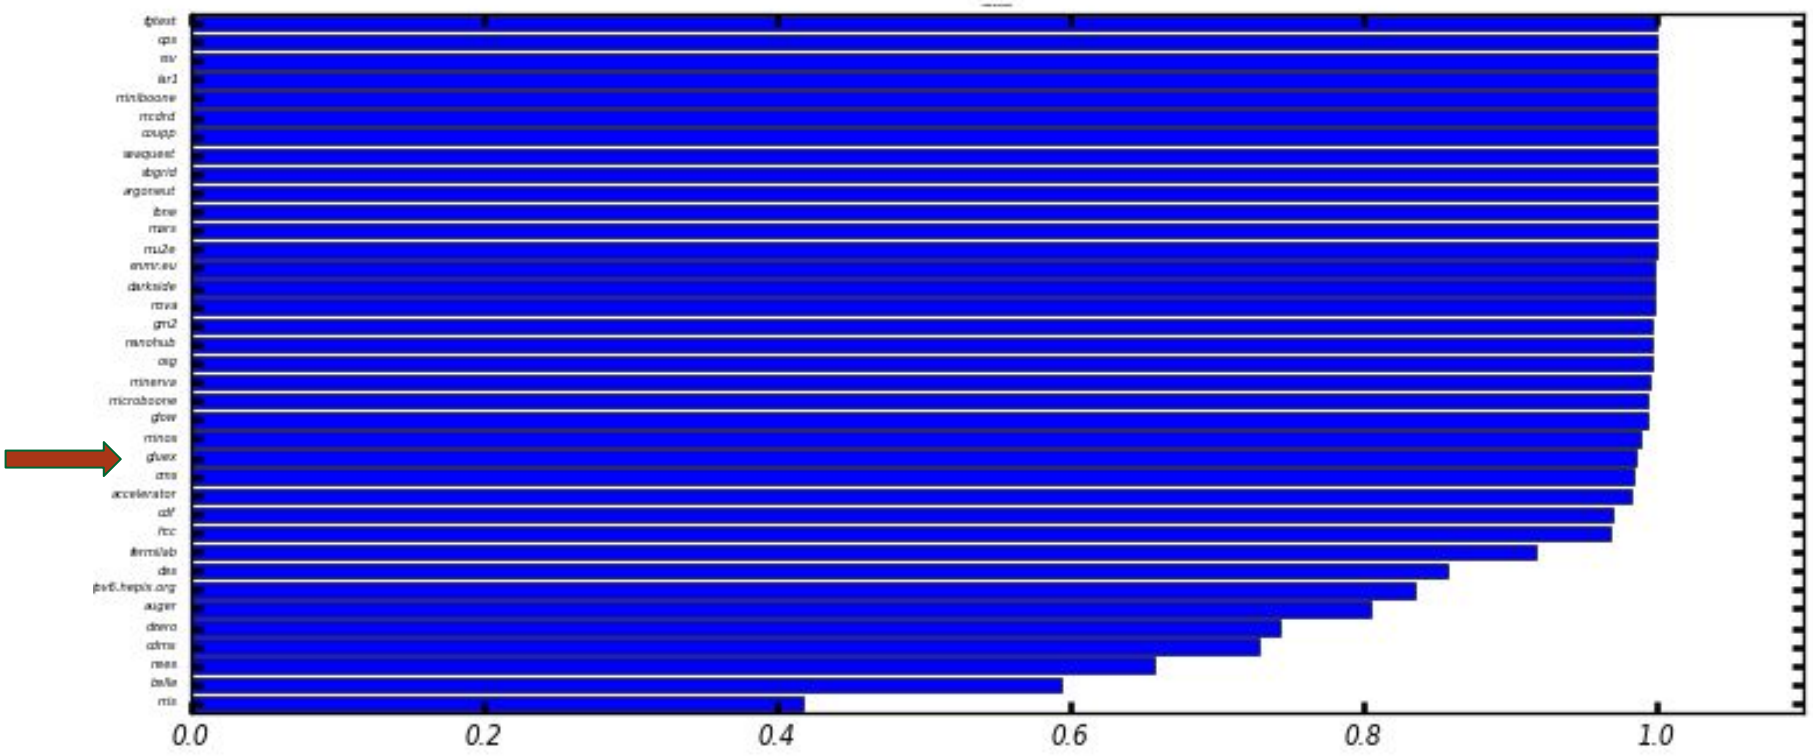

Richard Jones, Jefferson Lab Computing Round Table, July 11, 2017

Gluex activity on osg 2014-2016

Open Science Grid

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

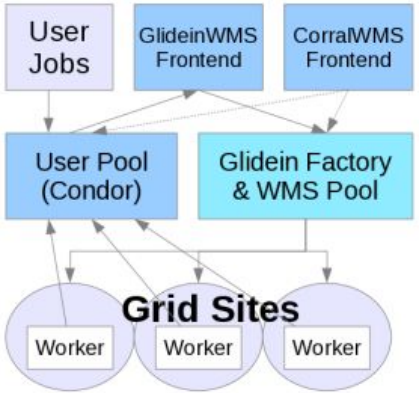

- ➢ located at JLab
- ➢ supported by JLab IT staff
- ➢ GlueX to be among the first users
- ➢ only out-flow of work is currently envisioned
- ➢ server configuration recommended, tested by OSG expert
- ➢ server installed, configured in 2Q 2017, testing by GlueX is now underway.

Gluex @ – the reboot

New infrastructure for osg @ jlab:

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

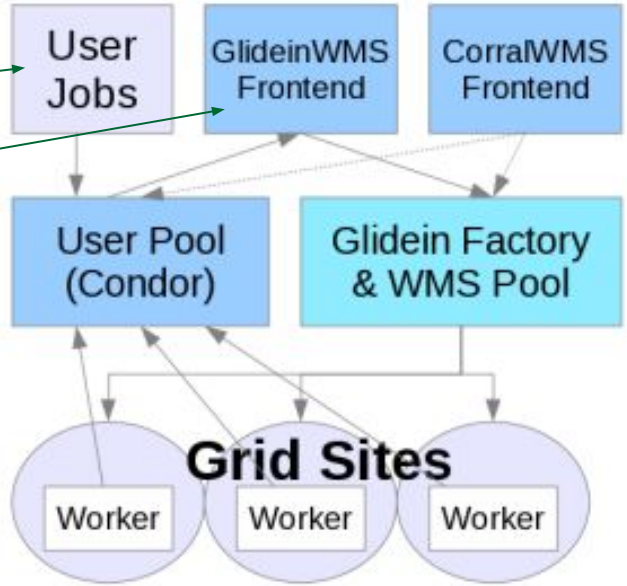

Gluex @ – the reboot

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**

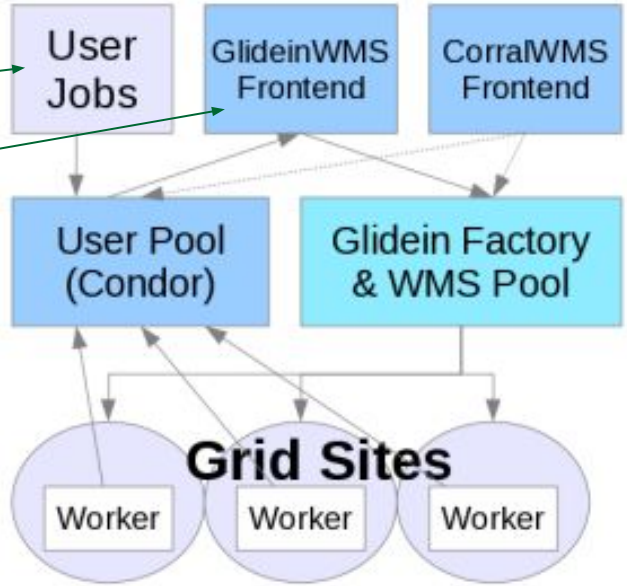

GlueX @  - opportunity cost

- 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

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

Backup slides

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

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

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

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.