Open Science Grid

Open Science Grid All-Hands Meeting

Introduction

April 9, 2014

Lothar Bauerdick
OSG Executive Director

OSG in Numbers: Running Well

- OSG Delivers up to 2 Million CPU hours every day
  - ★ almost 700M hours of Distributed High-Throughput Computing per year, of which ~90M were provided as “opportunistic resources”
  - ★ about 60% go to LHC, 20% to other HEP, 20% to many other sciences
- OSG has a footprint on ~120 campuses and labs in the U.S.
- OSG transfers ~1 PetaByte of data every day
- Supports active community of 20+ multi-disciplinary research groups

Prepare for expected growth in 2015

Open Science Grid

LHC Run2

★ LHC plans to ~double CPU resources over coming 2 years

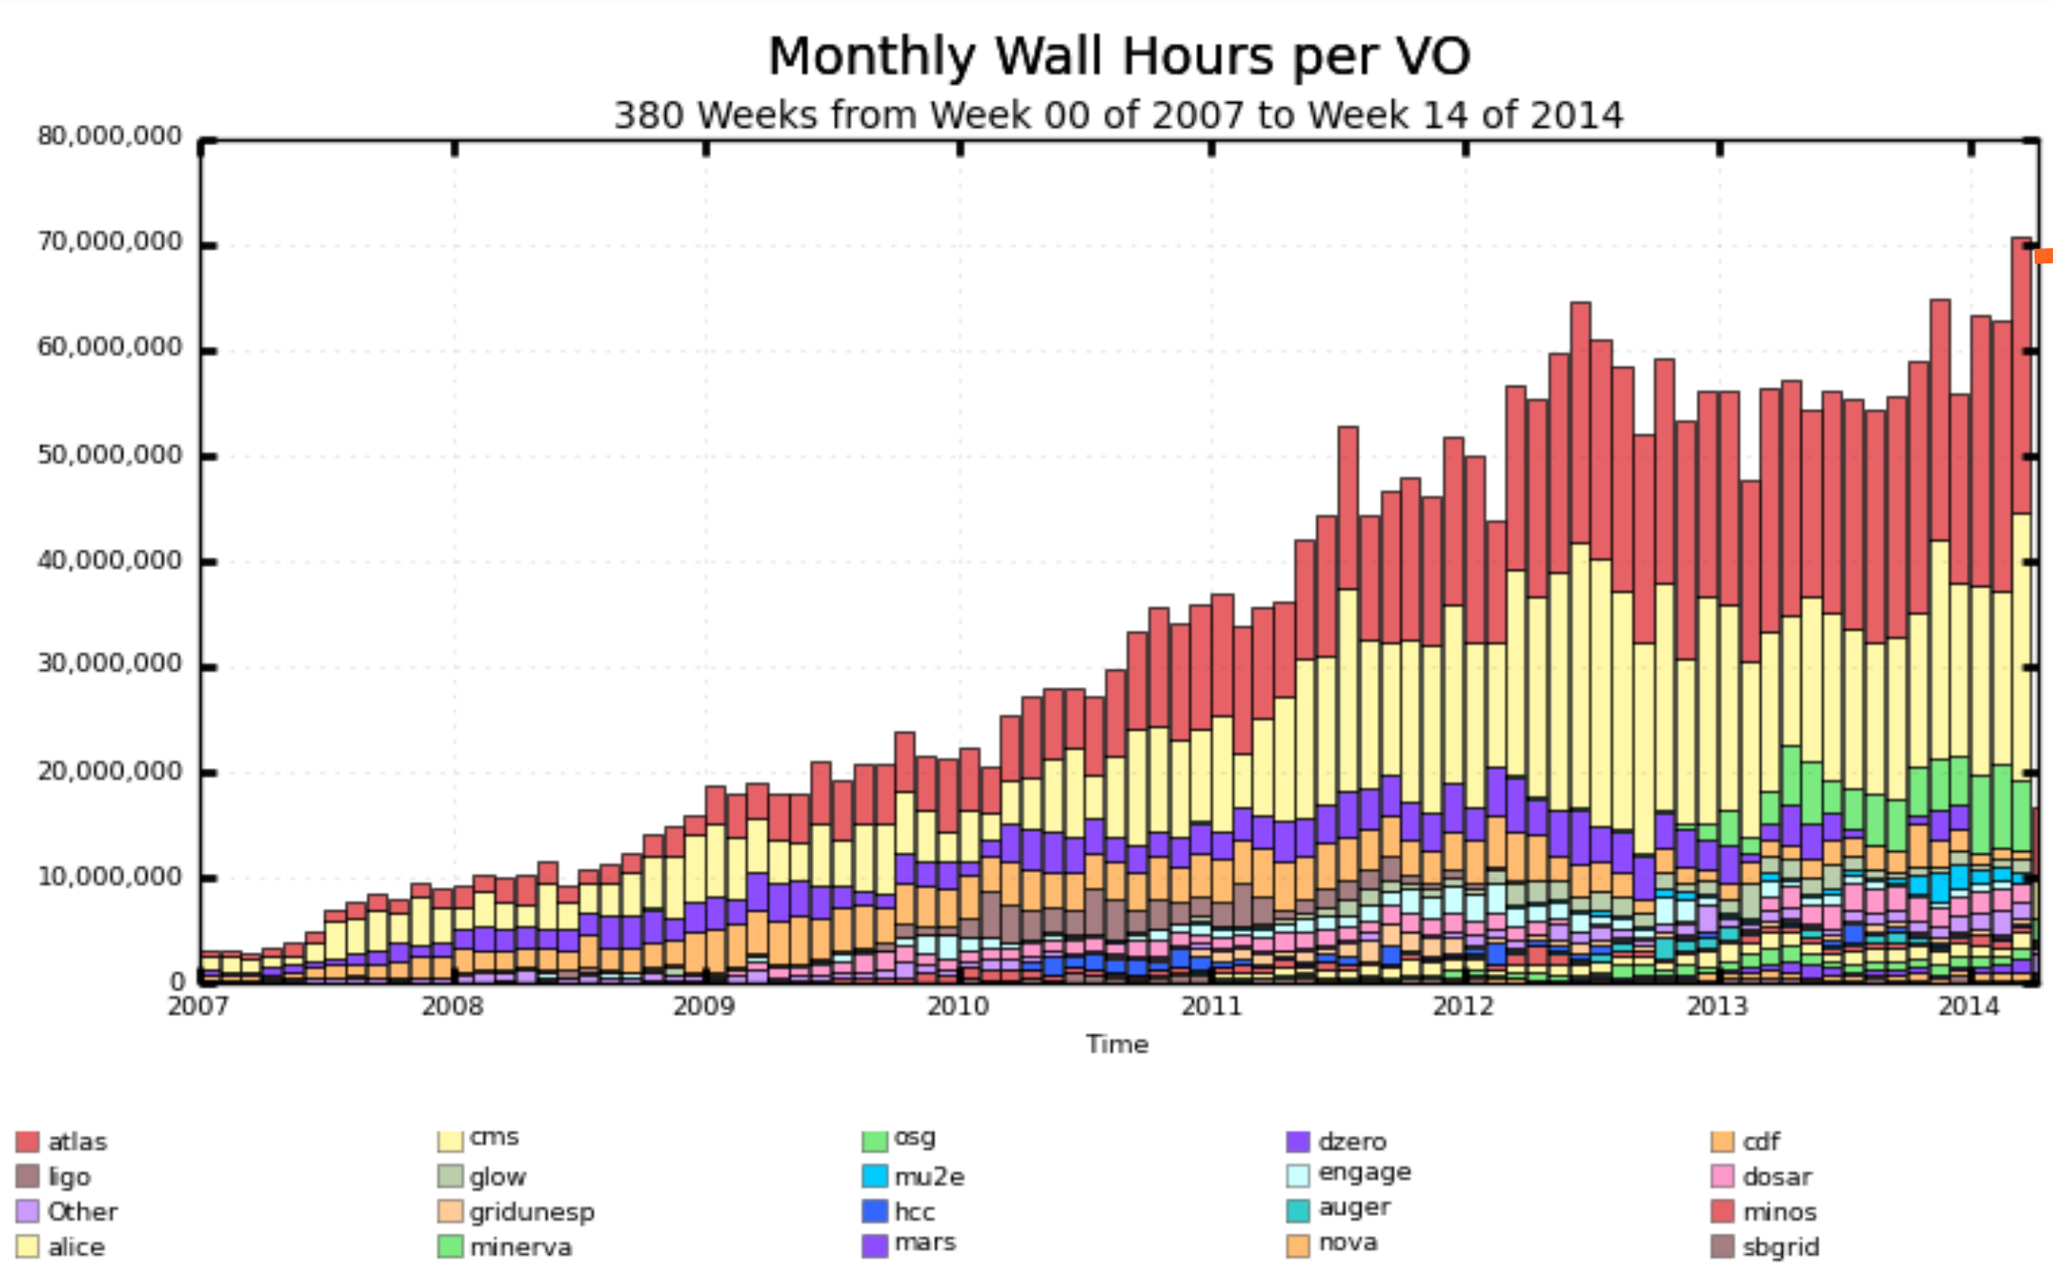

Open Science Grid

~Proportional Increase of Opportunistic Use

- ★ Opportunistic use at ~>25%
- ★ Mostly “Campus users”
- ★ Diverse group of sciences

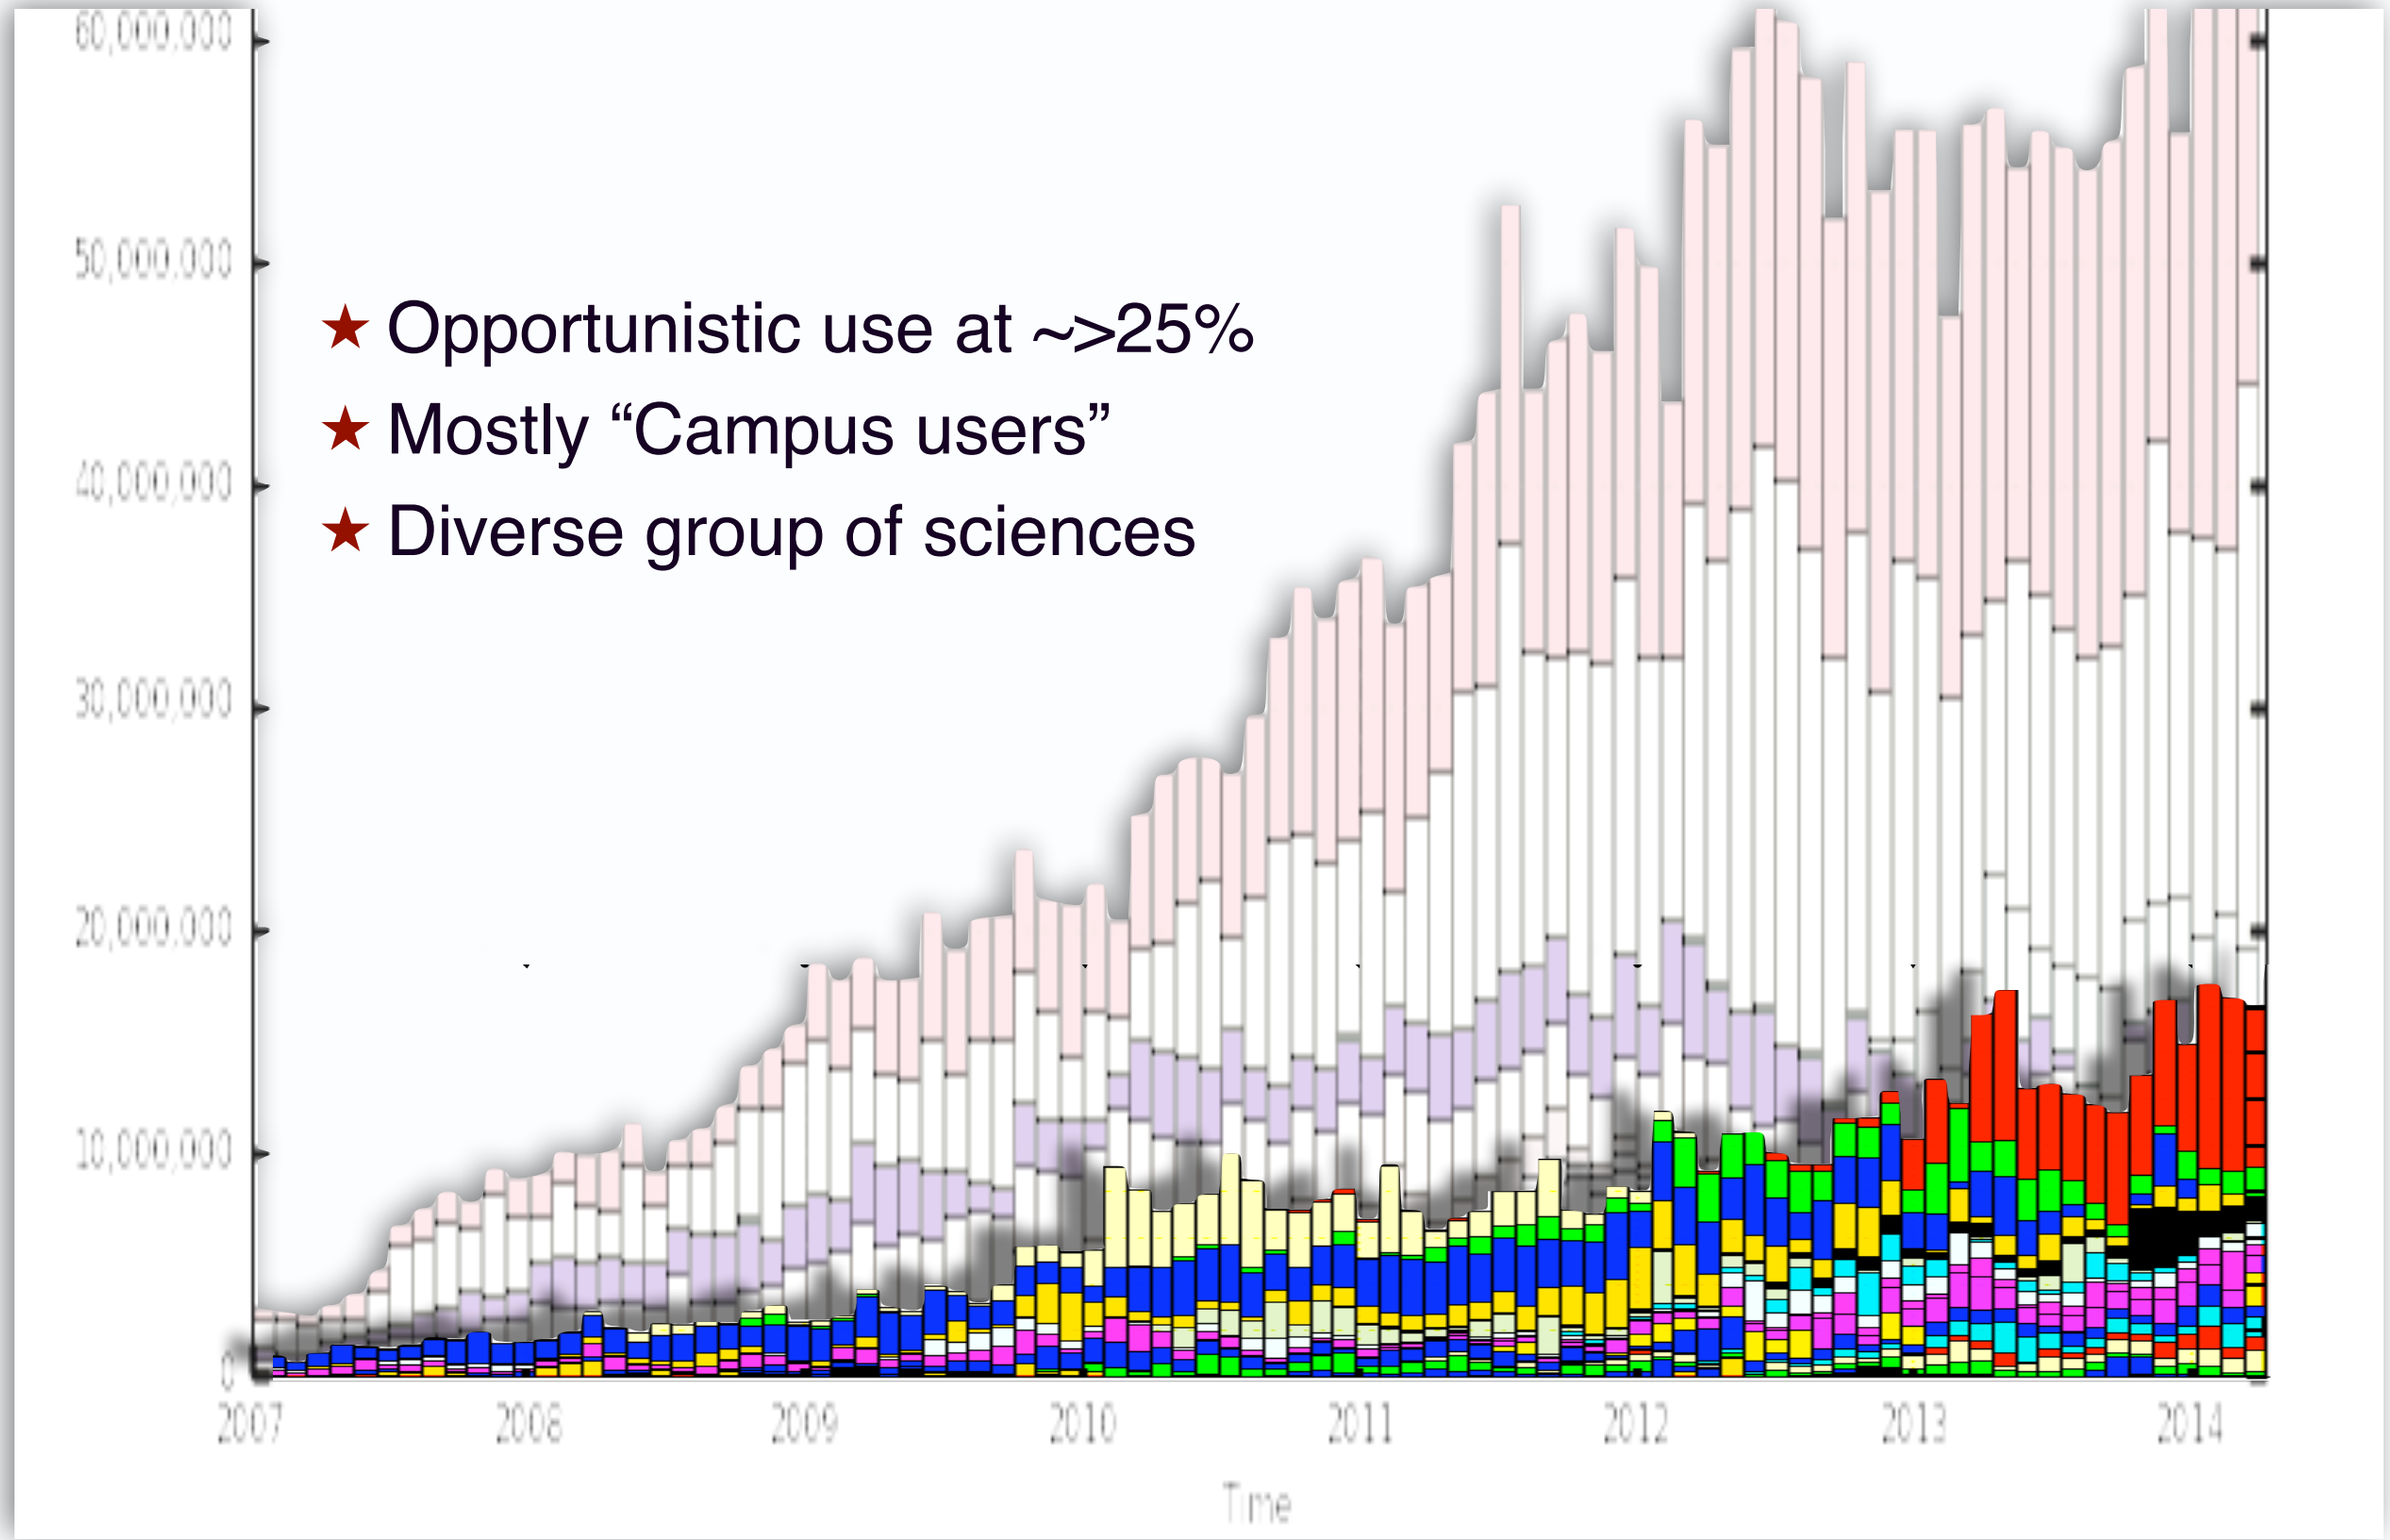

OSG Services and Facility Ecosystem

Open Science Grid

VO-centric Model

Campus User Model

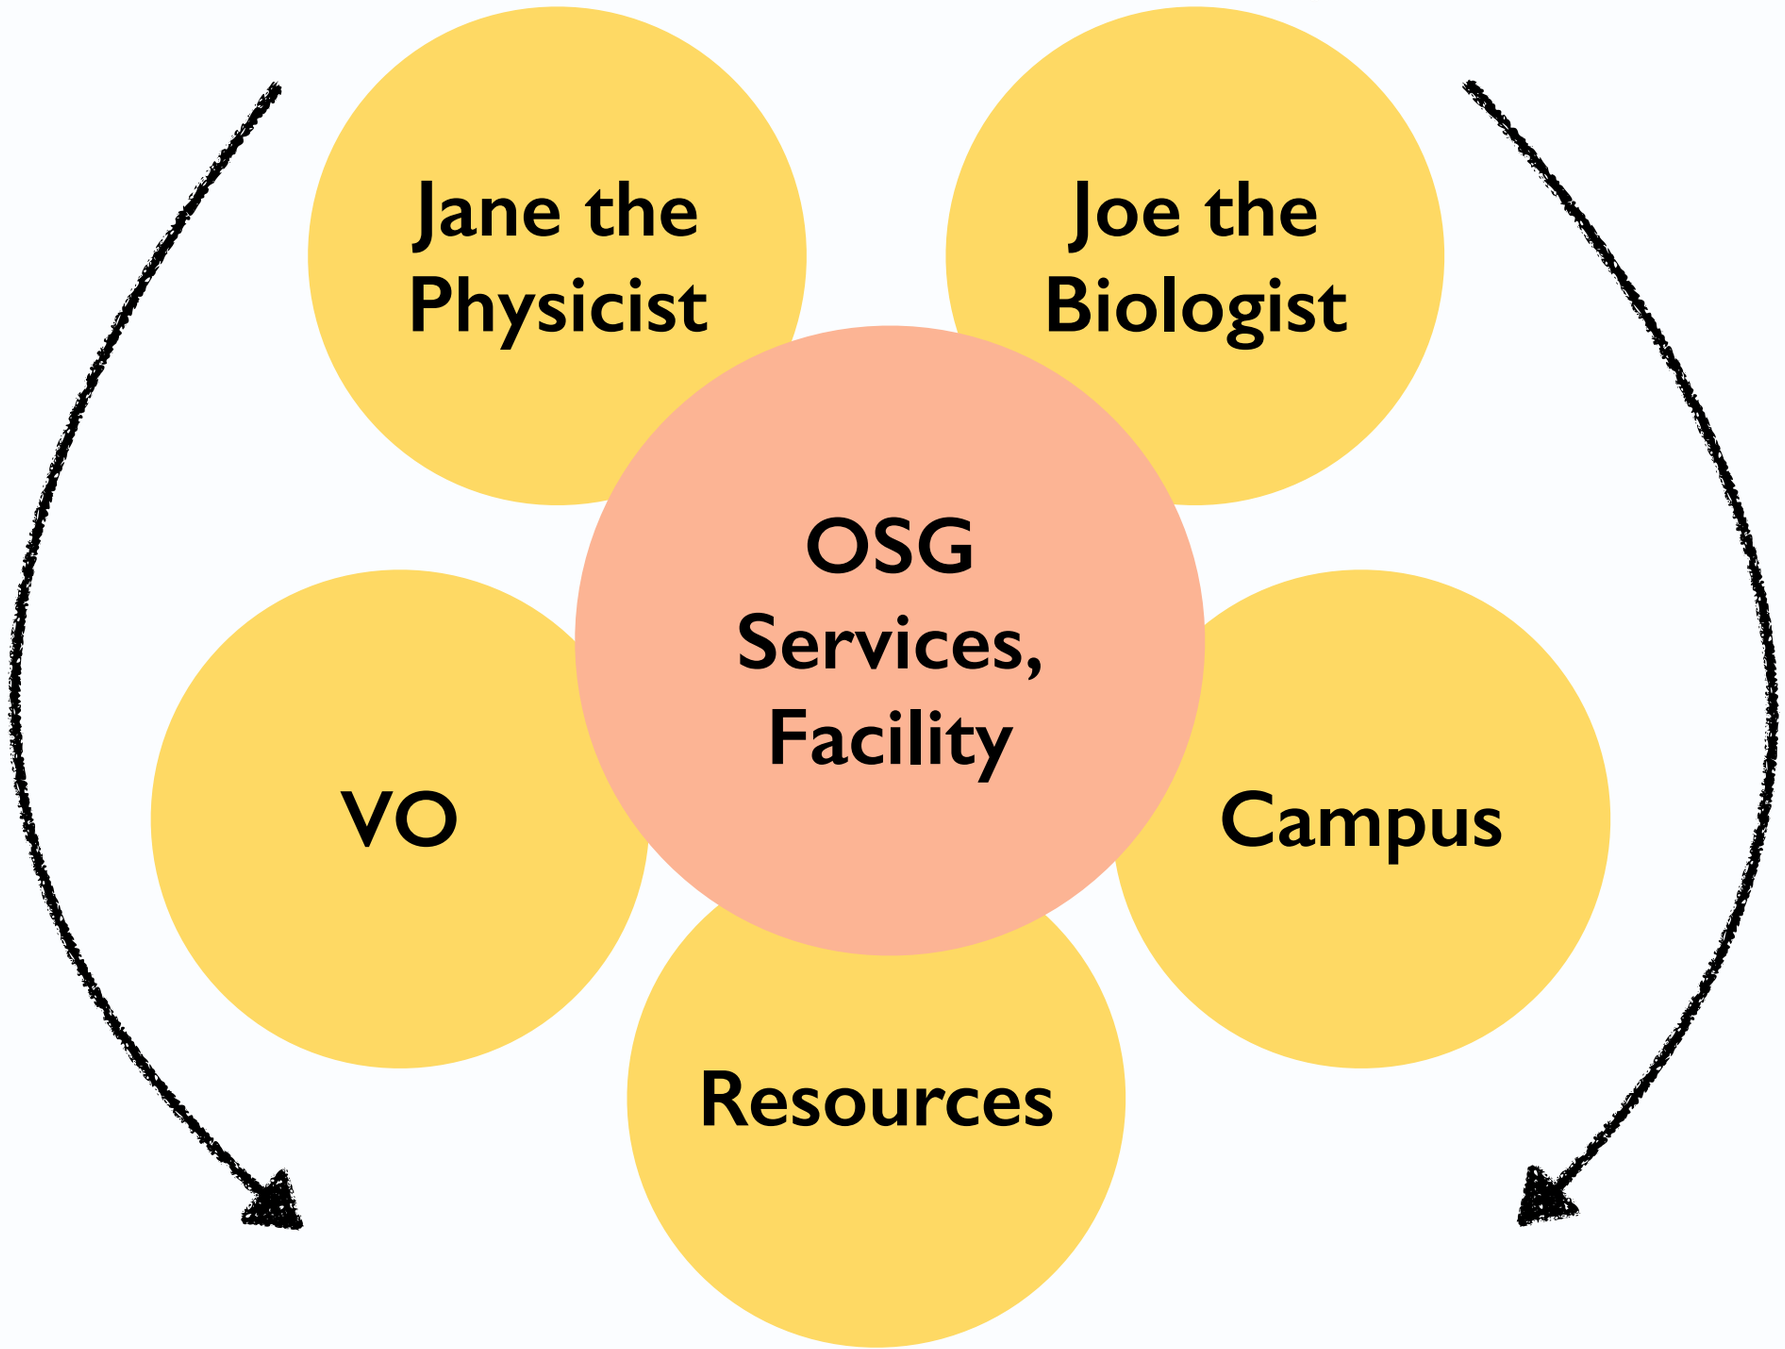

The OSG Open DHTC Facility
based on harvesting otherwise idle resources

- OSG VO:
  - ★ XD Service Provider
    - about 30M hours
  - ★ OSG Connect users
  - ★ a number of science applications
  - ★ individual science groups, who also come in through the engage VO
- Users flowing into OSG from campus grids:
  - ★ glow campus grid at U.Wisconsin
  - ★ sbgrid at Harvard Medical School
  - ★ hcc campus grid at Omaha/Lincoln

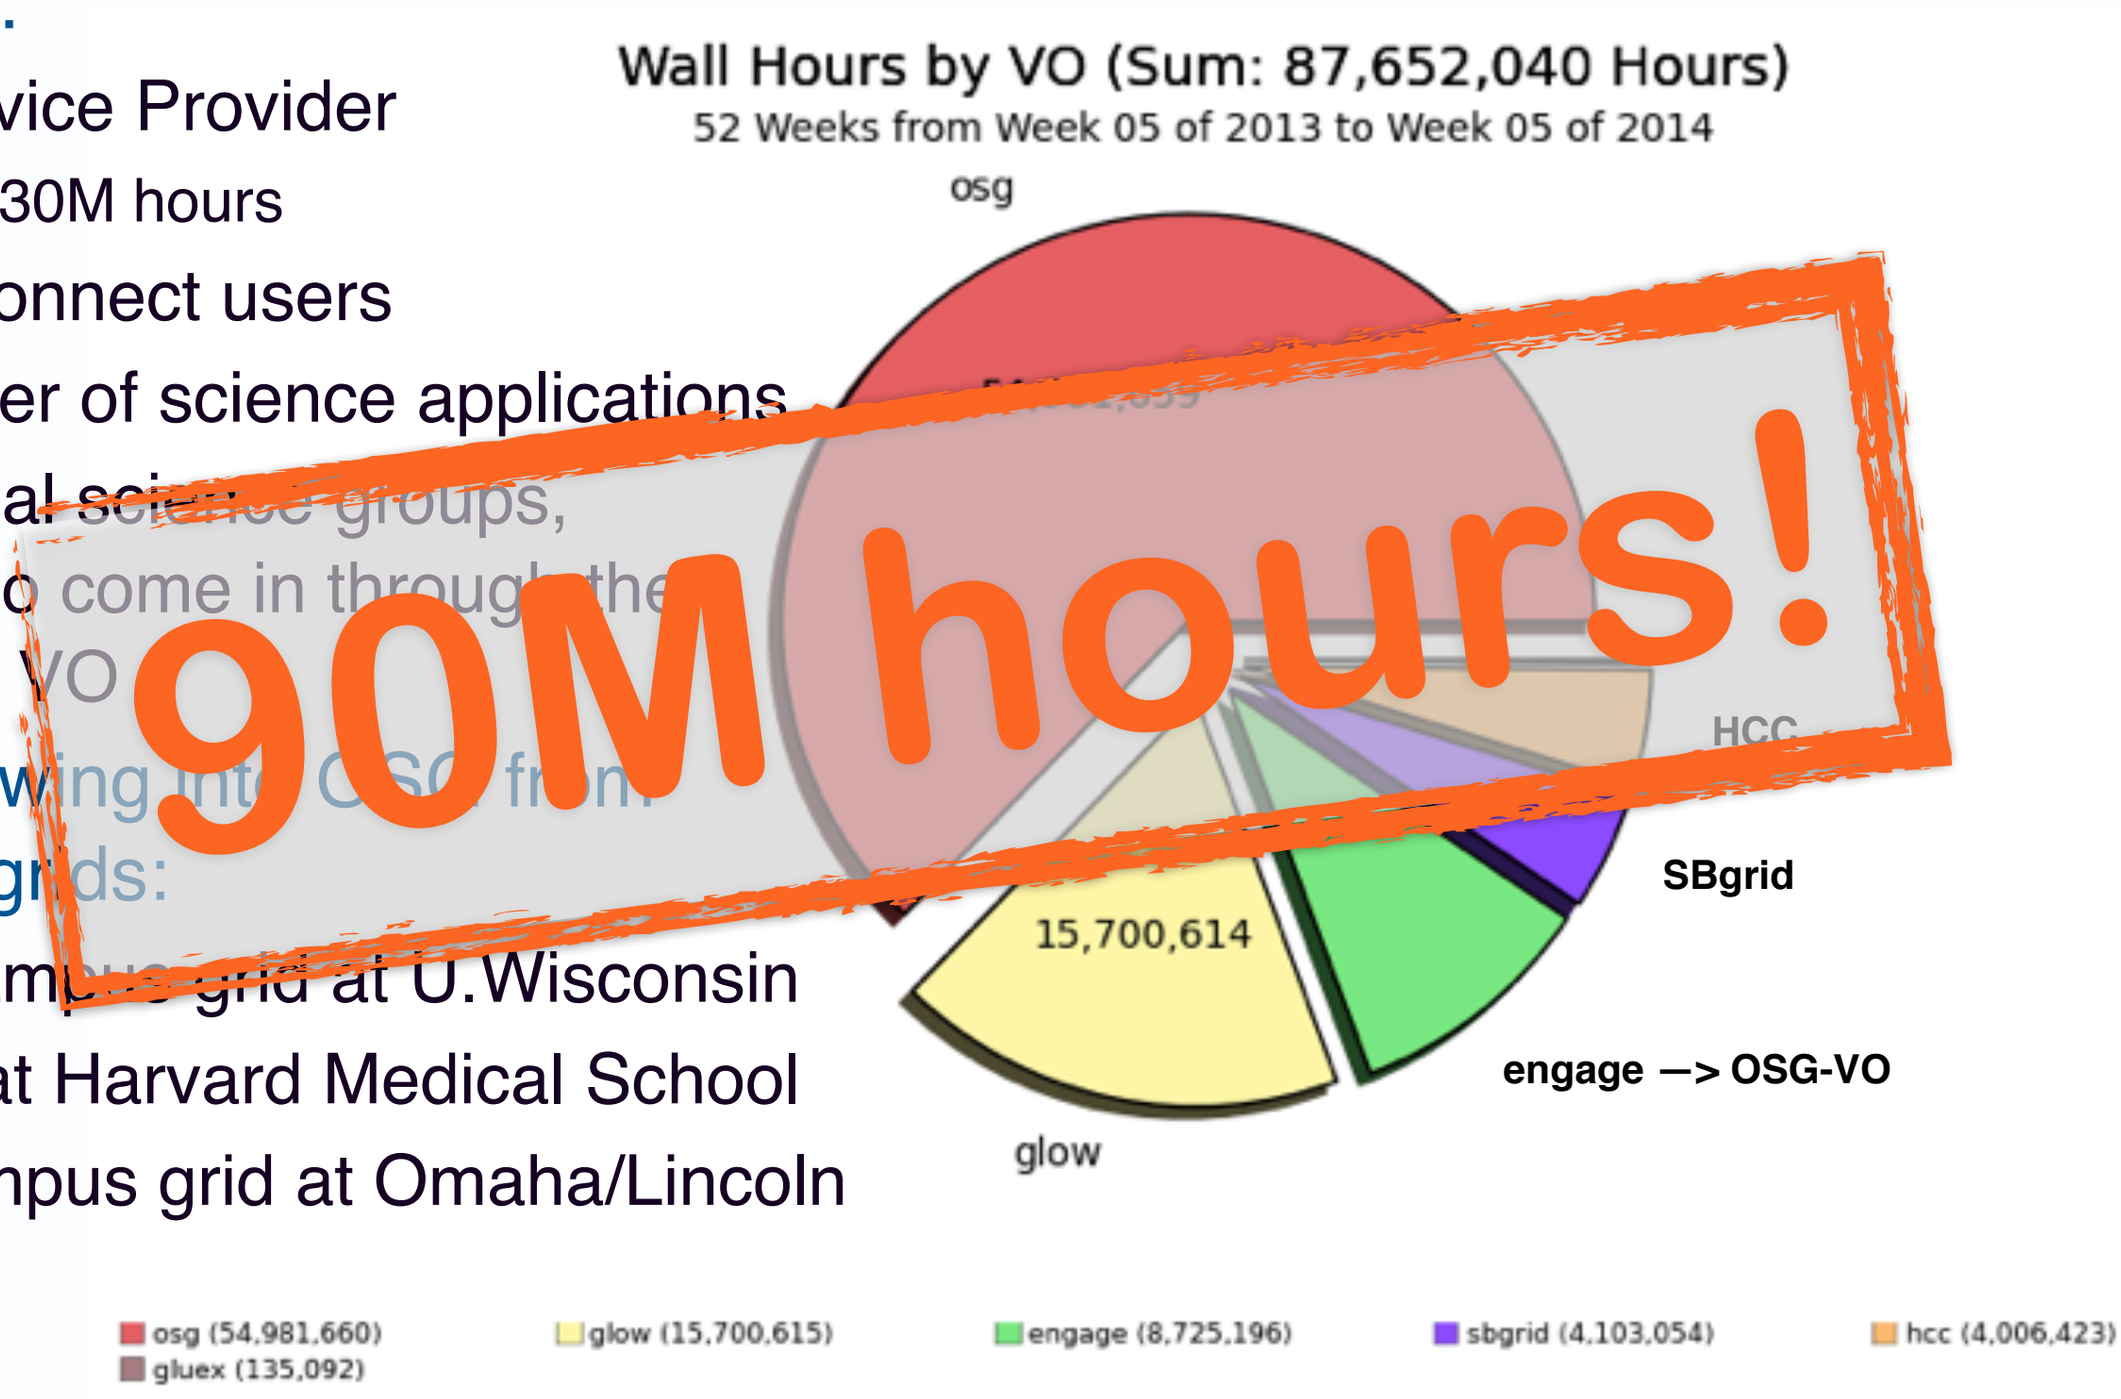

90M hours!

OSG Project Facts

- OSG Project is reaching the half-time point of its current 5 year run
  - ★ will have an agency review of our work plan for the 2nd half of the project

Program	Funds/Year
NSF OCI	$1,000k
NSF MPS	$2,750k
DOE OHEP	$1,600k
DOE NP	$50k
Total	**$5,400k**

Area Manpower	FTE
Technology	8.6
Release Mgmt	1
Operations	9.2
Campus Grids	1.5
Networking	0.4
Security	2.1
User Support	2.45
Project Office & Communications	1.4
Total	26.65

OSG Staff	Responsibility
Lothar Bauerdick (Fermilab)	Executive Director
Miron Livny (Wisconsin)	Technical Director and PI
Chander Sehgal (Fermilab)	Project Manager, User Support Lead
Rob Quick (Indiana)	Operations Lead
Rob Gardner (Chicago)	Campus Grids Lead
Brian Bockelman (Nebraska)	Technologies Lead
Mine Altunay (Fermilab)	Security Officer
Ruth Pordes (Fermilab)	Council Chair
Tim Cartwright (Wisconsin)	Software Lead
Tim Theisen (Wisconsin)	Software Releases
Shawn McKee (Michigan)	Network Area Lead
Michael Ernst (BNL)	ATLAS planning
Frank Würthwein (UCSD)	CMS planning, Resource Manager

Main Areas OSG Delivers On

- Production/Operations
  - provide the OSG platform/eco system of services, sites, software for DHTC
    - infrastructure services, operations support, cyber security and incident response etc to enable VOs to run DHTC workflows and data systems across OSG sites (à la LHC)
  - a production quality HTC facility for a large & diverse community of researchers
    - built on harvesting resources opportunistically from OSG sites
    - delivering amongst others as a XD Service Provider, through XRAC allocations
  - provide other added values:
    - user and host certificates, software distribution services, network monitoring, …
- User Support
  - consulting on technologies, architectures and user support
  - spreading knowledge on HTC as a science problem solver
- Technologies and Software
  - developing concepts and blueprints, deliver an evolving software stack
  - packaging, system testing, patching
- Campus Grids
  - OSG Connect service
  - Campus Infrastructure Community

Preparing to formulate project plan for coming years

- OSG plan for the 2nd half of the project, to be reviewed by agencies
  - develop a convincing work plan, put us on a trajectory for sustaining OSG
- Ran an internal review on March 14, 2014 at Fermilab
  - review committee from the OSG eco system to advise the ET/AC: Jim Bottum, Ewa Deelman, Ian Fisk, Mark Neubauer (thank you!)
  - AC prepared excellent presentations outlining status and plans
    - focus on impact to stakeholders, key metrics, challenges to work on in year3+
    - available at https://indico.fnal.gov/conferenceDisplay.py?confId=8099
    - also attached are summaries of reviewer comments
- Endorsement of the general directions, e.g. OSG connect,
  - input on where to focus technical work, e.g. more dynamic provisioning…
  - guidance: engage stakeholders more, provide technical leadership, run focus meetings and workshops, engage CIO/campuses, improve web, outreach materials, work with NSF to name OSG in solicitations, …
- Next steps is to outline a formal plan in preparation for agency review
  - staff retreat with AC and ET on May 14/15 at Madison

Agenda for the Morning Session

- Technical Director and PI
  - Miron Livny
- Council Chair
  - Ruth Pordes
- OSG Production - The Foundations for 2M+ CPU Hours/Day
  - Rob Quick
- OSG Technologies Update
  - Brian Bockelman

  - break

- dvdt: Accelerating Science
  - Ewa Deelman
- Panel: OSG in 2017
  - Ken Bloom, with panelists