



Open Science Grid

User Support Report

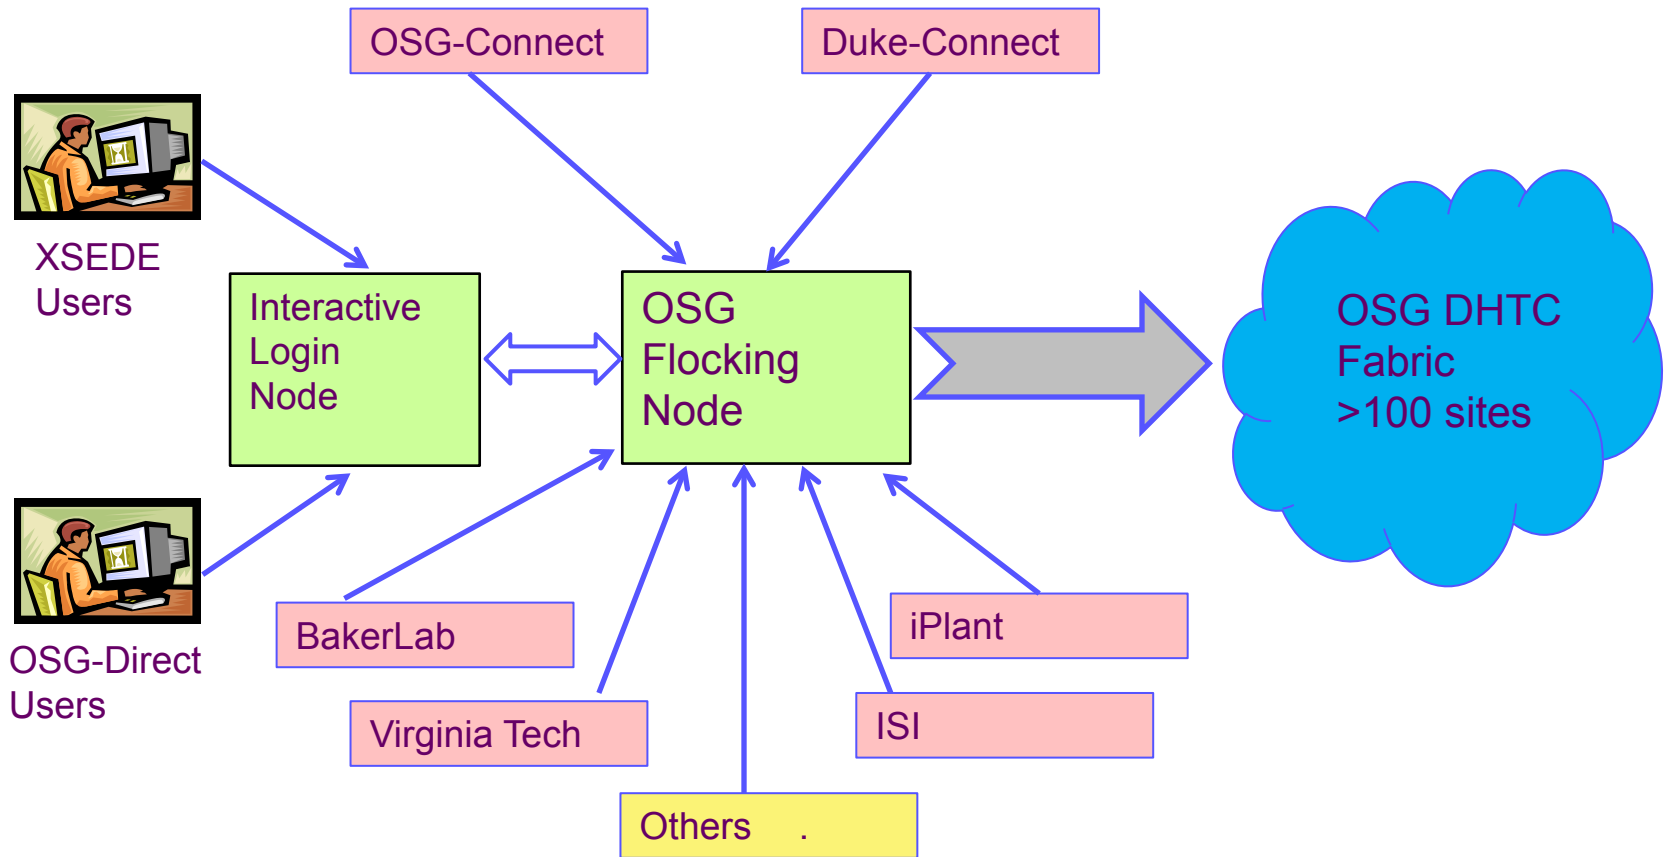
Bo Jayatilaka, Mats Rynge, Chander Sehgal

OSG Council Meeting

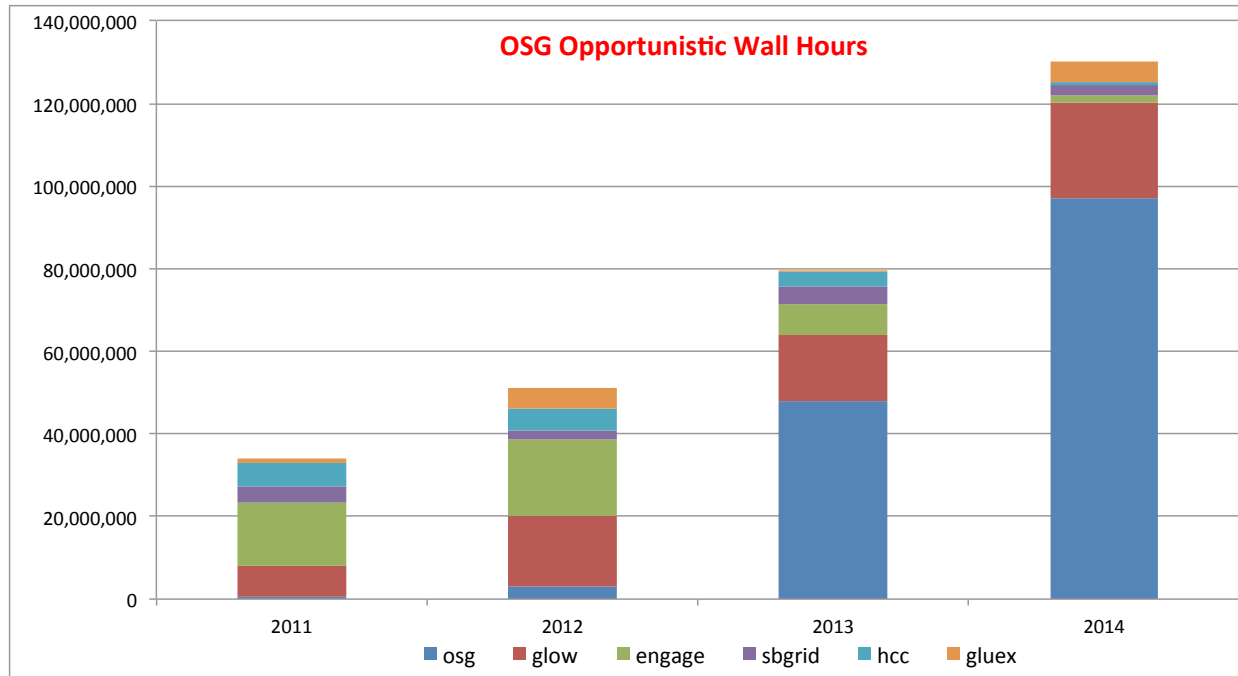
January 14, 2015



OSG Flocking



Opportunistic Trends



- Opportunistic VOs received **130M** wall hours in 2014
 - 16% of total hours in the OSG
 - Continues 50%+ annual growth rate since 2011




Opportunistic Trends

Opportunistic breakdown

VO	2011	2012	2013	2014
osg	400,568	3,168,025	47,931,106	97,009,409
glow	7,801,442	16,924,657	15,896,802	23,223,587
engage	15,011,328	18,610,782	7,673,167	1,680,835
sbgrid	4,046,324	2,271,005	4,194,011	2,558,192
hcc	5,741,379	5,017,649	3,693,243	639,626
gluex	998,031	5,257,905	106,317	4,942,329
Opportunistic Hours	33,999,072	51,250,023	79,494,646	130,053,978
Annual Growth Rate	N/A	51%	55%	64%
All OSG Hours	510,671,422	651,365,203	649,812,810	798,145,827
Opportunistic Percent	7%	8%	12%	16%

OSG VO breakdown

	2011	2012	2013	2014
OSG-XD	0	2,129,345	24,948,630	45,326,731
OSG-Connect	0	0	541,167	6,478,164
OSG-Direct	400,568	1,038,680	22,441,309	45,204,514



- Availability in opportunistic cycles fueled by lulls in LHC computing, onboarding of primarily opportunistic sites (e.g. Syracuse) and upgrade of OSG VO infrastructure
- Expect new primarily opportunistic sites in 2015 (e.g., FIU)

OSG VO scope and usage

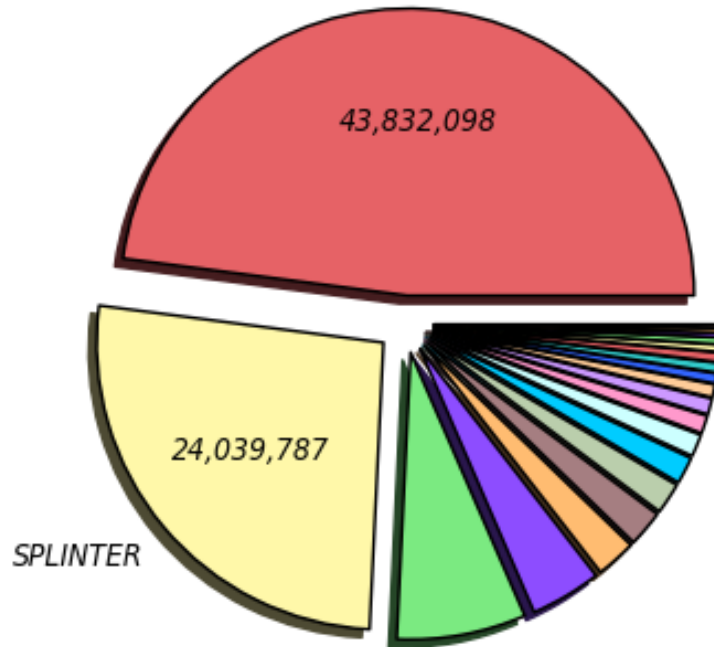
- The OSG VO does not own any computing resources and only exists to harvest unused cycles at OSG sites (Opportunistic cycles) and make them available to researchers who are not already affiliated with an OSG VO.
- For the 12 months ending 31-Dec-2014, the OSG VO harvested 92.5M hours (from sites by using gWMS) and delivered them to various submit hosts to enable the computing of researchers

Submit Host	Wall Hours
XD-Login (XD and OSG Direct)**	84,896,342
OSG-CONNECT **	6,446,676
BakerLab	1,139,098
ISI	68,308



OSG VO Wall Hours

Wall Hours by VO (Sum: 91,437,184 Hours)
365 Days from Week 00 of 2014 to Week 52 of 2014
TG-IBN130001



- TG-IBN130001 (43,832,099)
- PROTFOLDING (2,063,096)
- PHENO (1,149,479)
- UMICH (508,368)
- TG-PHY120014 (332,824)

- SPLINTER (24,039,788)
- ALGDOCK (1,866,887)
- UPRRP-MR (918,016)
- IU-GALAXY (491,700)
- KNOWLEDGESYS (218,451)


- DUKE-QGP (6,661,506)
- DETECTORDESIGN (1,608,401)
- TG-DMR130036 (816,052)
- CONNECTTRAIN (476,917)
- EIC (215,995)

- RIT (3,655,333)
- Other (1,390,647)
- ICECUBE (661,850)
- ORBITER (357,187)
- ERRORSTUDY (172,590)



Top sites (December)


SITE	INSTITUTION (ORG)	RECENT	RECENT	PRIOR	PRIOR
USCMS-FNAL-WC1	Fermi National Accelerator Laboratory (USCMS Tier 1)	1	3,642,690	1	3,611,416
UCSDT2	University of California San Diego (USCMS Tier 2)	2	2,830,802	2	2,552,017
MIT_CMS	Massachusetts Institute of Technology (USCMS Tier 2)	3	2,337,065	6	1,279,463
CIT_CMS_T2	California Institute of Technology (USCMS Tier 2)	4	2,206,811	5	1,287,979
FNAL_FERMIG RID	Fermi National Accelerator Laboratory (Campus Grid gateway)	5	2,093,407	3	1,622,445
Nebraska	The Holland Computing Center at the University of Nebraska-Licoln (USCMS Tier2)	6	1,861,905	8	949,050
BNL-ATLAS	Brookhaven National Laboratory (USATLAS Tier 1)	7	1,019,035	33	21,811
Tusker	The Holland Computing Center at the University of Nebraska-Omaha (HCC Campus Grid)	8	967,614	9	794,830
MWT2	University of Chicago, Indiana University, University of Illinois, Urbana-Campaign (USATLAS Midwest Tier 2)	9	816,660	4	1,589,290
Purdue-Hadoop	CMS Hadoop cluster at Purdue University	10	717,378	37	424





Increased Contribution from BNL

Cumulative Hours Spent on Jobs By Facility
365 Days from Week 00 of 2014 to Week 52 of 2014



■ BNL_ATLAS_1 (957,276) ■ BNL_ATLAS_2 (131,871) ■ BNL_ATLAS_3 (1,247)

Total: 1,090,395 Hours, Average Rate: 0.03 Hours/s



XD Allocations 1/1/14-12/31/14

Open Science Grid

Project Name	PI	Institution	Field of Science	Wall Hours	Allocation
TG-IBN130001	Donald Krieger	University of Pittsburgh	Biological Sciences	43,832,099	2,000,000
TG-DMR130036	Emanuel Gull	University of Michigan	Materials Science	816,052	2,000,000
TG-PHY120014	Qaisar Shafi	University of Delaware	Physics and astronomy	332,824	1,010,000
TG-CHE130091	Paul Siders	University of Minnesota; Duluth	Chemistry	83,353	100,000
TG-CHE140110	John Stubbs	University of New England	Chemistry	67,108	270,000
TG-OCE130029	Yvonne Chan	University of Hawaii; Manoa	Ocean Sciences	36,738	33,120
TG-MCB090163	Michael Hagan	Brandeis University	Molecular and Structural Biosciences	35,589	272,000
TG-CHE140094	John Stubbs	University of New England	Chemistry	27,009	23,000
TG-CHE130103	Jeremy Moix	Massachusetts Institute of Technology	Chemistry	24,461	100,000
TG-IRI130016	Joseph Cohen	University of Massachusetts; Boston	Information Robotics and Intelligent	20,401	200,000
TG-TRA120014	Pol Llovet	Montana State University	Evolutionary Sciences	19,479	200,000
TG-IBN130008	Jorden Schossau	Michigan State University	Biological Sciences	14,411	200,000
TG-OCE140013	Yvonne Chan	University of Hawaii; Manoa	Ocean Sciences	6,055	60,000
TG-DEB140008	Robert Toonen	University of Hawaii; Manoa	Biological Sciences	4,147	31,000
TG-CHE140098	Paul Siders	University of Minnesota; Duluth	Chemistry	2,263	205,000
TG-DMR140072	Adrian Del Maestro	University of Vermont	Materials Science	2,068	20,000
TG-MCB120070	Joseph Hargitai	Albert Einstein College of Medicine	Molecular and Structural Biosciences	378	100,000
TG-MCB140160	David Rhee	Albert Einstein College of Medicine	Molecular and Structural Biosciences	149	200,000
TG-SEE140006	Sheila Kannappan	University of North Carolina; Chapel Hill	Physics and astronomy	46	20,000
TG-PHY110015	Pran Nath	Northeastern University	Physics and astronomy	37	1,000,000
TG-CDA100013	Mark Reed	University of North Carolina; Chapel Hill	Mathematical Sciences	6	200,000
TG-CCR120041	Luca Clementi	San Diego Supercomputer Center	Computer Science	3	50,000
TG-CCR130001	Ruth Marinshaw	Stanford University	Training	2	200,000
				45,324,676	8,494,120

Research Allocation

Startup Allocation

Educational Allocation

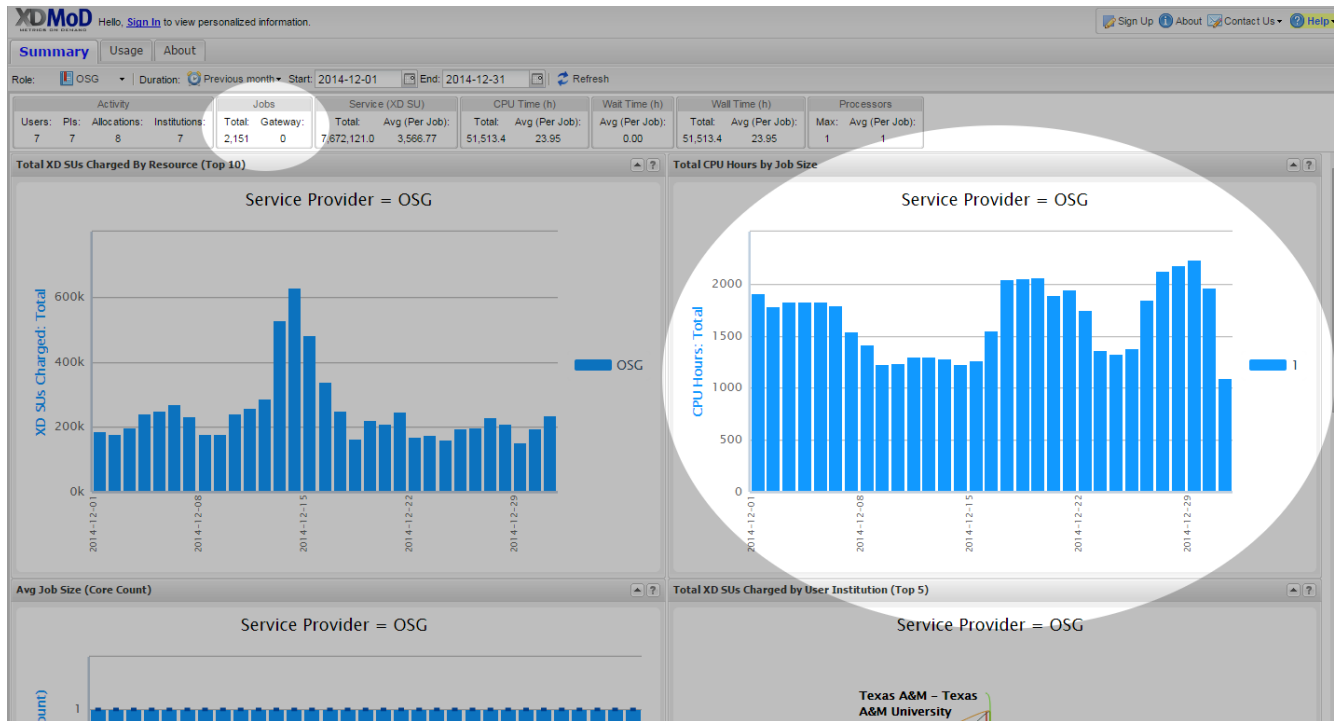
Campus Champion Allocation



OSG Direct Projects

Project Name	PI	Institution	Field of Science	Wall Hours
SPLINTER	Robert Quick	Indiana University	Medical Sciences	24,039,788
Duke-QGP	Steffen A. Bass	Duke University	Nuclear Physics	6,661,506
RIT	P. Stanislaw Radziszowski	Rochester Institute of Technology	Computer Science and Engineering	3,655,333
DetectorDesign	John Strologas	University of New Mexico	Medical Sciences	1,608,401
Pheno	Stefan Hoeche	SLAC	High Energy Physics	1,149,479
UPRRP-MR	Steven Massey	Universidad de Puerto Rico	Bioinformatics	918,016
UMich	Paul Wolberg	University of Michigan	Microbiology	508,368
IU-GALAXY	Robert Quick	Indiana University	Bioinformatics	491,700
EIC	Tobias Toll	Brookhaven National Laboratory	High Energy Physics	215,995
DeerDisease	Lene Jung Kjaer	Southern Illinois University	Biological Sciences	131,937
HL-LHC-TP	Meenakshi Narain	Brown University	High Energy Physics	90,870
BNLPET	Martin Purschke	Brookhaven National Laboratory	Medical Sciences	46,549
SoyKB	Dong Xu	University of Missouri	Plant Biology	45,143
dVdT	Ewa Deelman	University of Southern California	Computer Science and Engineering	29,185
PO-LBNE	Maxim Potekhin	Brookhaven National Laboratory	High Energy Physics	18,334
OSG-Staff	Chander Sehgal	Fermilab	Computer Science and Engineering	4,140
UNC-RESOLVE-photometry	David Stark	UNC Chapel Hill	Physics and astronomy	26
				39,614,768

XDMoD Issues



- Two issues identified: total number of jobs, and total number of CPU hours
- Problems stems from sending summary records from Gratia database to XSEDE central database
- Initiated discussion with two teams: XSEDE accounting (data propagation and central database), and XDMoD developers (visualization)




Conclusions

- Opportunistic availability (and OSG VO utilization) outpacing goals set during last year's review.
- Re-establishing VO forum for opportunistic VOs
 - Currently discussion between OSG and GLOW
 - Regular contact point between User Support and Fermilab IF efforts as well
- Difficult to maintain current growth rate without new users
 - Number of running projects (~30) relatively flat over the past year
 - New users sought via ACI-REFs
 - Increasing access for GLOW and Fermilab VOs across OSG



Backup: Resource Contributions

Cumulative Hours Spent on Jobs By Facility 365 Days from Week 00 of 2014 to Week 52 of 2014



- Other (26,742,617)
- UCSDT2-D (5,253,226)
- CIT_CMS_T2B (4,509,798)
- MWT2_CĒ_UIUC (2,800,812)
- red-gateway2 (2,305,444)
- FNAL_CDFOSG_2 (2,141,763)
- SMU_HPC (1,929,112)

- USCMS-FNAL-WC1-CE3 (11,073,515)
- UCSDT2-C (5,048,875)
- FNAL_GPGRID_3 (4,060,936)
- red-gateway1 (2,424,410)
- NWICG_NDCMS (2,281,689)
- Sandhills (2,034,163)
- Crane-CE1 (1,873,011)

- Tusker-CE1 (7,489,887)
- CIT_CMS_T2 (4,592,239)
- FNAL_GPGRID_OPP_3 (3,946,338)
- FNAL_CDFOSG_1 (2,419,420)
- GLOW-OSG (2,233,353)
- GridUNESP_CENTRAL (1,974,710)

Total: 97,135,329 Hours, Average Rate: 3.08 Hours/s