GlueX Simulation on the Open Science Grid

lessons learned from the Data Challenge December, 2012

Richard Jones, University of Connecticut

Purpose

- Test the current simulation and reconstruction tools
  1. bggen – pythia-based background Monte Carlo
  2. hdgeant – geant3-based physics simulation, base detector
  3. mcsmear – efficiency and resolution models
  4. hd-ana – reconstruction of tracks, neutrals
  5. REST plugin – summary of reconstruction results
- Develop the ability to manage production and data storage at rates approaching GlueX demands
- Produce a large sample of background simulation data

startup running at 10^7 tags/s, 2000 events/s to tape
goal: 10 billion events, **60 days** at startup intensity

GlueX Context

- steady-state cpu goal (from SW readiness review June, 2012)

  9000 Intel cores (2012) – includes only offline needs

  not all cores are equal: 1 Intel core == 2.x AMD core

- plan combines lab and member university resources to achieve this goal
- GlueX is pursuing grid technology to enable aggregation of distributed resources for GlueX production
- Gluex exists as a Virtual Organization on the Open Science Grid since 2009 – result of Physics Information Frontier grant PHY9876044.

OSG Context

- Open Science Grid – founded 2004
- primary driver supporting LHC experiments in N/S America
- over 75,000 cores, running a distribution of Linux
- sites at 72 institutions including 42 universities, 90 sites (US, Brazil)
- centrally managed and operated by full-time staff (GOC @ I.U.)

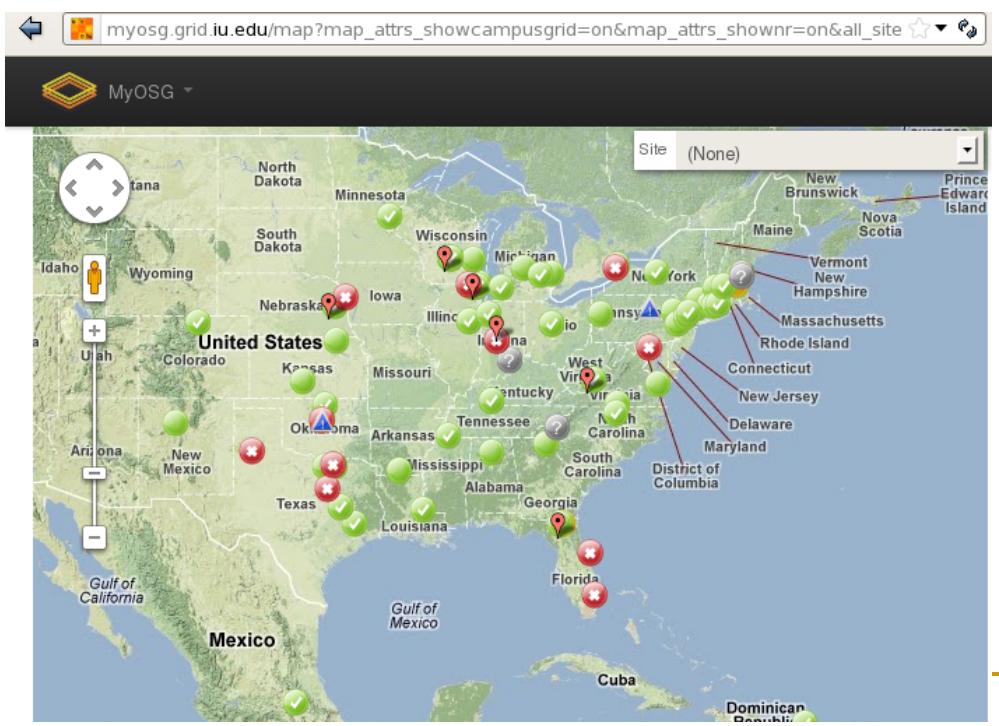

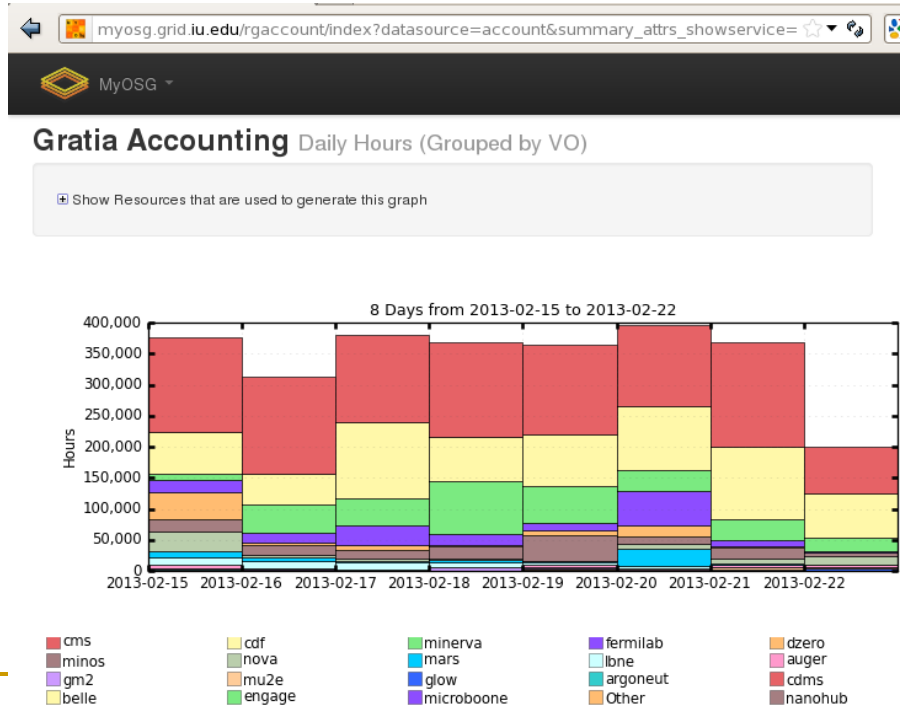

GlueX Data Challenge 1.1

- total of 5,561,650000 events *successfully* generated
  - 4,240,600,000 on the OSG
  - 959,950,000 at Jefferson Lab
  - 361,100,000 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

Achievements

- 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

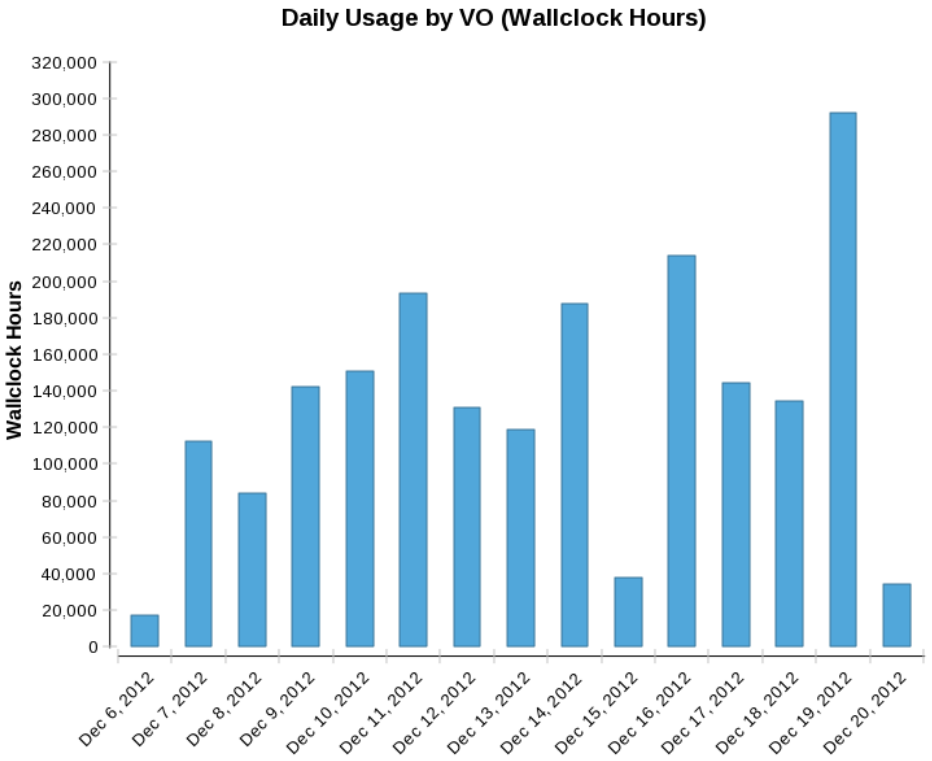

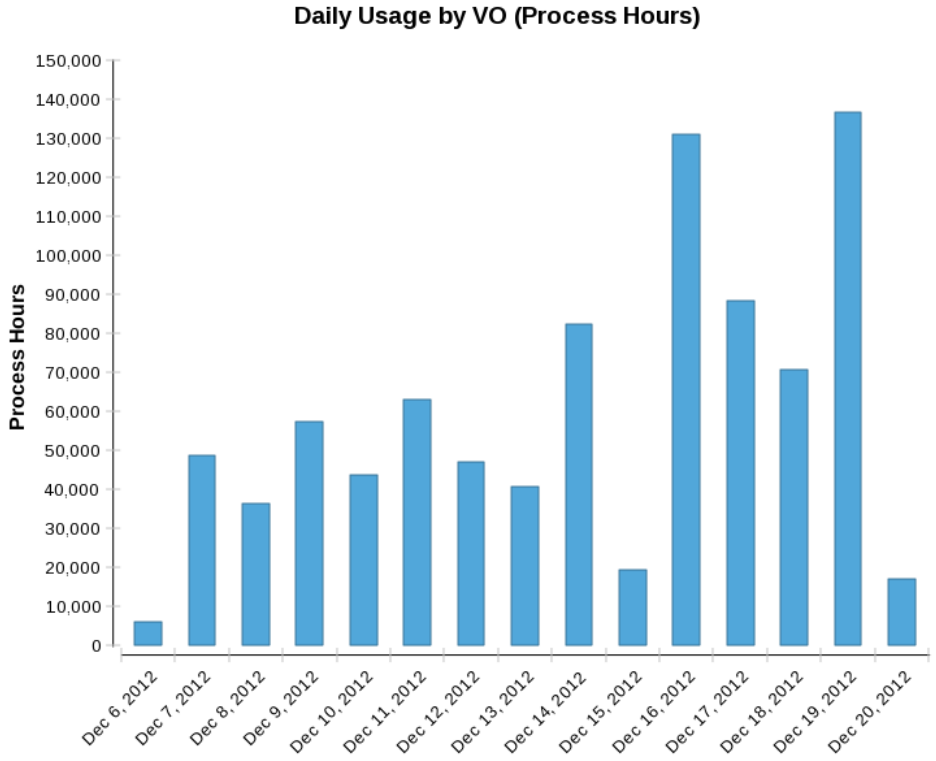

Production inefficiency

- 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.

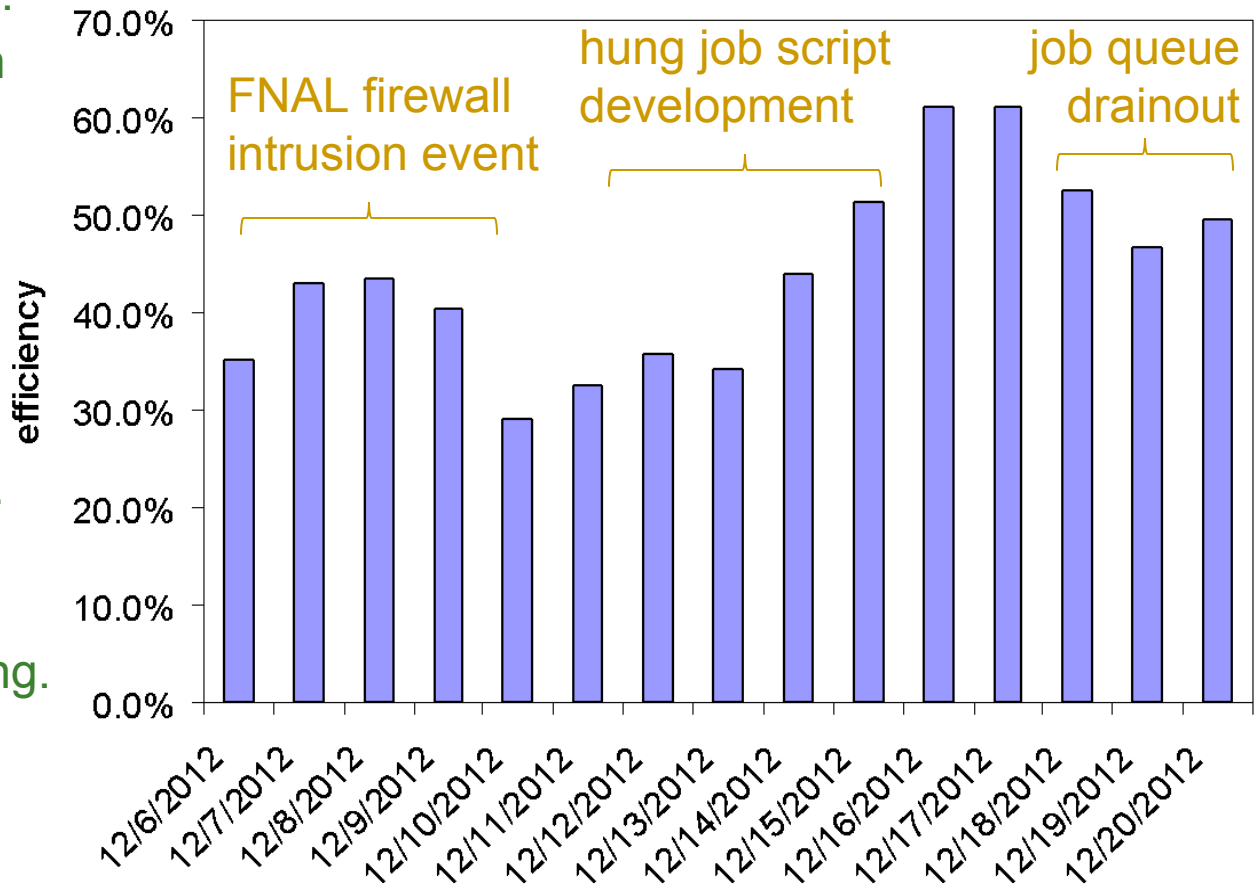

Path for growth

- Congratulations to Gluex from OSG production managers
- R. Jones asked to present to OSG Council, Sept. 2012
  - Council members inquired about GlueX computing plan
  - Response: total scale (9000 cores x 5 years approved)
    - Q: How much might be carried out on OSG?
    - A: Up to 5000 x 300 x 24 hrs / yr = 36M hr/yr
  - Follow-up: By when? What is your schedule for ramping up resources?
  - Response: by 2117, so far there is no detailed roadmap.

- Step 1: move existing resources into the grid framework
- Step 2: carry out new data challenges to test effectiveness
- Step 3: devise a plan for growing resources to the required level

Access and support for Gluex users

- Support for resource consumers (10 users registered)
  - howto get a grid certificate — Getting a Grid Certificate
  - howto access data from DC
  - howto test your code on osg — HOWTO get your jobs to run on the Grid
  - howto run your skims on osg
- Support for resource providers (UConn, IU, CMU, …?)
  - NOT a commitment to 100% occupation by OSG jobs
  - OSG site framework assumes that the local admin retains full control over resource utilization (eg. supports priority of local users)
  - UConn site running for 2 years, new site at IU being set up
  - MIT members have shared use of local CMS grid infrastructure
  - Potential interest to configure CMU site as a Gluex grid site

Case for further challanges

- Existing data set is still only 30 days @ low intensity.
- Existing simulation is missing features: em background, noise hits in BCal, hadronic pile-up.
- Most obvious bugs in framework have been identified in dc1.1, should verify solutions and look for new sub-leading effects.
- New simulation based on G4 is expected by Fall, will enable better estimates of material effects, hadronic backgrounds.
- As new sites are added, demonstrating ramp-up of capability to do GlueX simulations on the grid.