



Open Science Grid

Production Analyst's view of IF on OSG

Bo Jayatilaka
Fermilab

Intensity Frontier Workshop: OSG AHM

Northwestern University

~~March 23, 2015~~

March 24, 2015



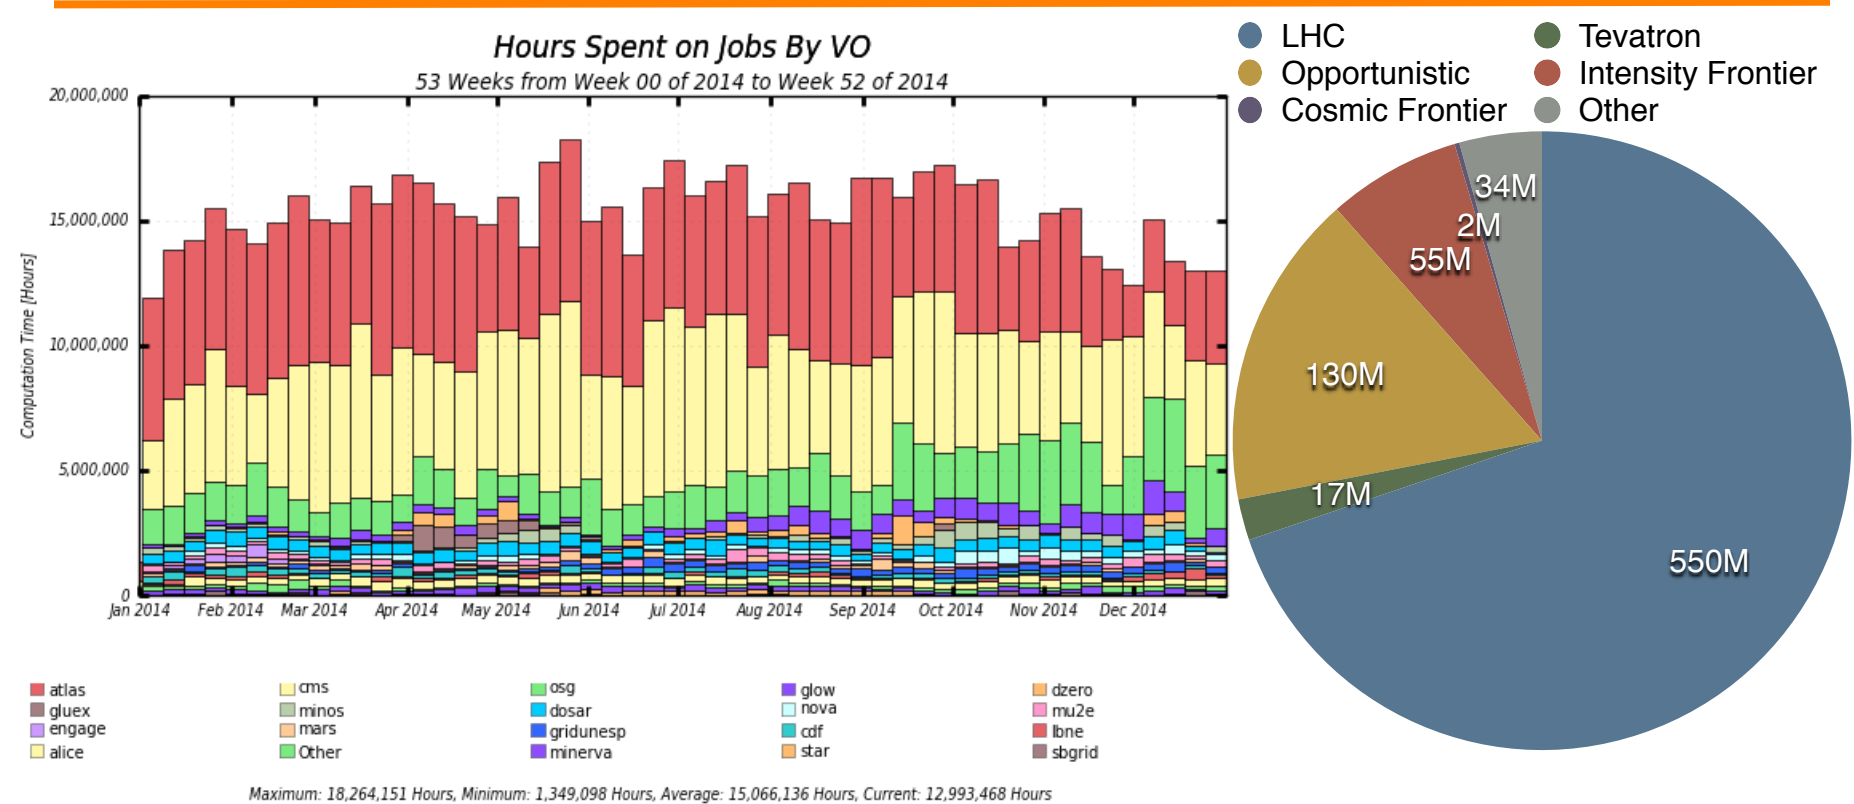


It's now Tuesday

If you were here Monday, skip to slide 10



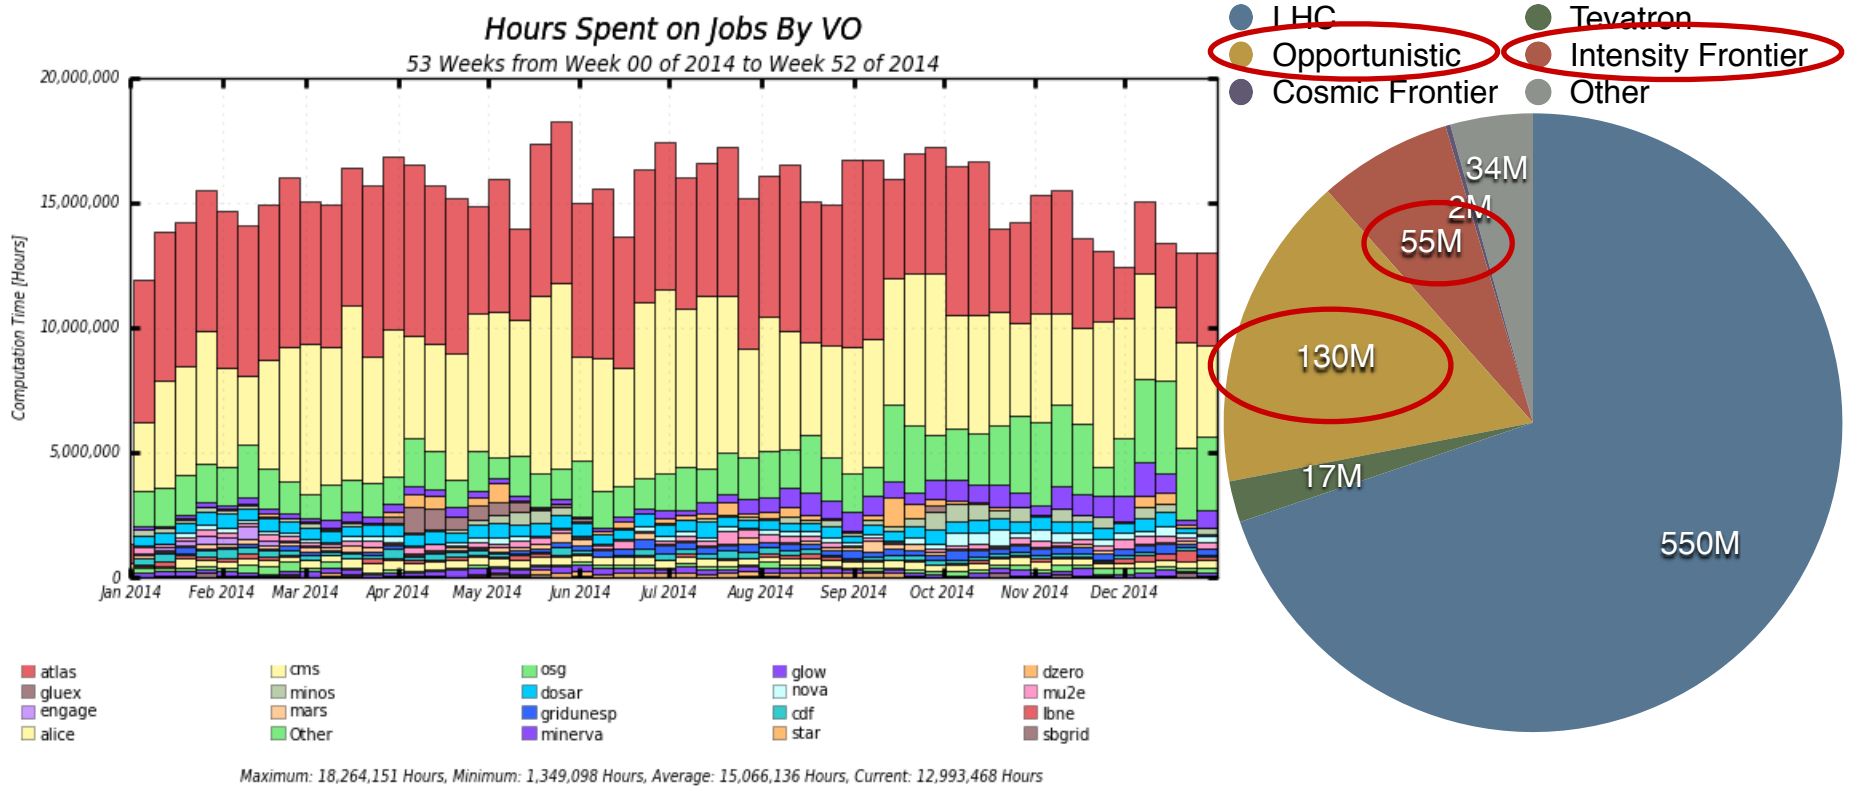
OSG as a whole



- **~800M wall hours**
 - 541M from ATLAS and CMS (~67%)




OSG as a whole



- ~800M wall hours
 - 541M from ATLAS and CMS (~67%)



Intensity Frontier on the OSG



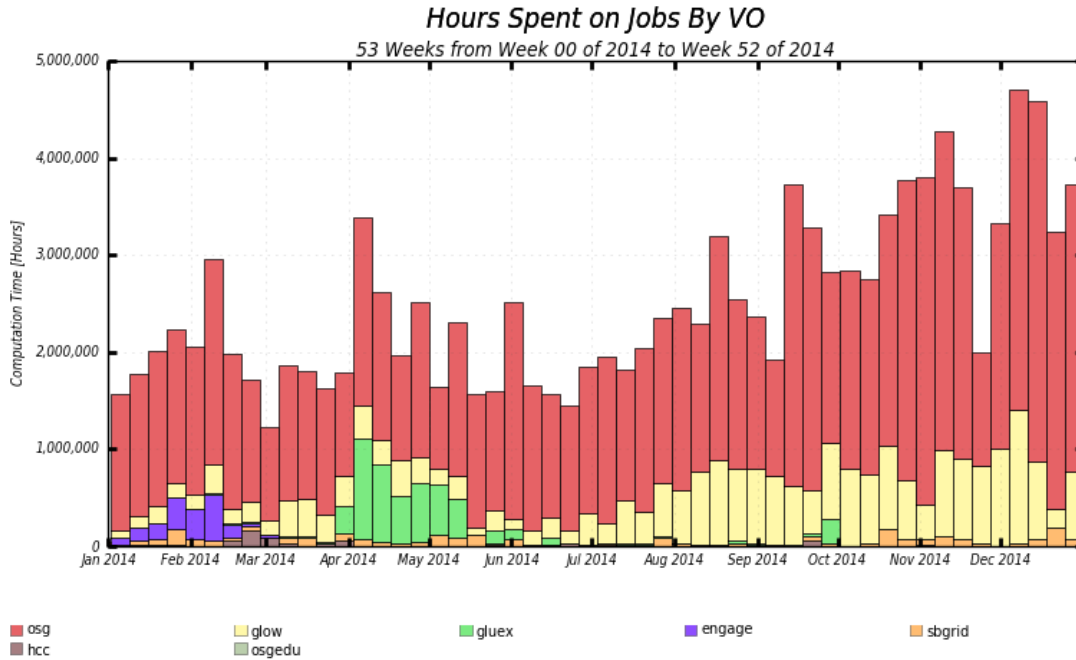
- FIFE* experiments ran a total of **44M** wall hours accounted by Gratia in 2014
 - **42.6M** of those were run at Fermilab
 - Of the 1.64M off-site hours, 1.62M were by NOvA
- Other intensity frontier experiments: LBNE (5.2M hours), Belle (604k hours)

*"FIFE" = *argoneut, cdms, coupp, darkside, gm2, lar1, lar1nd, lariat, microboone, minerva, miniboone, minos, mu2e, nova, sequest,*




OSG as a sharing ecosphere

Open Science Grid



Maximum: 4,699,677 Hours, Minimum: 45,027 Hours, Average: 2,457,438 Hours, Current: 3,728,689 Hours

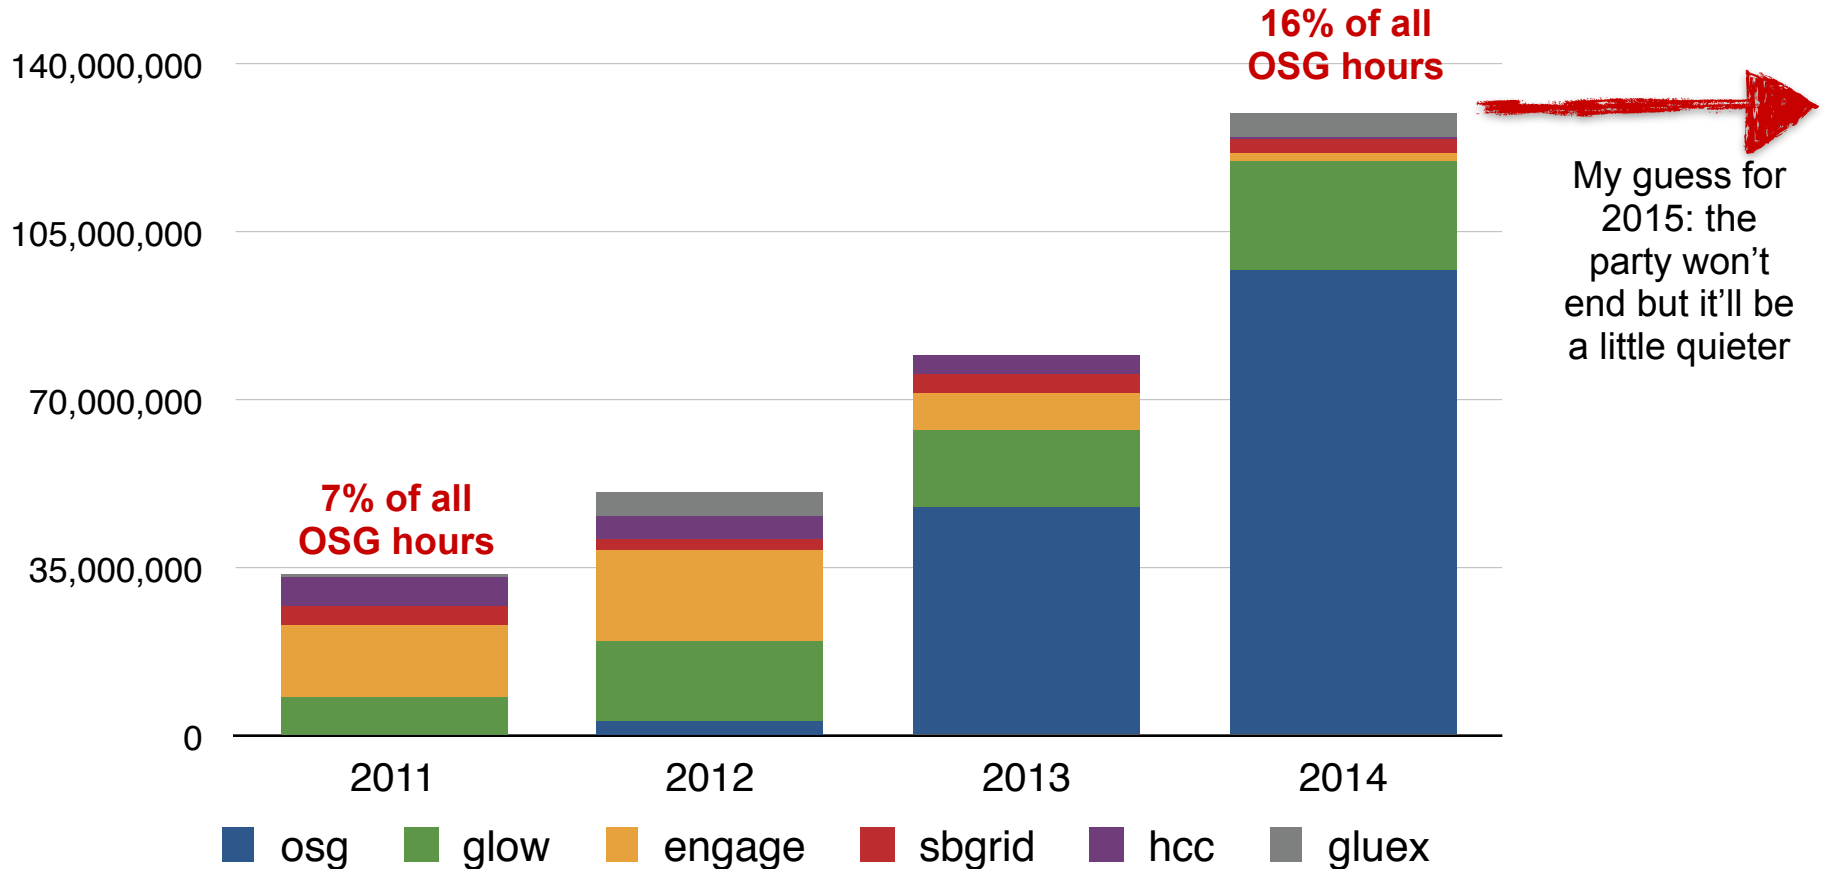


- OSG is committed to sharing any potentially idle resources
 - Update to bylaws: *Consortium members recognize that the OSG is a sharing ecosystem and strive to maximize the sharing of computing resources, software, and other assets to enable science.*
- Primarily opportunistic VOs (osg, glow, gluex, engage, sbgrid, hcc) received **130M wall hours** in 2014
 - 16% of all OSG hours



Growth of the opportunistic ecosphere

>50% growth each year





Where do these hours come from?

Open Science Grid

Top 20 Opportunistic Sites on the OSG in 2014

Site	engage	glow	hcc	osg	sbgrid	Total Opportunistic	Total hours	Percent opportunistic
CIT_CMS_T2		3,416,415	29,763	9,102,039	511,981	13,060,198	28,180,174	46%
FNAL_FERMIGRID	16,874		108,624	12,711,968		12,837,466	82,519,460	16%
UCSDT2	485,006	1,253,139	13,972	10,713,436	252,019	12,717,572	25,256,162	50%
USCMS-FNAL-WC1	136	822,270	5,268	11,073,515	281,014	12,182,203	91,395,332	13%
Tusker	284,971	3,369,581	24,104	7,490,094		11,168,750	14,751,109	76%
Nebraska	35,966	3,333,533	91,792	6,680,426		10,141,717	35,277,858	29%
MWT2	454,547	2,282,363	36,151	6,219,981		8,993,042	61,095,144	15%
MIT_CMS		57	78,228	4,585,284	809,369	5,472,938	24,356,265	22%
Sandhills	23,112	1,748,946	21,573	2,034,164		3,827,795	4,232,466	90%
Crane	66,077	1,492,771	11,247	1,873,012		3,443,107	3,977,168	87%
Purdue-Hadoop		644,142	8,559	2,385,959	26,429	3,065,089	10,888,174	28%
UConn-OSG	76,977	1,081,891	43,159	1,608,189		2,810,216	3,174,455	89%
AGLT2	18,473	1,033,200	10,168	1,327,397	219,954	2,609,192	35,204,179	7%
GridUNESP_CENTRAL	5,587	474,813	5,062	1,974,711	61,271	2,521,444	12,784,909	20%
GLOW		109,332	14,679	2,233,313		2,357,324	27,170,835	9%
NWICG_NDCMS				2,281,689		2,281,689	2,728,851	84%

- Apart from Fermigrid, IF experiments made **no substantial use** of any of these sites



What we'd like to see

- Intensity frontier experiments take full advantage of opportunistic resources on the OSG
- Promote the **Fermilab VO** as an equal priority opportunistic VO as OSG and GLOW
 - Individual experiment VOs still used at sites which have a special arrangement with a particular experiment (e.g., μ BooNE for SLAC, NOvA for Harvard)
- Have Fermilab IF experiments submit offsite and opportunistically using the Fermilab VO
 - Especially useful for low-IO jobs (MC generation, CPU-intensive analysis jobs, not requiring reading from SAM)
 - Have the end-analyzer/user see a flat landscape of opportunistic resources to run their jobs on



Production Support

- Newly formed area within OSG
 - Effectively “user” support but at the **VO and site level**
 - Formed out of previous User Support group (lead by BAJ)
- Mission for Production Support
 1. Grow the opportunistic ecosphere of the OSG
 2. **Help experiments/VOs access opportunistic resources**
 3. Be a catalyst for projects and technologies that help VOs run opportunistically
- Bottom line: we encourage VOs to use opportunistic resources and will help make that a practical reality in any way we can



Why this isn't easy (from the OSG side)

- OSG sites are not a monolithic block (even within an experiment/organization)
 - Goal: have at least every site supporting the OSG VO also support the Fermilab VO
 - We are ~50% of the way there today
 - Every site will be accompanied by teething pains
 - OSG Production Support and FIFE should make these teething pains as invisible as possible to both sites and end-users
- The environment simply can't perfectly replicate the FermiGrid experience
 - But it should be close
- **Data movement** (~10GB and above) is a challenge
 - StashCache project (caching via XRootD) would provide one solution



Data access

- How do the (Fermilab) IF experiments handle data storage and access?
 - *Very well for jobs running at Fermilab*
 - Enstore tape system (49PB) with dCache
 - Most user jobs access data via a system called SAM (updated version of data access protocol used at Tevatron experiments)
- Above is for **event data**
 - Plenty of computing that experiments do that don't use event data (e.g. Monte Carlo simulation)
- These jobs still will need to access **auxiliary data** in many cases



Auxiliary data

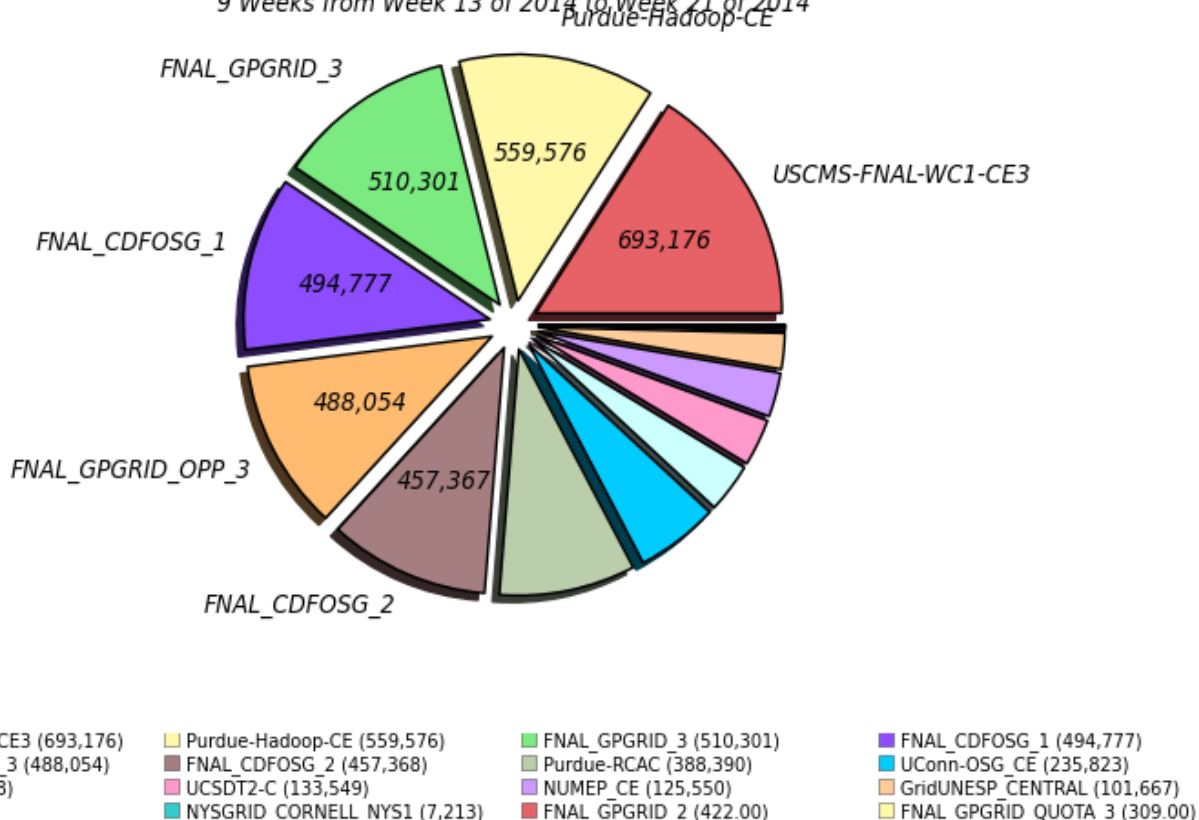
- Event data tends to be different for each job
 - Perhaps the same dataset but events themselves tend to be processed once per job
- **Auxiliary data**
 - Data that are shared amongst numerous jobs
 - Cache hit rate is low (not a shared object library but data)
 - *e.g.* GENIE flux files, minbias overlay, event samples for MVA training
 - CVMFS is a poor technical match for this (but has been used!)
 - Dataset size can be $O(100\text{GB})$
 - POSIX access is useful for accessing such files



GlueX challenge by the numbers

Wall Hours by Facility (Sum: 4,344,368 Hours)

9 Weeks from Week 13 of 2014 to Week 21 of 2014




- 4.1M hours opportunistically over ~6 weeks



Can we do this for IF experiments?

WMS Hours Spent on Jobs By Facility (Glidein)
14 Days from 2015-03-10 to 2015-03-23



MIT_CMS SU-OG-CE Crane-CE1 OSC_OSG_CE MWT2
GLOW BNL_ATLAS_1 Tusker AGLT2 UCSDT2

Maximum: 55,571 , Minimum: 1.37 , Average: 11,870 , Current: 55,571

- mu2e usage outside of Fermilab in the past two weeks



Conclusion

- The OSG now provides ~140M hours/year of opportunistic cycles to users from a variety of research areas (see Wednesday's talk "Science on the OSG in 2014")
- OSG would like to see the intensity frontier community take advantage of these resources
 - Fermilab IF community can take advantage of common submission tools and a single VO operating on many OSG sites
 - FIFE project offers common tools and methods for this access
 - Newly formed Production Support group will assist any VO seeking to utilize opportunistic resources.




Open Science Grid

Backup



Opportunistic in 2015 so far

Hours Spent on Jobs By VO
11 Weeks from Week 00 of 2015 to Week 11 of 2015



Maximum: 4,998,485 Hours, Minimum: 2,894,461 Hours, Average: 3,625,110 Hours, Current: 3,153,405 Hours

Wall Hours by VO (Sum: 39,876,209 Hours)
11 Weeks from Week 00 of 2015 to Week 11 of 2015



- Total of **39.9M** hours to-date in 2015 for primarily opportunistic VOs
 - ~24% of all OSG hours