The State of OSG 2016

Frank Würthwein

OSG Executive Director

UCSD/SDSC

One Billion hours a year

100 Million Core hours
in the past 30 days

Over the last 12 months
200 Million jobs consumed
1 Billion hours of computing
involving **1.5 Billion** data transfers
to move **223 Petabytes**

This aggregate was accomplished by
federating 114 clusters
that contributed 1h to 100M hours each

http://display.grid.iu.edu

In the last 24 Hours	
543,000	Jobs
2,087,000	CPU Hours
5,823,000	Transfers
588	TB Transfers
In the last 30 Days	
16,096,000	Jobs
96,656,000	CPU Hours
155,744,000	Transfers
20,046	TB Transfers
In the last 12 Months	
198,318,000	Jobs
1,087,534,000	CPU Hours
1,525,604,000	Transfers
223,000	TB Transfers

OSG Magic

We create a uniform environment across a heterogeneous set of resources that is distributed globally

Submit locally – Run Globally

OSG supports computing across different types of resources

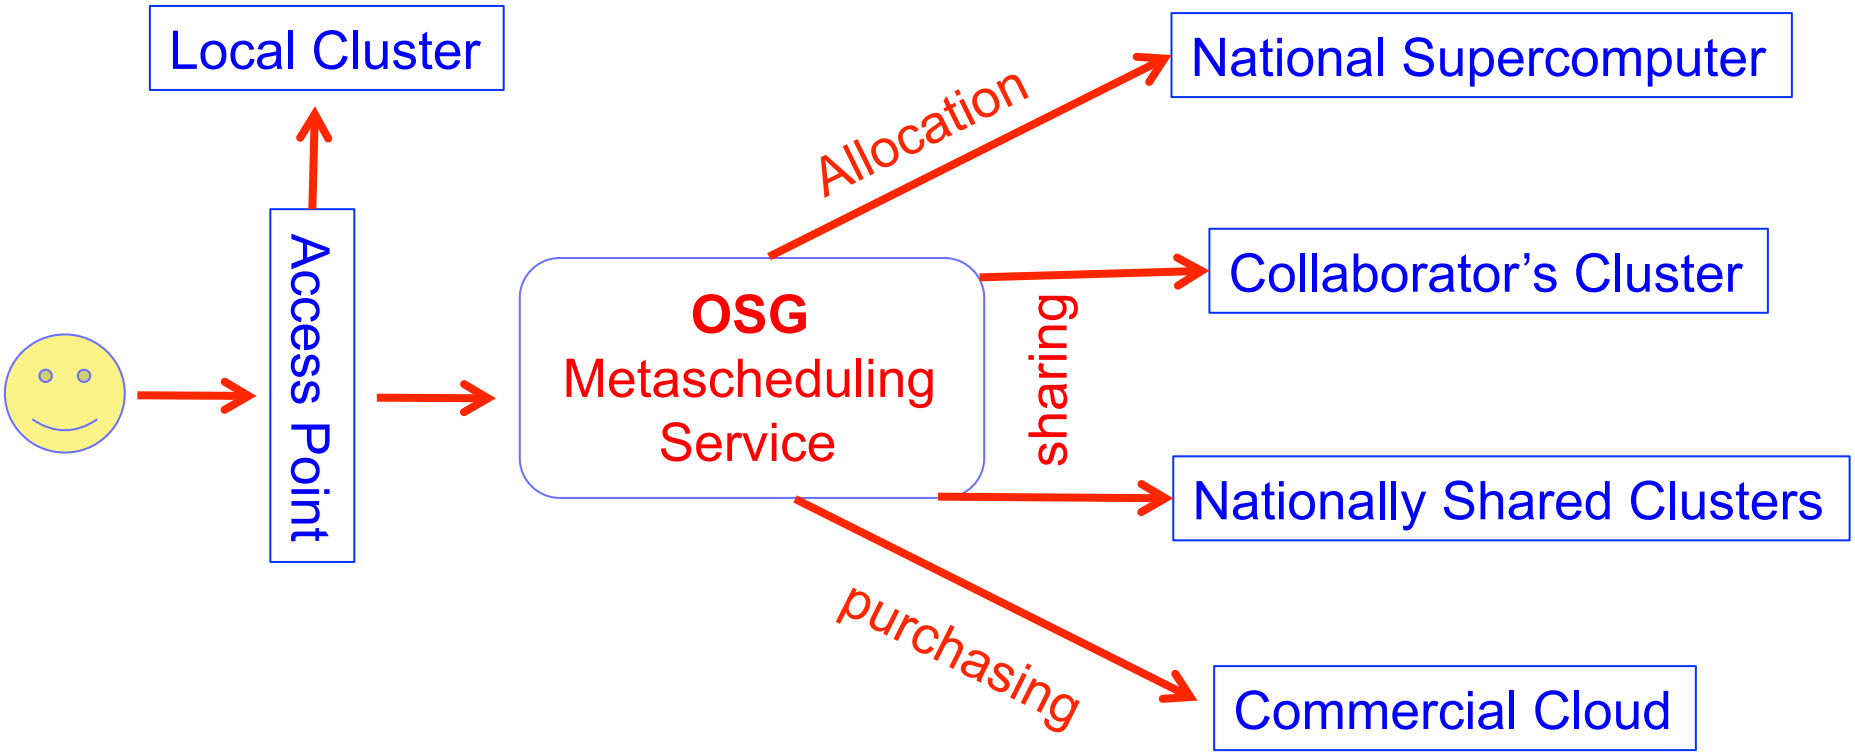

Seamless Integration is the key to our success !!!

OSG is Open to All

- Open to providers at all scales
  - from small colleges to large national labs
- Open to user communities at all scales
  - from individual students to large research communities
    - domain science specific and across many campuses
    - campus specific and across many domain sciences
- Open to any business model
  - sharing, allocations, purchasing
  - preemption is an essential part of operations

Toolset to match diversity of scale

- OSG-Connect
  - OSG hosts the service on OSG hardware
- OSG Cluster in a box
  - OSG manages services on hardware placed inside Campus Science DMZs
- OSG-CE et al.
  - OSG provides software that campuses use to instantiate & operate services

In all cases, seamless integration is key !

OSG since Inception

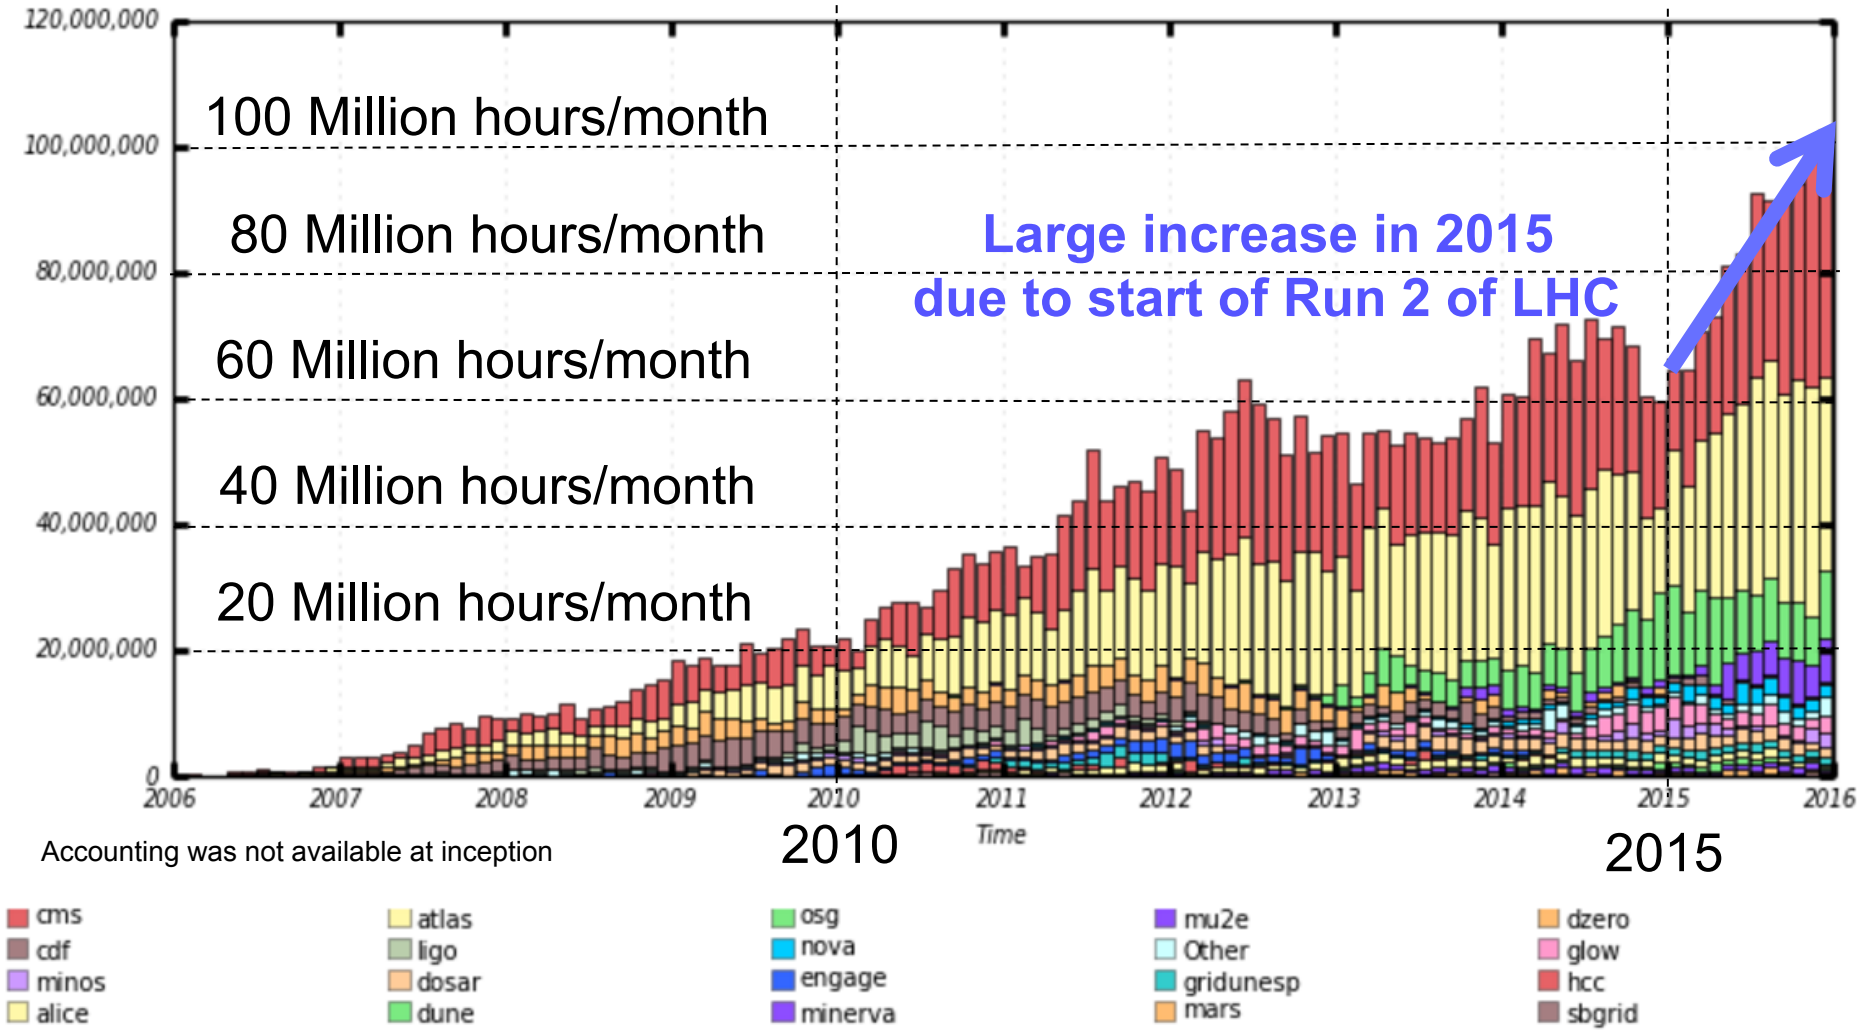

Accounting was not available at inception

OSG since Inception

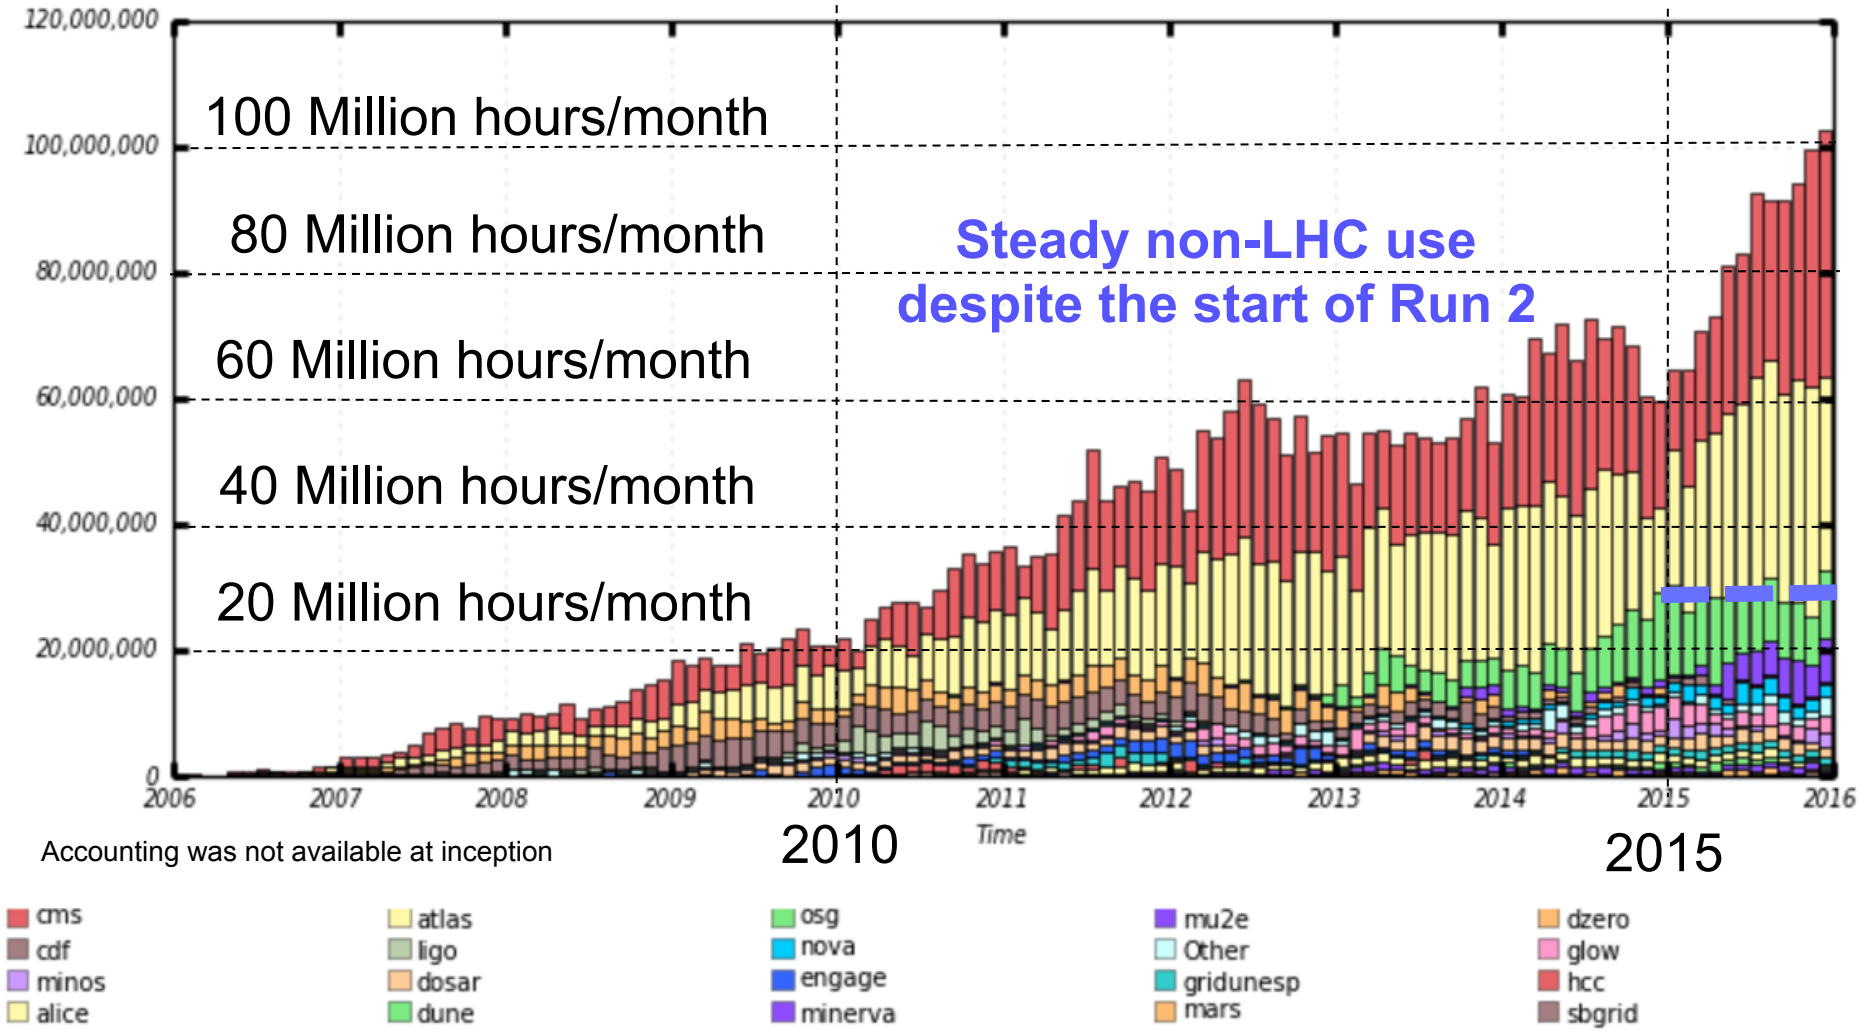

Accounting was not available at inception

Universities and National Labs

This includes Universities that have no LHC involvement

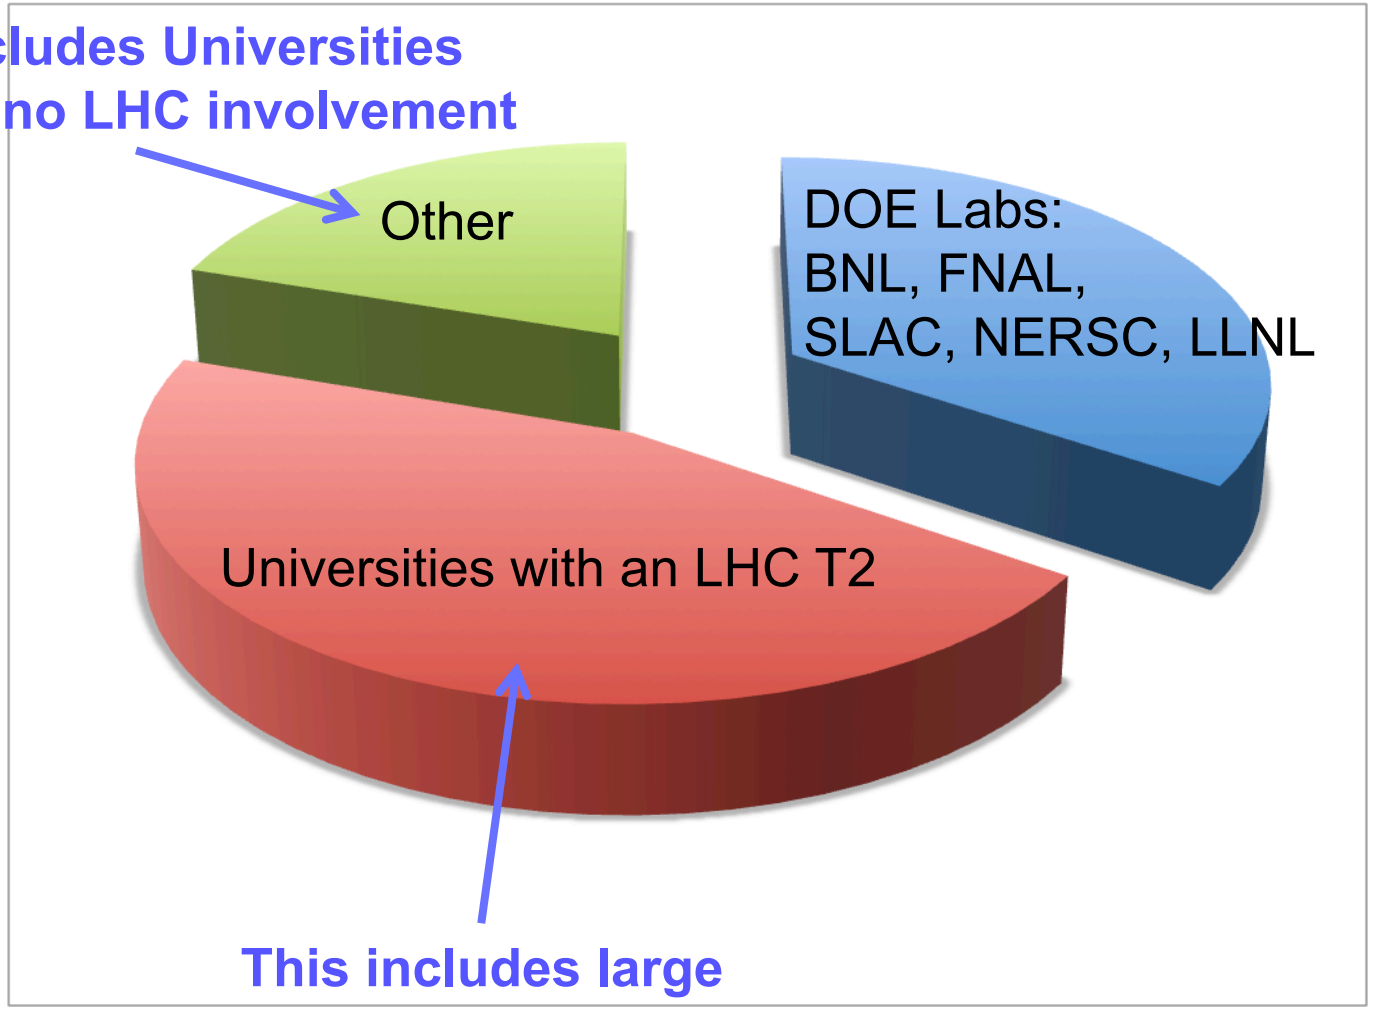

This includes large Campus clusters independent of the LHC.

Science on OSG

OSG Hours 2015 by Science Domain

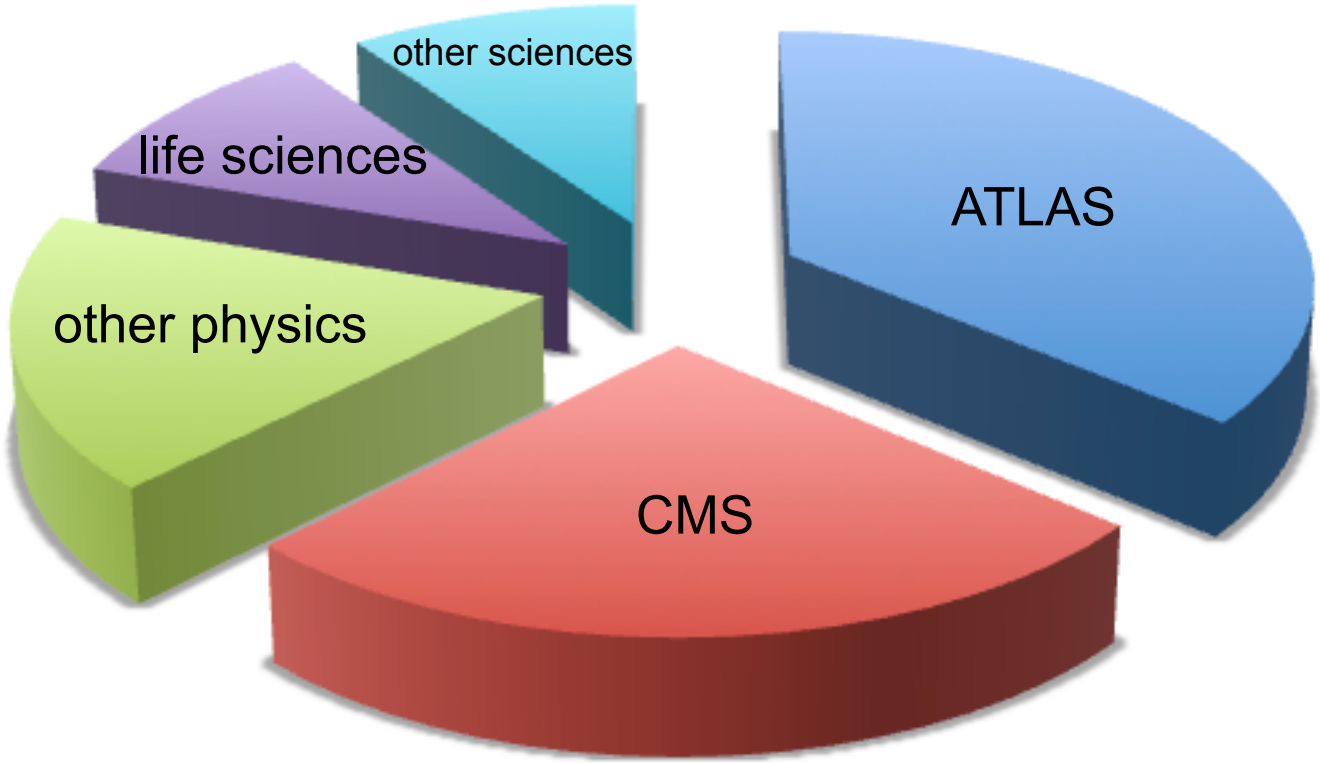

Science other than LHC makes up ~34% of the OSG hours

Science other than Physics makes up ~20% of the OSG hours

LHC continues to be the dominant force in OSG

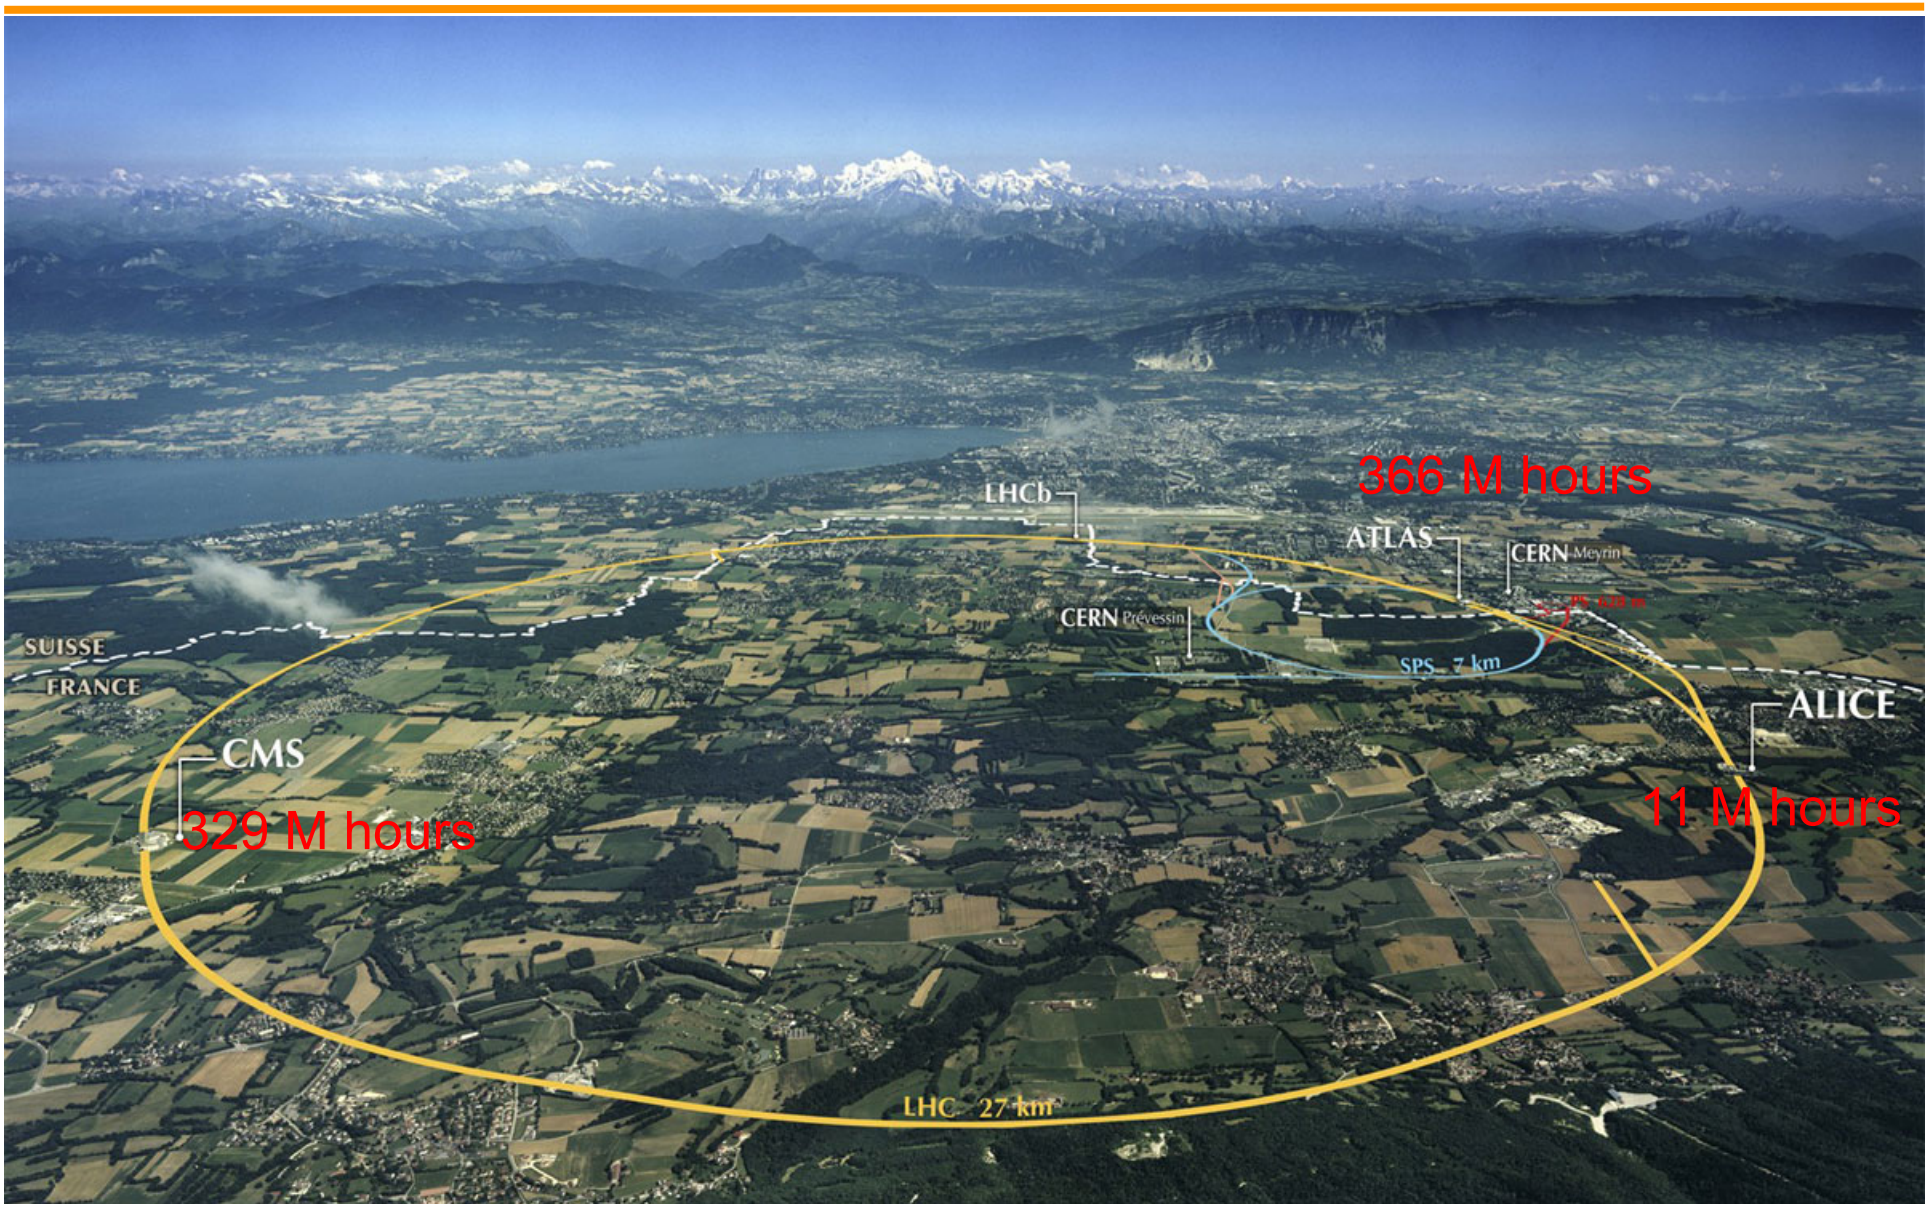

Particle Physics beyond the LHC

Particle physics other than the LHC scaled up from 2.5M to more than 50M within the last year

Experiments range from those built in 1980s to fresh proposals

LArIAT

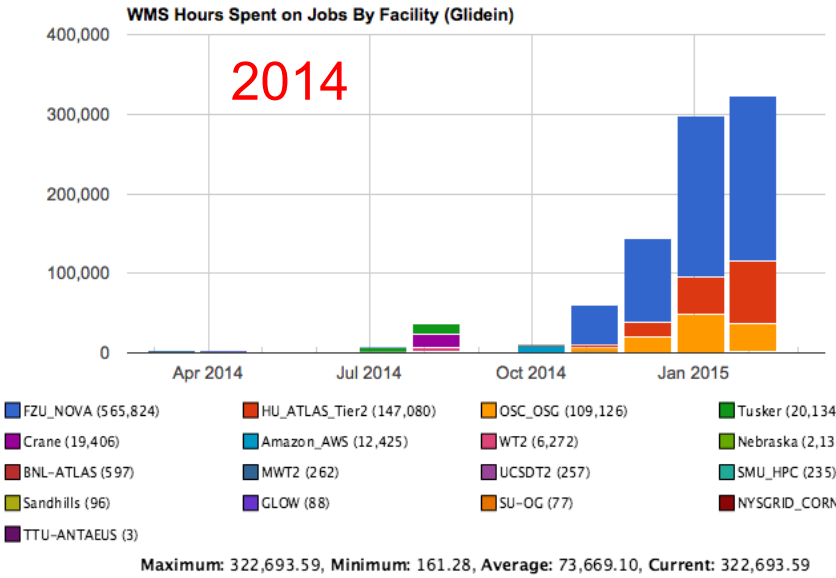

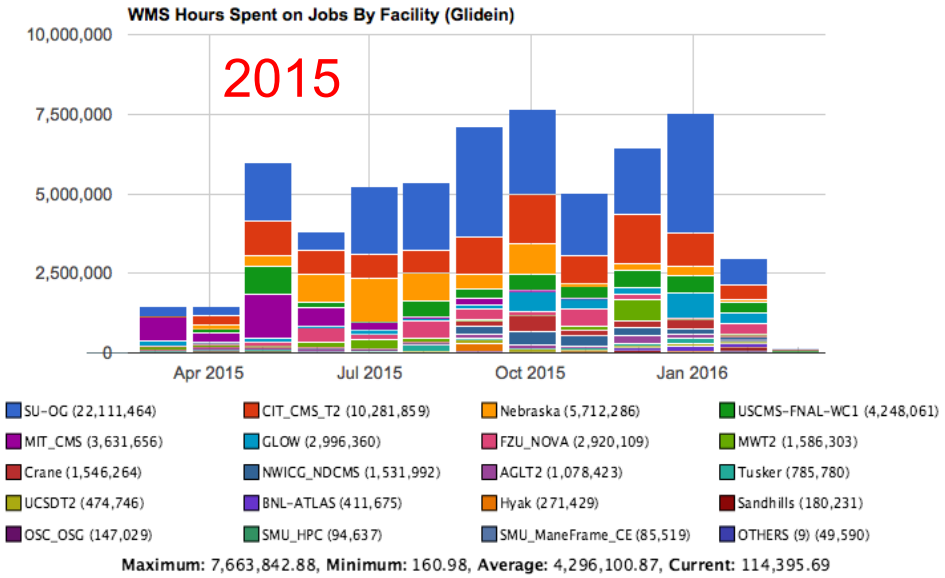

Nuclear Physics beyond the LHC

Open Science Grid

UCSD

Theory meets experiment on the OSG …

… for 50 Million hours in 2015

Steffen Bass (Theory)

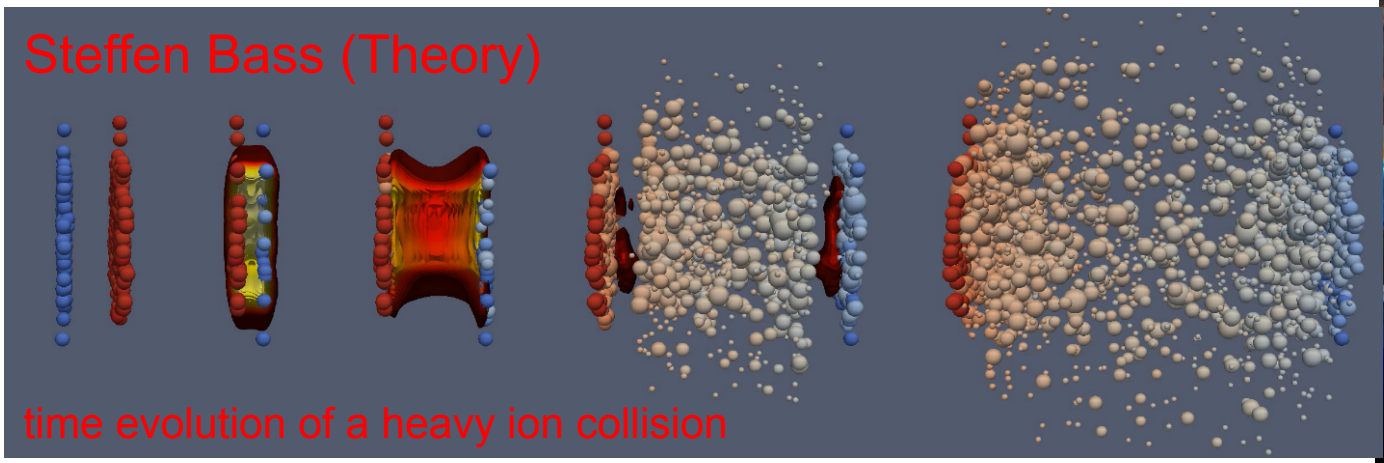

time evolution of a heavy ion collision

LIGO recording a chirp from a long time ago in a galaxy far far away

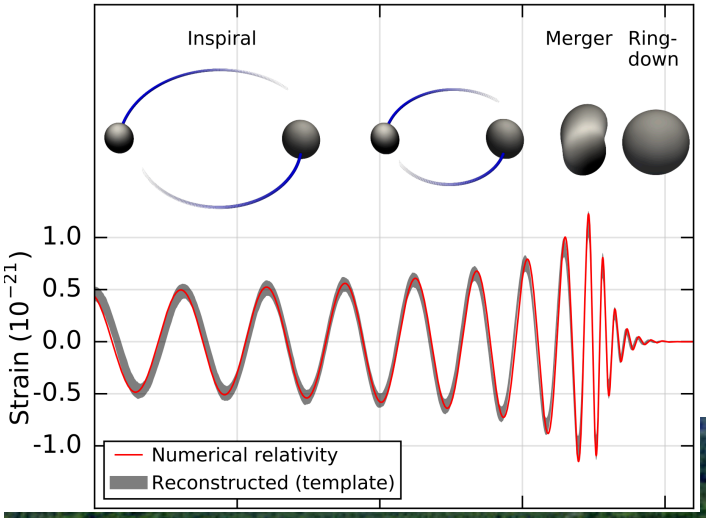

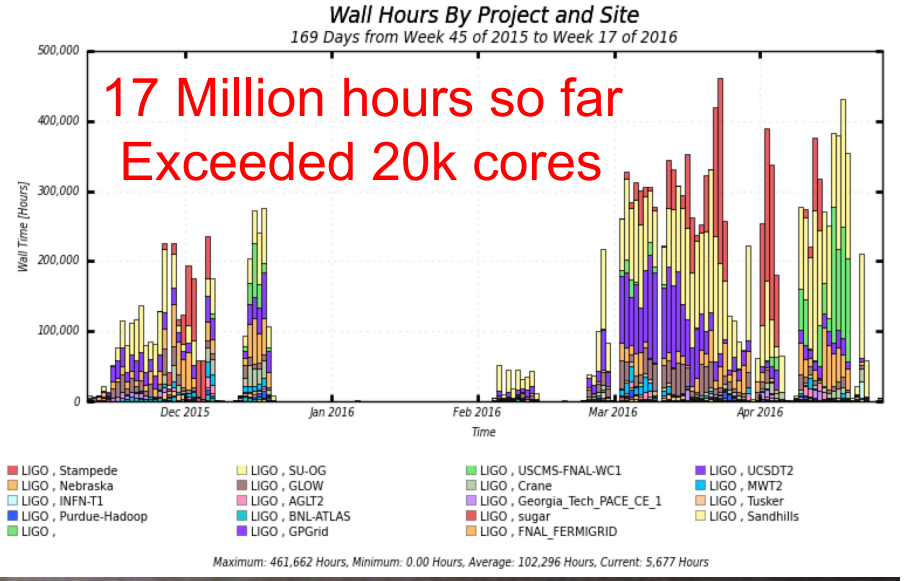

LIGO runs transparently across OSG, Stampede, Comet, and clusters at collaborating institutions in Europe and US

IceCube at the South Pole

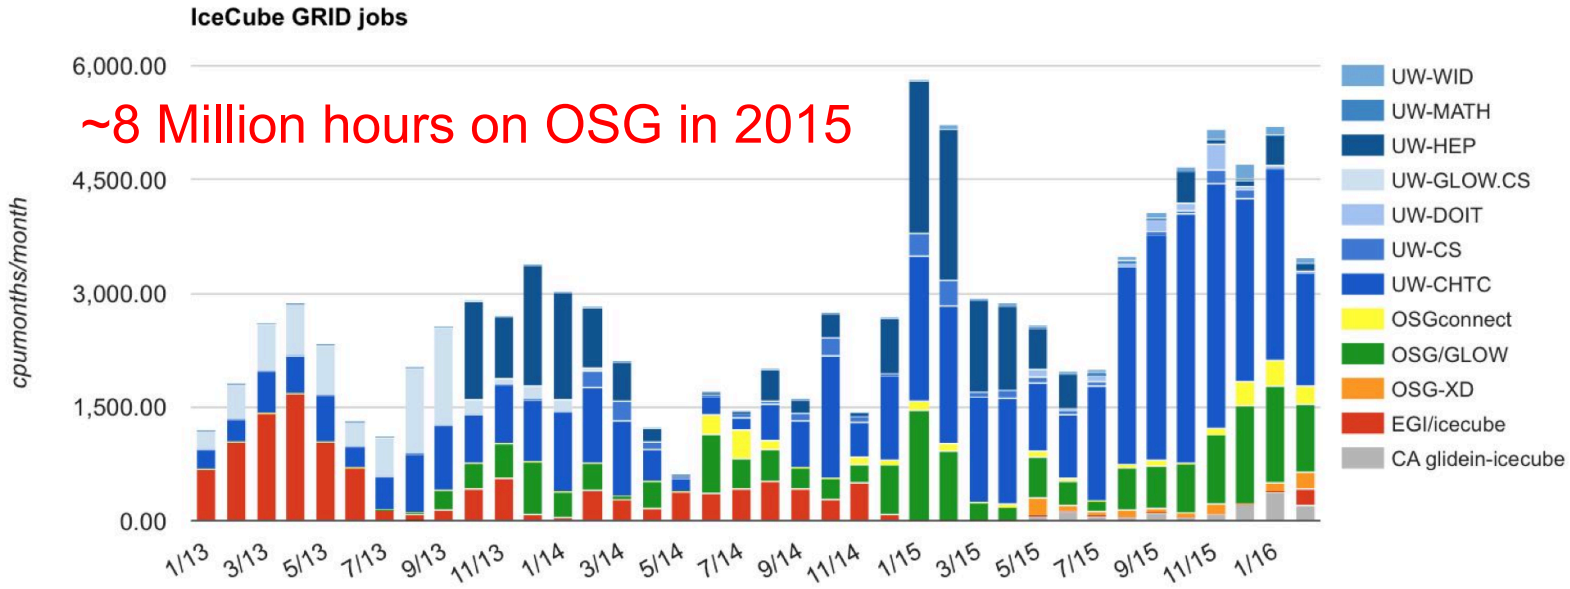

IceCube is pioneering GPU Processing on OSG

The Tusker cluster at UNL is the first campus cluster to share their GPUs on OSG.

Row Labels	Count Jobs	Sum of wall_time	Hours
glow	26774	116169618	32269.33833
hcc	110	98480	27.35555556
osg	2006	18890593	5247.386944
Grand Total	**28890**	**135158691**	37544.08083

← IceCube is pioneering GPU capability on OSG via GLOW VO

~ 40,000 hours of GPU computing in 2015.

GPU use on Comet via OSG started. Other XD resources expected to follow.

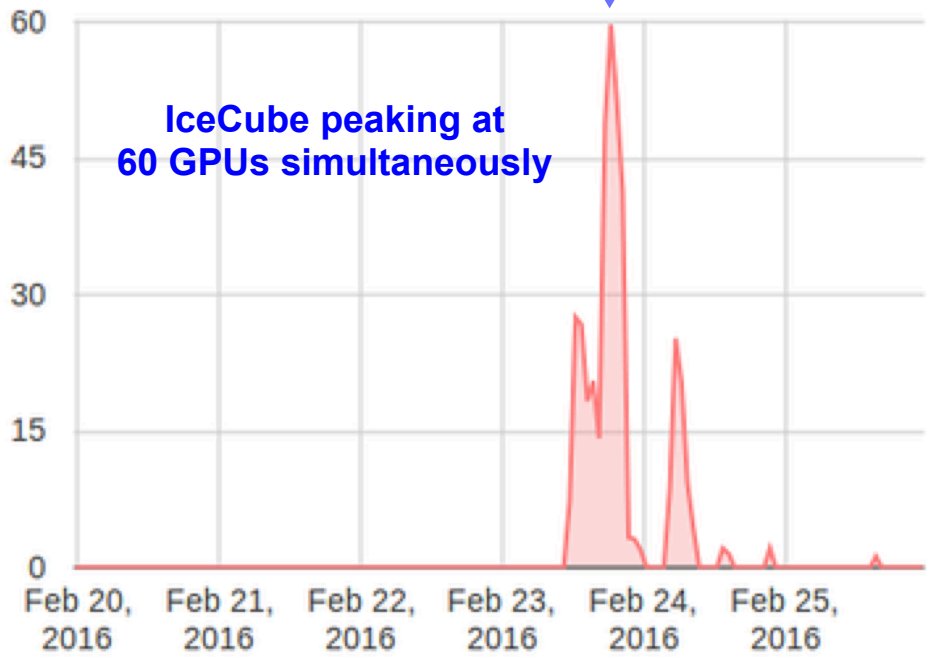

AMS on the Space Station

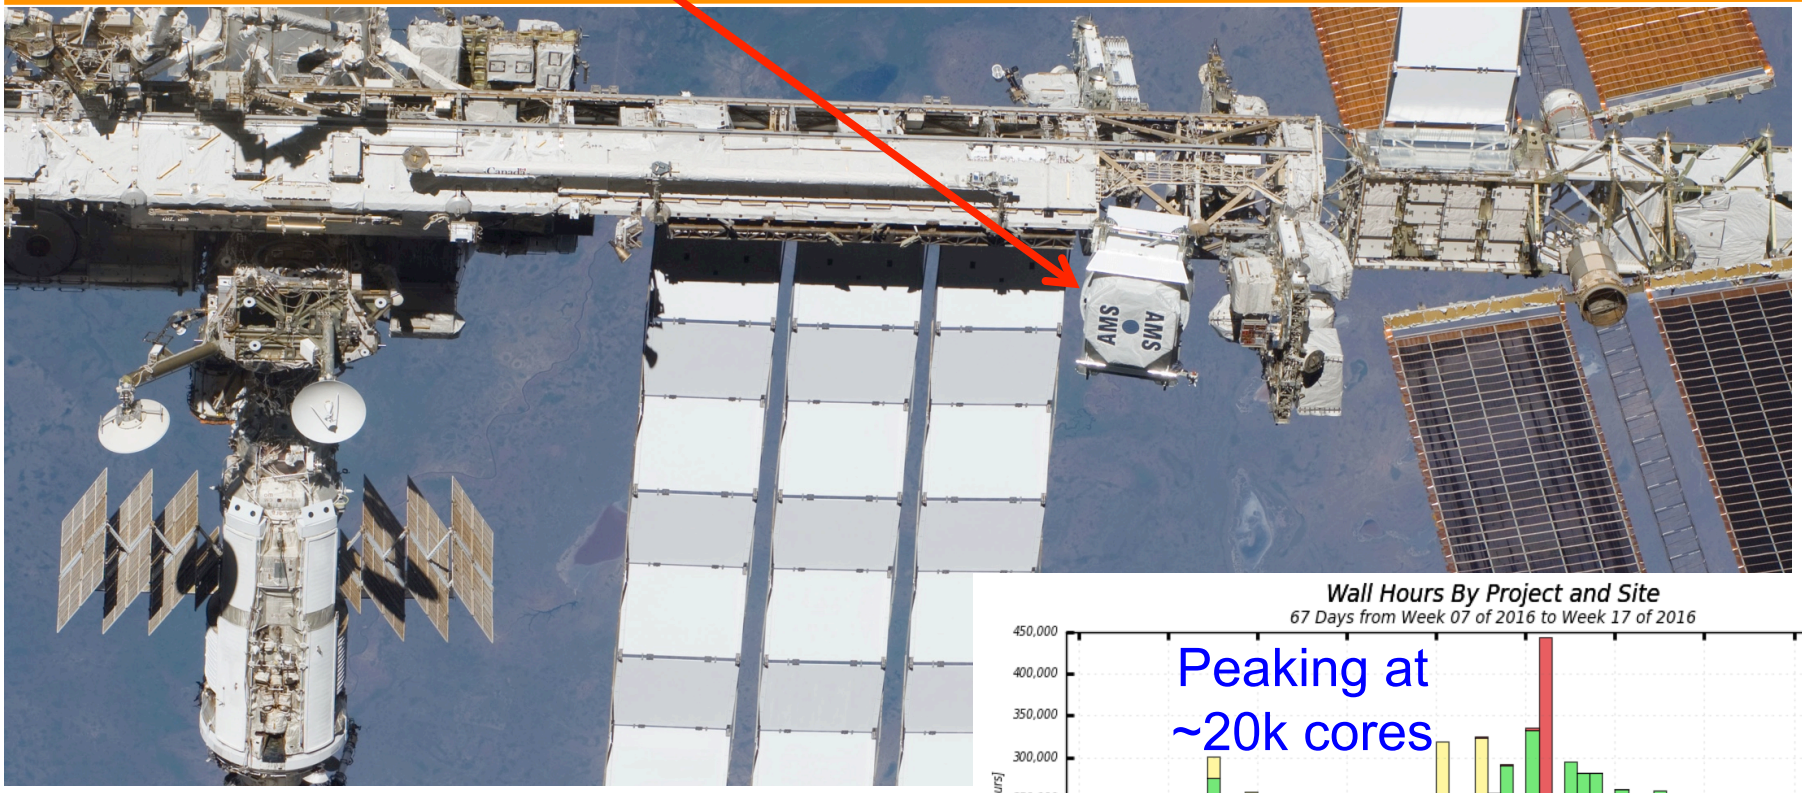

OSG helping MIT to integrate Earth, Atmospheric, and Planetary Science and Laboratory of Nuclear Science computing with OSG and XD resources to produce AMS simulations.

> 10 Million hours so far

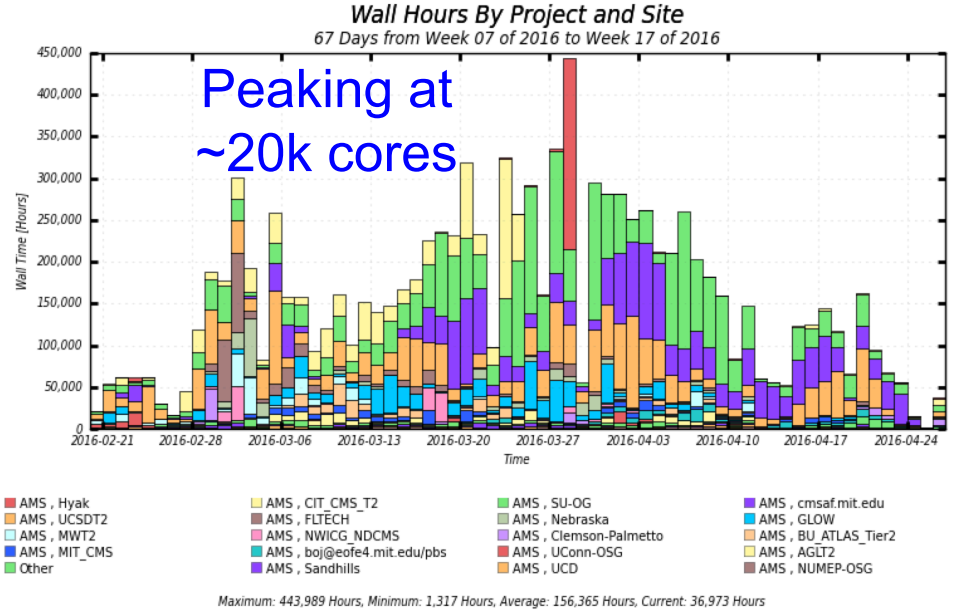

3rd Largest Science Gateway across NSF operates exclusively on OSG

SBGrid uses OSG as "compute engine" for its portals.

3.4 Million hours in 2015

SBGrid CONSORTIUM

Supported Software | Computing Resources | Get Help | About SBGrid | Join us

Home › Portal Applications

SBGrid Portals

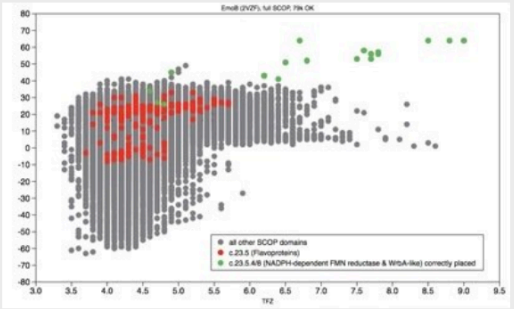

Wide Research Molecular Replacement

Wide Research Molecular Replacement (WSMR) performs a molecular replacement search with the application Phaser against approximately 100,000 SCOP domains. The search results are presented in the form of an LLG/TFZ graph and computations usually converge within 4-14 days. Before you submit your computations you will need to create an SBGrid portal account (request an account here) and upload your mtz file. Protein sequence files are not required. For a full description of this method please see our publication Stokes-Rees I and Sliz P, PNAS 2010. 107(50):21476-81.

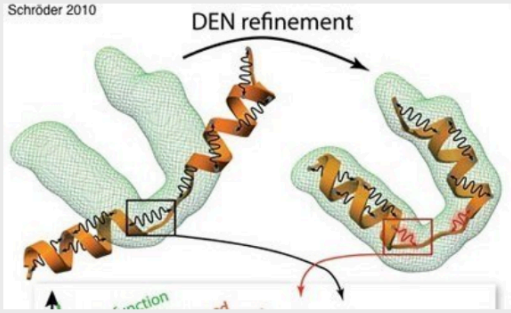

Deformable Elastic Network

Deformable Elastic Network (DEN) restraints are a powerful tool for refining structures from low-resolution X-ray crystallographic data sets. Our DEN web service provides access to resources for running computationally intensive DEN refinements in parallel on the Open Science Grid, a US cyberinfrastructure. Refinements combined with full parameter optimization that would take many thousands of hours on standard computational resources, can be completed in several hours using the DEN portal. For a full description of this service please see our publication O'Donovan D et. al, Acta Cryst. D. 2012. 68:261-7.

OSG Activities Targeted at Individuals

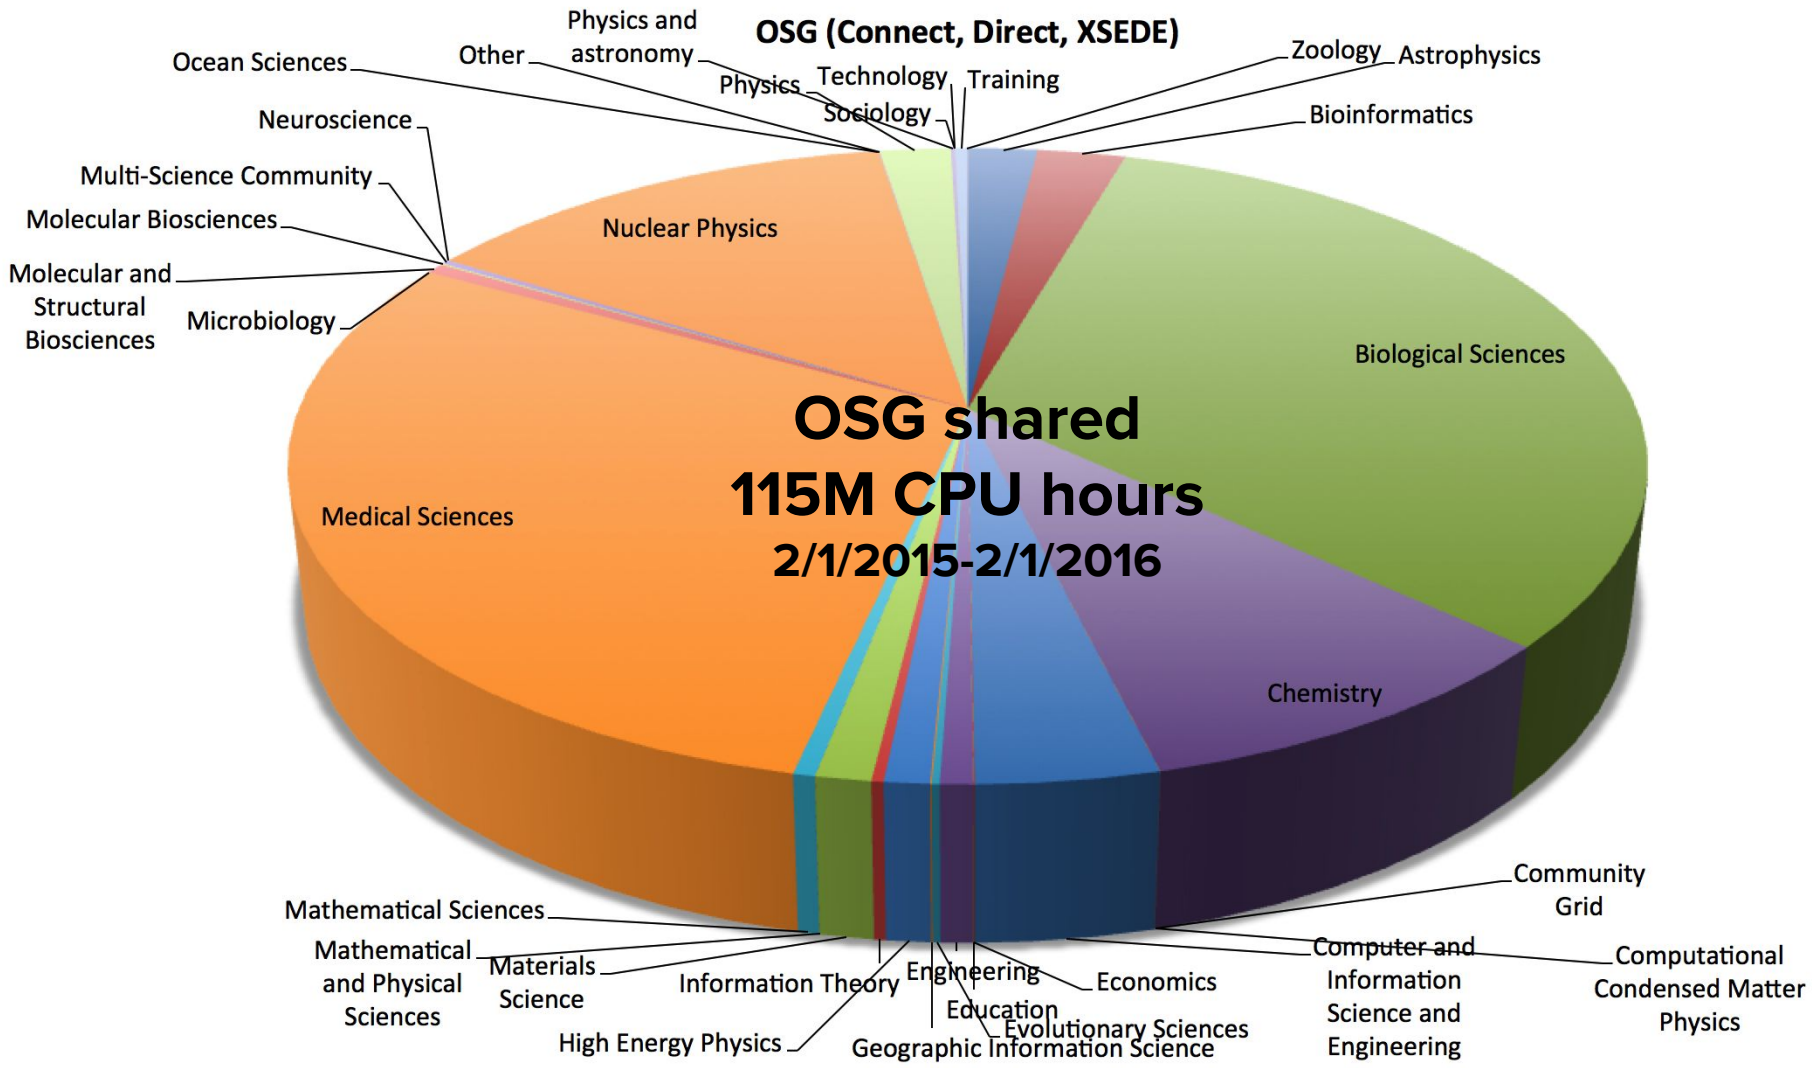

Diversity in many Dimensions

Open Science Grid

UCSD

OSG-Connect Projects 2015/02/01 - 2016/02/01

Project Name	PI	Institution	Feild of Science	Wall Hours
AlGDock	David Minh	Illinois Institute of Technolog	Chemistry	5240846
numfpi	Jerry Tessendorf	Clemson University	Computer and Information Scien	3679098
FFValidate	Vijay Pande	Stanford University	Chemistry	3556135
z2dqmc	Snir Gazit	University of California Berke	Physics	2241066
BioGraph	Alex Feltus	Clemson University	Biological Sciences	1322596
EvolSims	Oana Carja	University of Pennsylvania	Biological Sciences	1286555
CentaurSim	Nathan Kaib	Northwestern University	Astrophysics	1050911
IceCube	Francis Halzen	University of Wisconsin	Astrophysics	901601
PainDrugs	Pei Tang	University of Pittsburgh	Medical Sciences	785932
SourceCoding	Ahmad Golmohammadi	New Mexico State University	Engineering	753884
Errorstudy	Christopher Richards	USDA Agricultural Research	Molecular and Structural Biosci	573448
microphases	Patrick Charbonneau	Duke University	Chemistry	516830
ConnectTrain	Robert William Gardner Jr	University of Chicago	Training	377906
OSG-Staff	Chander Sehgal	Fermilab	Computer and Information Scien	329686
mab	Vivek Farias	Massachusetts Institute of Te	Information Theory	277125
ProtEvol	Premal Shah	University of Pennsylvania	Evolutionary Sciences	163930
uchicago	Robert William Gardner Jr	University of Chicago	Multi-Science Community	156912
RicePhenomics	Harkamal Walia	University of Nebraska Lincc	Biological Sciences	122319
SouthPoleTelescope	John Carlstrom	University of Chicago	Astrophysics	118223
cms-org-nd	Robert William Gardner Jr	University of Notre Dame	High	2
MS-EinDRC	Jacob Pessin	Albert Einstein College of Me	Medi	8
FutureColliders	Sergei Chekanov	Argon National Lab	High	4
SDEalgorithms	Harish S. Bhat	University of California; Merc	Math	0
duke-swcstaff	Robert William Gardner Jr	Duke University	Multi-	90
NSNM	Vadim Apalkov	Georgia State University	Physi	3
EvolvingAI	Jeff Clune	University of Wyoming	Comp	60
Swift	Michael Wilde	University of Chicago	Comp	03
scicomp-analytics	Robert William Gardner Jr	University of Chicago	Multi-	1
DemandSC	Fernando Luco	Texas A&M University; Unive	Econ	72
atlas-org-uchicago	Robert William Gardner Jr	University of Chicago	High	13
KnowledgeLab	James Evans	University of Chicago	Socic	4
Paniceae-trans	Jacob Washburn	University of Missouri	Evolu	0
SNOplus	Joshua R Klein	University of Pennsylvania	High	6
ERVmodels	Fabricia Nascimento	University of Oxford	Zoology	13151
z2qmc	Snir Gazit	University of California Berke	Physics	9598
ContinuousIntegration	Robert William Gardner Jr	University of Chicago	Technology	8537
PathSpaceHMC	Frank Pinski	University of Cincinnati	Computational Condensed Matt	8434
UserSchool2015	Robert William Gardner Jr	University of Wisconsin - Ma	Education	8163
ExhaustiveSearch	Sam Volchenboum	University of Chicago	Bioinformatics	4889
PTMC	Derek Dolney	University of Pennsylvania	Molecular and Structural Biosci	3263
Phylo	Siavash Mirarab	UC San Diego	Bioinformatics	2315
hABCNWHI	Yvonne Chan	Iolani School	Biological Sciences	1822
NSLS2ID	Dean Andrew Hidas	Brookhaven National Labora	High Energy Physics	1486
MiniWorkshopUC15	Robert William Gardner Jr	University of Chicago	Computer and Information Scien	643
freesurfer				234514
SWC-OSG				73112
ProbTracx				8163
bipediffusi				643
DelhiWork				286
EHEC				244
ASPU				9
RADICAL	Shantenu Jha	Rutgers University	Computer and Information Scien	44
HTCC	Rob Quick	Indiana University	Community Grid	15
UserSchool2014	Tim Cartwright	OSG	Multi-Science Community	9
NeoflAnnot	Petra Lenz	University of Hawaii at Mano	Biological Sciences	5
atlas-wg-Exotics	Robert William Gardner Jr	ATLAS	High Energy Physics	4
OSGOpsTrain	Rob Quick	Open Science Grid	Community Grid	3
cms-org-fnal	Lothar Bauerdick	Fermi National Accelerator L	High Energy Physics	3
atlas-org-illinois	Mark Neubauer	University of Illinois	High Energy Physics	1
gem5	Dean Tullsen	University of California San I	Multi-Science Community	1
Total				**24253300**

OSG-Direct Projects 2015/02/01 - 2016/02/01

Project Name	PI	Institution	Feild of Science	Wall Hours
SPLINTER	Rob Quick	Indiana University	Medical Sciences	32746961
Duke-QGP	Steffen A. Bass	Duke University	Nuclear Physics	9591924
sPHENIX	Martin Purschke	Brookhaven Nation	Nuclear Physics	5589611
UPRRP-MR	Steven Massey	Universidad de Pue	Bioinformatics	1593405
IU-GALAXY	Rob Quick	Indiana University	Bioinformatics	1266000
BNL-PHENIX	Matthew Snowball	Brookhaven Nation	Nuclear Physics	519229
DetectorDesign	John Strologas	University of New M	Medical Sciences	376994
Pheno	Stefan Hoeche	SLAC	High Energy Physic	260755
DeerDisease	Lene Jung Kjaer	Southern Illinois Ur	Biological Sciences	157624
OSG-Staff	Chander Sehgal	Fermilab	Computer and Infor	138872
UNC-RESOLVE-pr	David Stark	UNC Chapel Hill	Physics and astron	34840
IBN130001-Plus	Donald Krieger	University of Pittsbi	Neuroscience	2002
Total				**52278216**

OSG-XD Projects 2015/02/01 - 2016/02/01

Project Name	PI	Institution	Feild of Science	Wall Hours
TG-IBN130001	Donald Krieger	University of Pittsbi	Biological Sciences	34454812
TG-DMR130036	Emanuel Gull	University of Michig	Materials Science	1269854
TG-CHE140110	John Stubbs	University of New E	Chemistry	1048149
TG-AST140088	Francis Halzen	University of Wisco	High Energy Physic	552688
	Gregory Snyder	Space Telescope S	Mathematical Scier	420545
	Paul Siders	University of Minne	Chemistry	318432
	Jennifer Lotz			128107
	Francis Halze			83901
	Emiliano Brin			77481
	Juliette Becke			68490
	David Rhee			39368
	Jon Pelletier			36641
	John Chrispe			21074
	Adrian Del Ma			18098
	Yvonne Chan			17884
	Qaisar Shafi	University of Delaw	Physics and astron	17594
	Neranjan Edirisingl	Georgia State Univ	Mathematical Scier	3933
	Suzanne Hawley	University of Washi	Mathematical and F	3020
TG-MCB140232	Alan Chen	SUNY at Albany	Molecular and Stru	598
TG-CCR140028	Shantenu Jha	Rutgers; the State	Computer and Infor	455
TG-TRA100004	Andrew Ruether	Swarthmore Colleg	Training	204
TG-MCB060061N	Jeffry D. Madura	Duquesne Universi	Molecular and Stru	49
TG-MCB140268	Graziano Vernizzi	Siena College	Molecular and Stru	19
TG-STA110011S	Stephen McNally	University of Tenne	Other	1
Total				**38581397**

A few projects selected from OSG Connect Training and Educational Projects

Project Name	PI	Institution	Feild of Science	Wall Hours
ConnectTrain	Robert William Gar	University of Chica	Training	377906
				694877

Campus Projects

Project Name	PI	Institution	Feild of Science	Wall Hours
Duke-SWC-OSG1	Mark DeLong	Duke University	Multi-Science Com	234514
uchicago	Robert William Gar	University of Chica	Multi-Science Com	156912
cms-org-nd	Kevin Lannon	University of Notre	High Energy Physic	116187
duke-swcstaff	Mark DeLong	Duke University	Multi-Science Com	73112
Total				**580725**

OSG shared
115M CPU hours
2/1/2015-2/1/2016

40M hours as Service Provider to the XD program

91 projects across 81 institutions in 2015

Openness, flexibility, and integration continue to be key to the success of OSG