Open Science Grid

OSG Council Mtg, July 15, 2014

LATBauerdick
OSG Executive Director

Open Science Grid

Update on Review

- Review of OSG Year 4-5 workplan by DOE and NSF
  - August 19-20, 2014, in the NSF Offices in Arlington, VA
  - run by Larry Price/DOE and Jim Shank/NSF
  - 1.5 days agenda with 3.5h presentation from OSG project
  - hopefully in attendance of other program managers, in particular on the NSF side
- Received request for detailed information
  - 1. A progress report
  - 2. A detailed plan for the work of OSG during years 4 and 5
  - 3. Demographics and funding tables for years 4 and 5
  - 4. An abridged curriculum vitae (CV) and professional summary (PS) for each “Permanent PhD” or “Computing Professional”
    - lots of details requested
  - Chander has/is sending request for information to institutional PIs to help wit this material, please respond quickly!

Open Science Grid

Year 3-5 Workplan: starting from a strong basis

- In developing the OSG year 3-5 work plan our starting position is a well established and well working project with a number of forward looking initiatives that we already have started, like
  - ★ the opportunistic “Open Facility” enabling OSG to provide value to campus researchers and being part of XD,
  - ★ the OSG Connect approach to connect researchers and campuses,
  - ★ the network monitoring,
  - ★ the work on trust relationships and replacing certificates with traceability,
  - ★ the software factory and release process,
  - ★ and many more.
- In developing the year 3-5 plan we used as guidelines:
  - ★ Continue doing brilliantly where we are doing brilliantly now!
    => Don't loose track of what we do well
  - ★ Expand on opportunities we see emerging from work we already have started, and solidify them towards doing brilliantly.
  - ★ Pick a few areas that we think are completely new to us, and where we can have a large impact with modest effort
- provided 2 docs to DOE, summarizing year 3 plans and year 1-2 achievements

OSG work plan: high-level goals

- In order to develop the concrete road maps/work plans we break down OSG goals and approaches into the following three focus areas, each of which presents a clear value proposition:

  - 1. Provide value to VOs and resource owners by enabling DHTC and resource sharing for their set of applications (the larger experiments and science collaborations)

  - 2. Provide opportunities for the "long tail of science" (individual scientists and groups on the campuses)

  - 3. Promote DHTC for science

- in the May staff retreat meeting, we defined and prioritized goals/ deliverables for these three focus areas, that inform the roadmaps and work plans for the OSG project areas

Production and Operations Related OSG Deliverables

- provide the OSG platform/eco system of DHTC services, sites, software **to enable VOs** to run workflows and data systems **across OSG sites**
  - infrastructure services, operations support, cyber security and incident response etc
  - main customer is the LHC, and other large experiments/VOs

- provide a production quality **HTC facility** built on **harvesting resources opportunistically** from OSG sites, for a large & diverse community of researchers and science platforms
  - delivering amongst others as a XD Service Provider, through XRAC allocations
  - delivering to science gateways e.g. for biology or medical applications etc

- provide **other added value**:
  - user and host certificates (OSG CA, the follow-up of DOEgrids CA)
  - software distribution services (OASIS based on CVMFS)
  - network monitoring and dashboard
  - …

OSG Technology Area Goals

- Focus Areas for **Stakeholder Needs**
  - finish transition to HTCondor-CE
  - integrate BOSCO/HTCondor-G+ssh into production grid
  - work on HTCondor scalability
  - modernize and solidify authz support
  - maintain software components to support VOs and allow sites to provide resources
  - improve software usability
- For the "**Long Tail of Science**"
  - more flexible application software delivery: OASIS
  - common input data project
  - improve support for user software
- Promote **use of DHTC** for science
  - reboot and re-organize Blueprints activity (Miron Livny)
  - start series of "Focus Meetings"
  - make software widely available, OSG user school

OSG Security Area Goals

- Phase out the PKI based Identity Management (IDM).
  - Decrease our dependence on certificates.
    Shrink the number of user and host certs
- Develop traceability and trust relationship models that can eliminate dependency on PKI
- Achieve certificate-free access to data delivery services
- Integrate OSG IDM with Federations, Campuses
- Focus on campus researchers. Enable use of campus identities.
  Focus on identity/security needs of campus researcher
- Make it easy for campus researcher to join OSG, run jobs and move data.

Open Science Grid

User Support Area Goals

- Effective service delivery for OSG and as Service Provider for XD
  - ★ grow number of XRAC requests
- Provide opportunistic access facility (“on-ramp”) for researchers
  - ★ modernize, iRods -> ROC, research/understand dynamics if opportunistic eco system, grow opportunistic pool
- Integrate Intensity Frontier experiments as prime stakeholders
  - ★ help to move applications into eco-system, like OASIS, data delivery etc
- Grow access to OSG DHTC for US researcher
  - ★ tutorials etc, partnership with ACI-REFs to provide support, XSEDE champions
- Provide integration support for new sites
- Provide OSG introduction and guidance to new communities interested in joining OSG

OSG "Size" is Increasing in Terms of CPU Resources

Open Science Grid

- ★ LHC plans increase in CPU resources at US T1 and T2 over coming 2 years

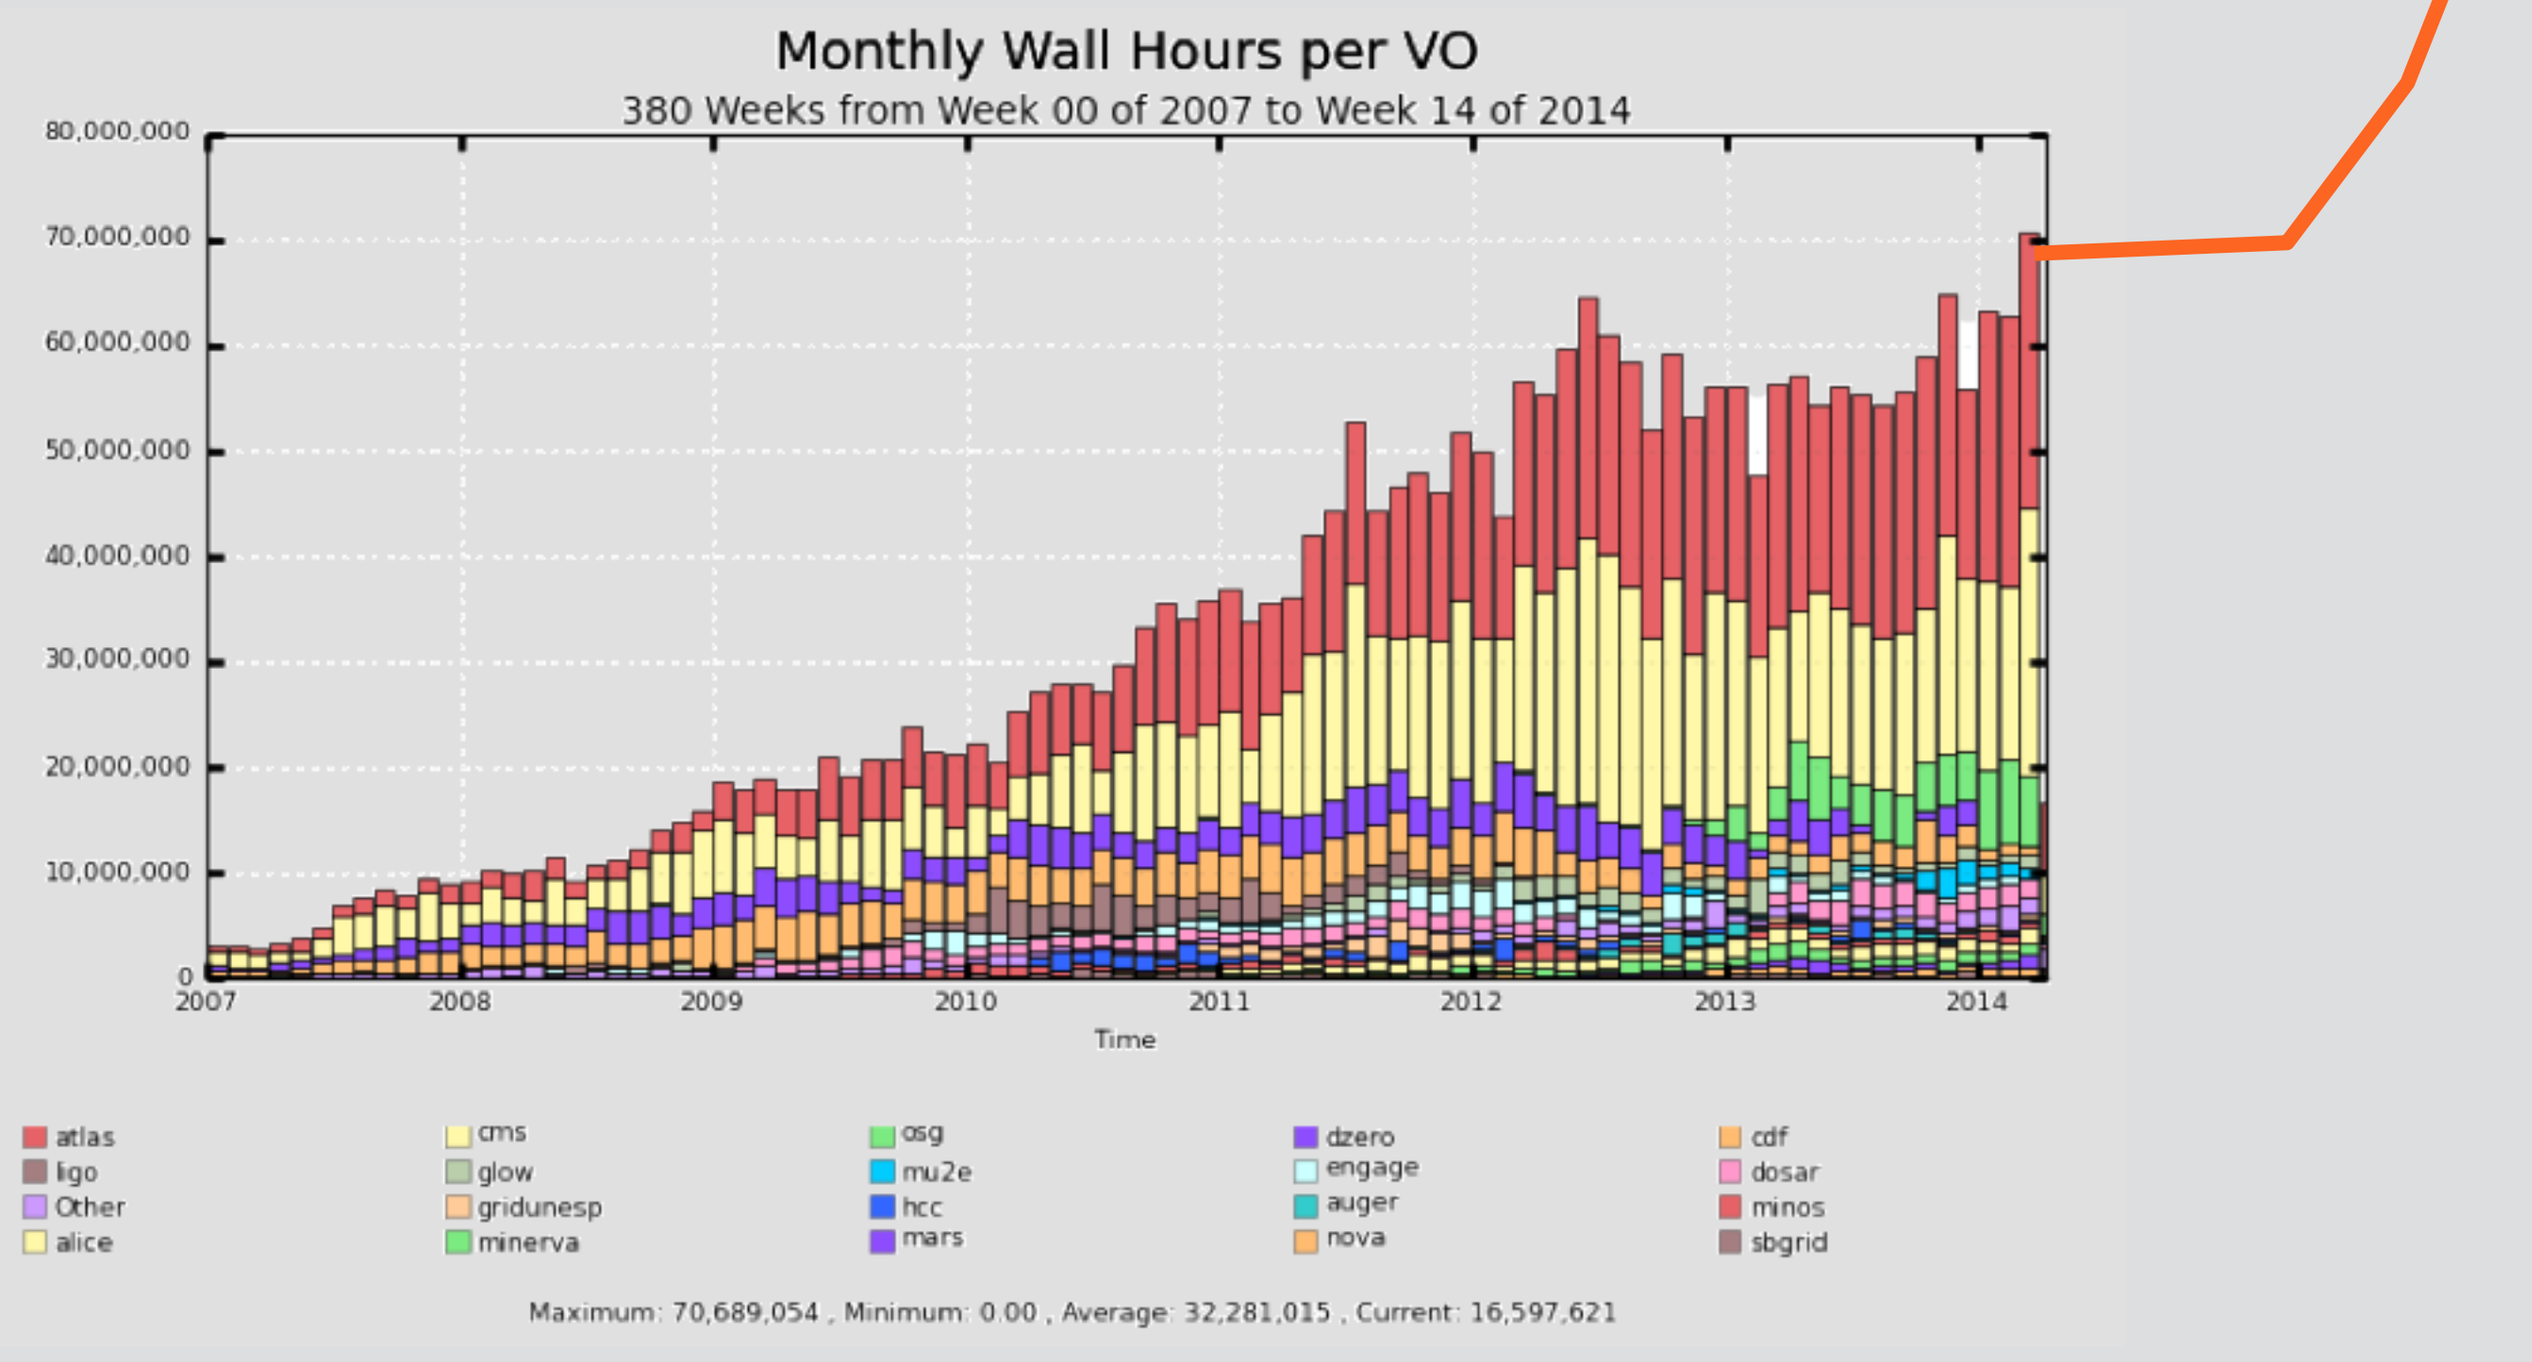

Open Science Grid

~Proportional Increase of Opportunistic Use?

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

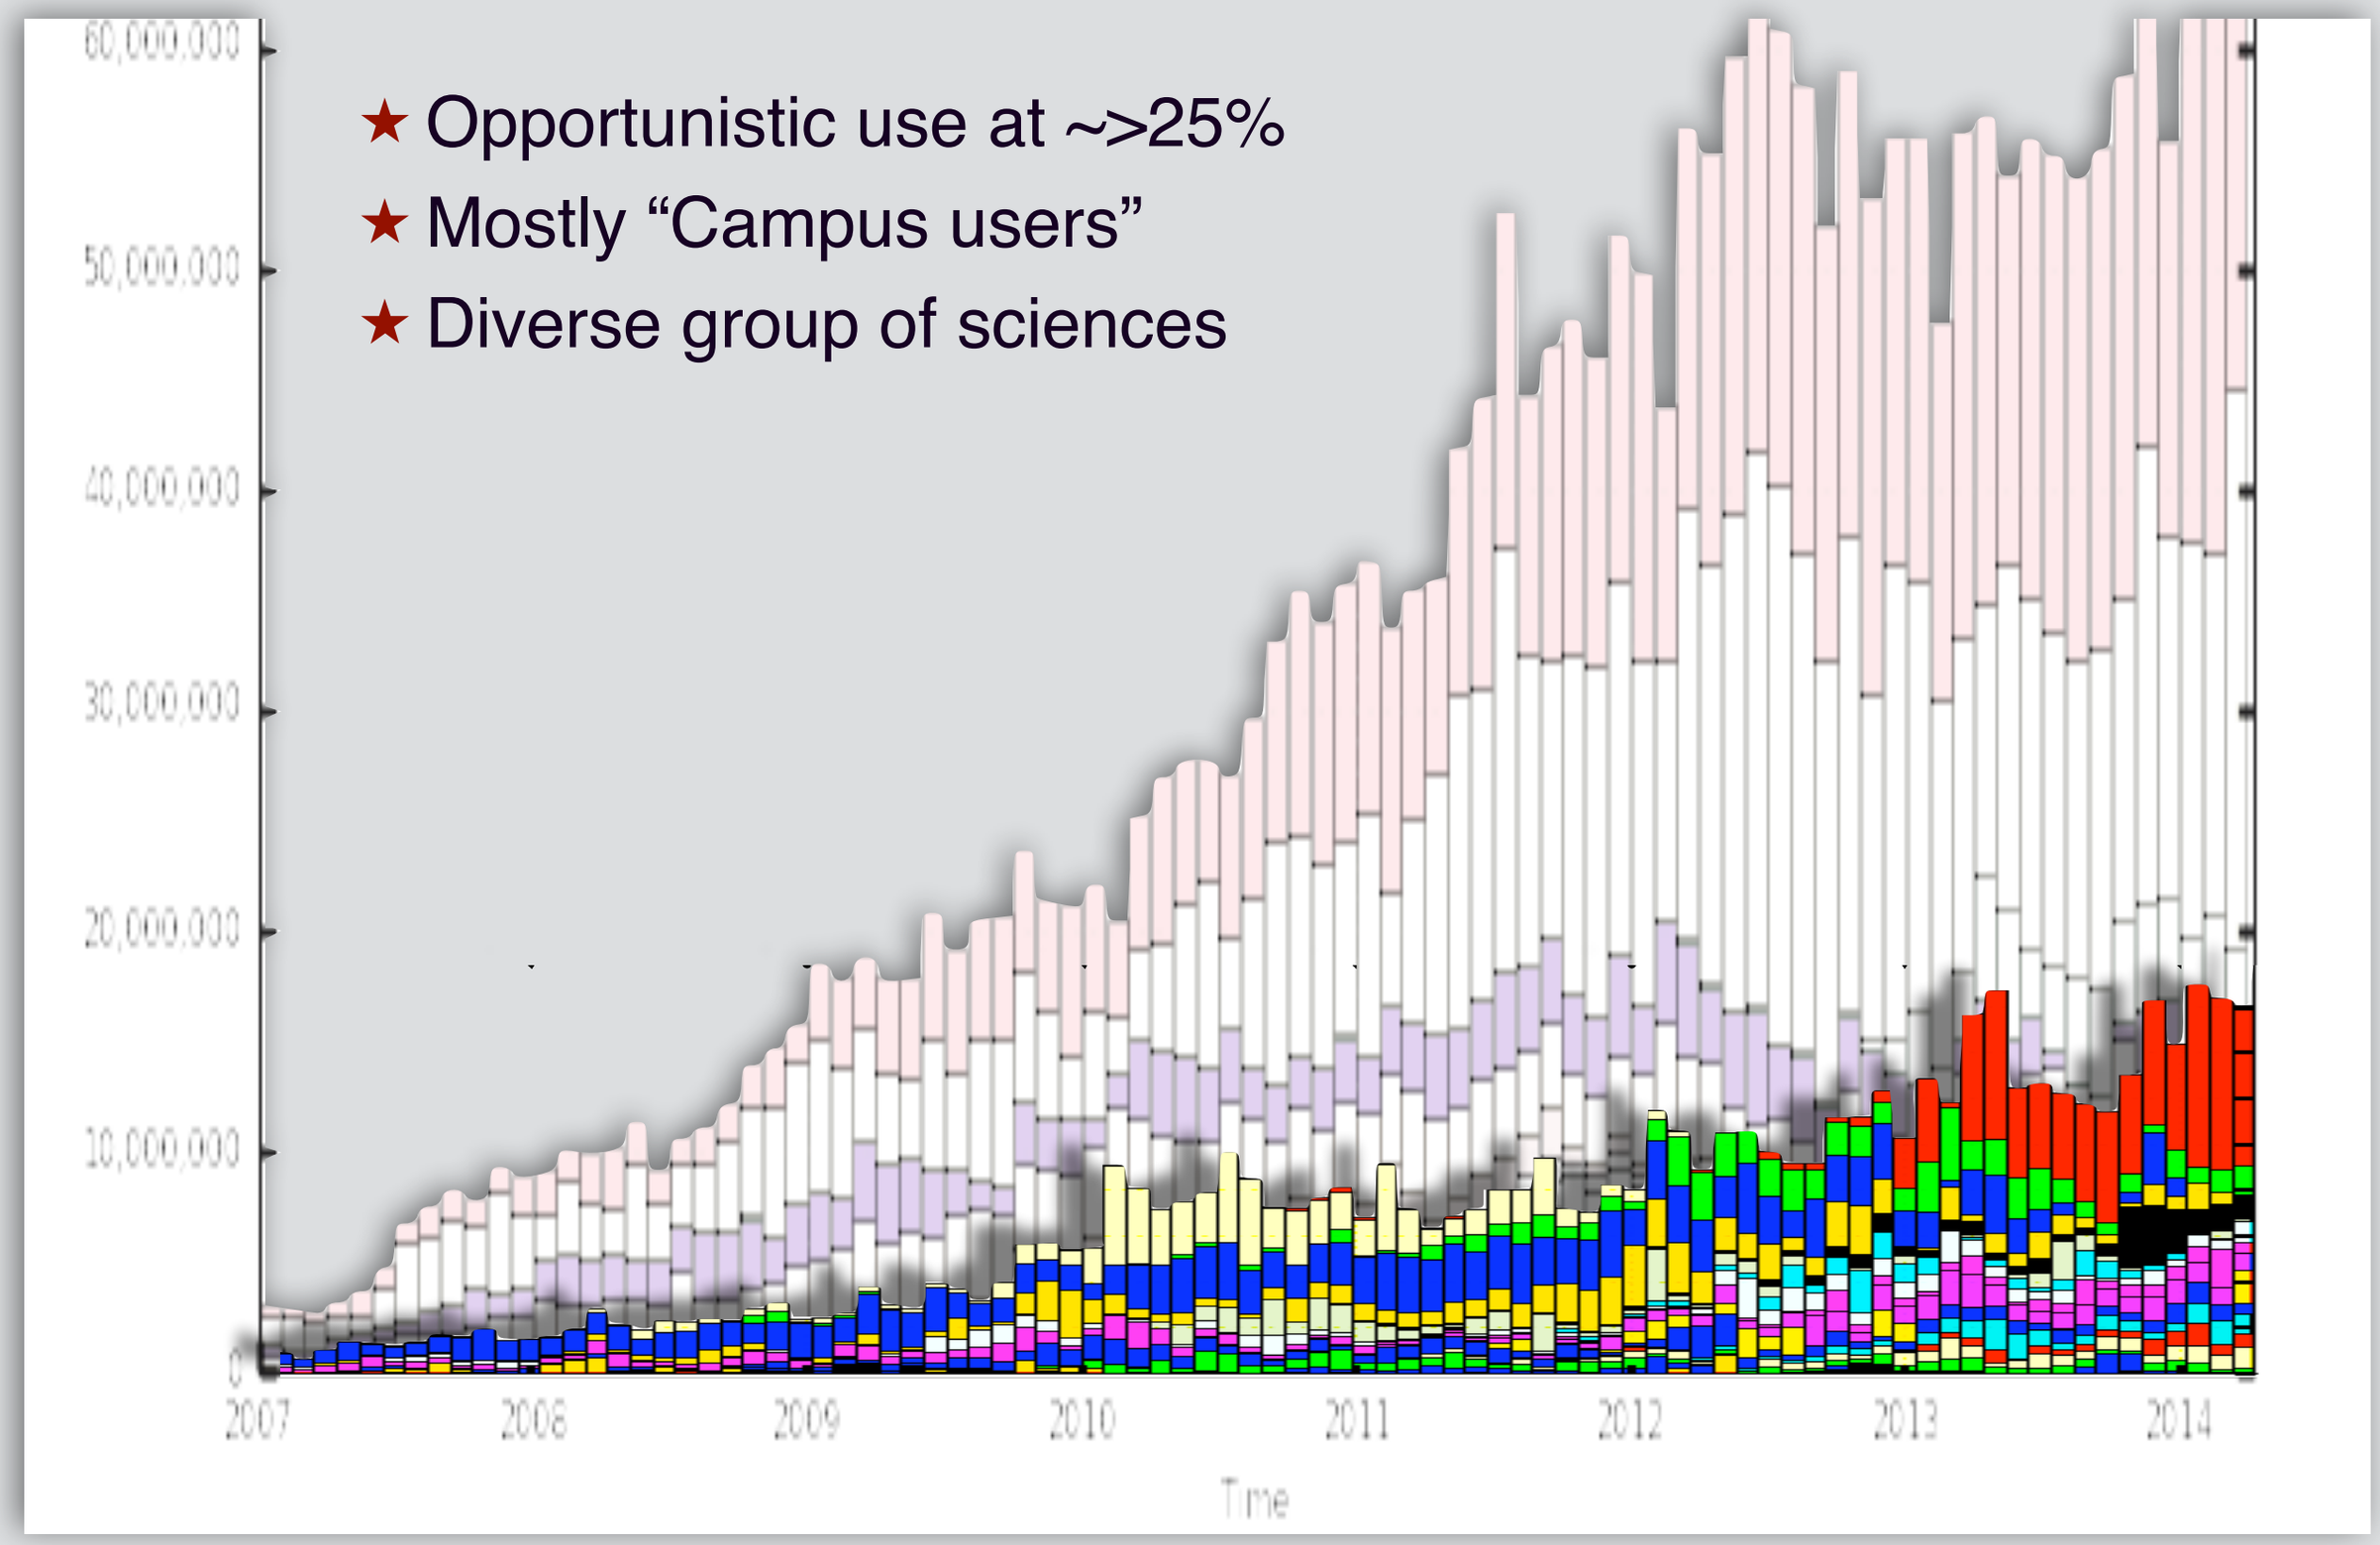

OSG Future Developments:

Evolving our approach to Campus Grids and Opportunistic Use

Enabling scientists and campuses for easy access to the OSG DHTC facility is a great value proposition that we will develop further!

OSG Connect: Focus on Services

- ★ Core idea: provide a set of services that OSG runs for campuses and users

✦ OSG Connect "encapsulates" basic infrastructure services:

- ★ from the network services (Globus online etc), to identity management services (bring you campus identity or your Globus ID, we'll map it to your OSG identity)
- ★ allow existing identities of researchers to be used across full set of services
- ★ allow research communities to select services or tools that they require, without additional complexity
- ★ setup easy access to a variety of facilities and resources

✦ Allows Campuses to create Campus Grids

- ★ provide campuses with the services to connect to the OSG and a useful set of services that help setting up a campus grid or extending an application into a campus grid
- ★ campuses don't have to build their own — OSG is running services for sites
- ★ initial set of services exists, two new campus grids being established
- ★ getting some traction also in Atlas and CMS as a solution for Tier-3 etc

OSG Connect as a Vehicle

- Extending the OSG eco system to the HPC community
  - OSG is a XD service provider,
    how about XD providing resources to OSG communities
  - interface to HPC installations using the CI Connect approach
  - Atlas is pioneering this with XSEDE (at Texas Stampede),
    - CMS is doing this with glideinWMS / BOSCO access to the NERSC DOE facilities
  - HPC as a data producer, OSG for user analysis
    - e.g. cosmology simulations
- Establishing Atlas and CMS Tier-3 centers this way
  - set up the list of services on “CI Connect” installation
  - then connect to campuses and resources
  - Is this also a model for future “VO facilities”, that bring in new sites?
- OSG Connect could also be a vehicle to establish new service offerings
  - e.g. in the future we might provide a Panda service
    for (non-Atlas) Panda users

Bringing in New Types of Resources

Open Science Grid

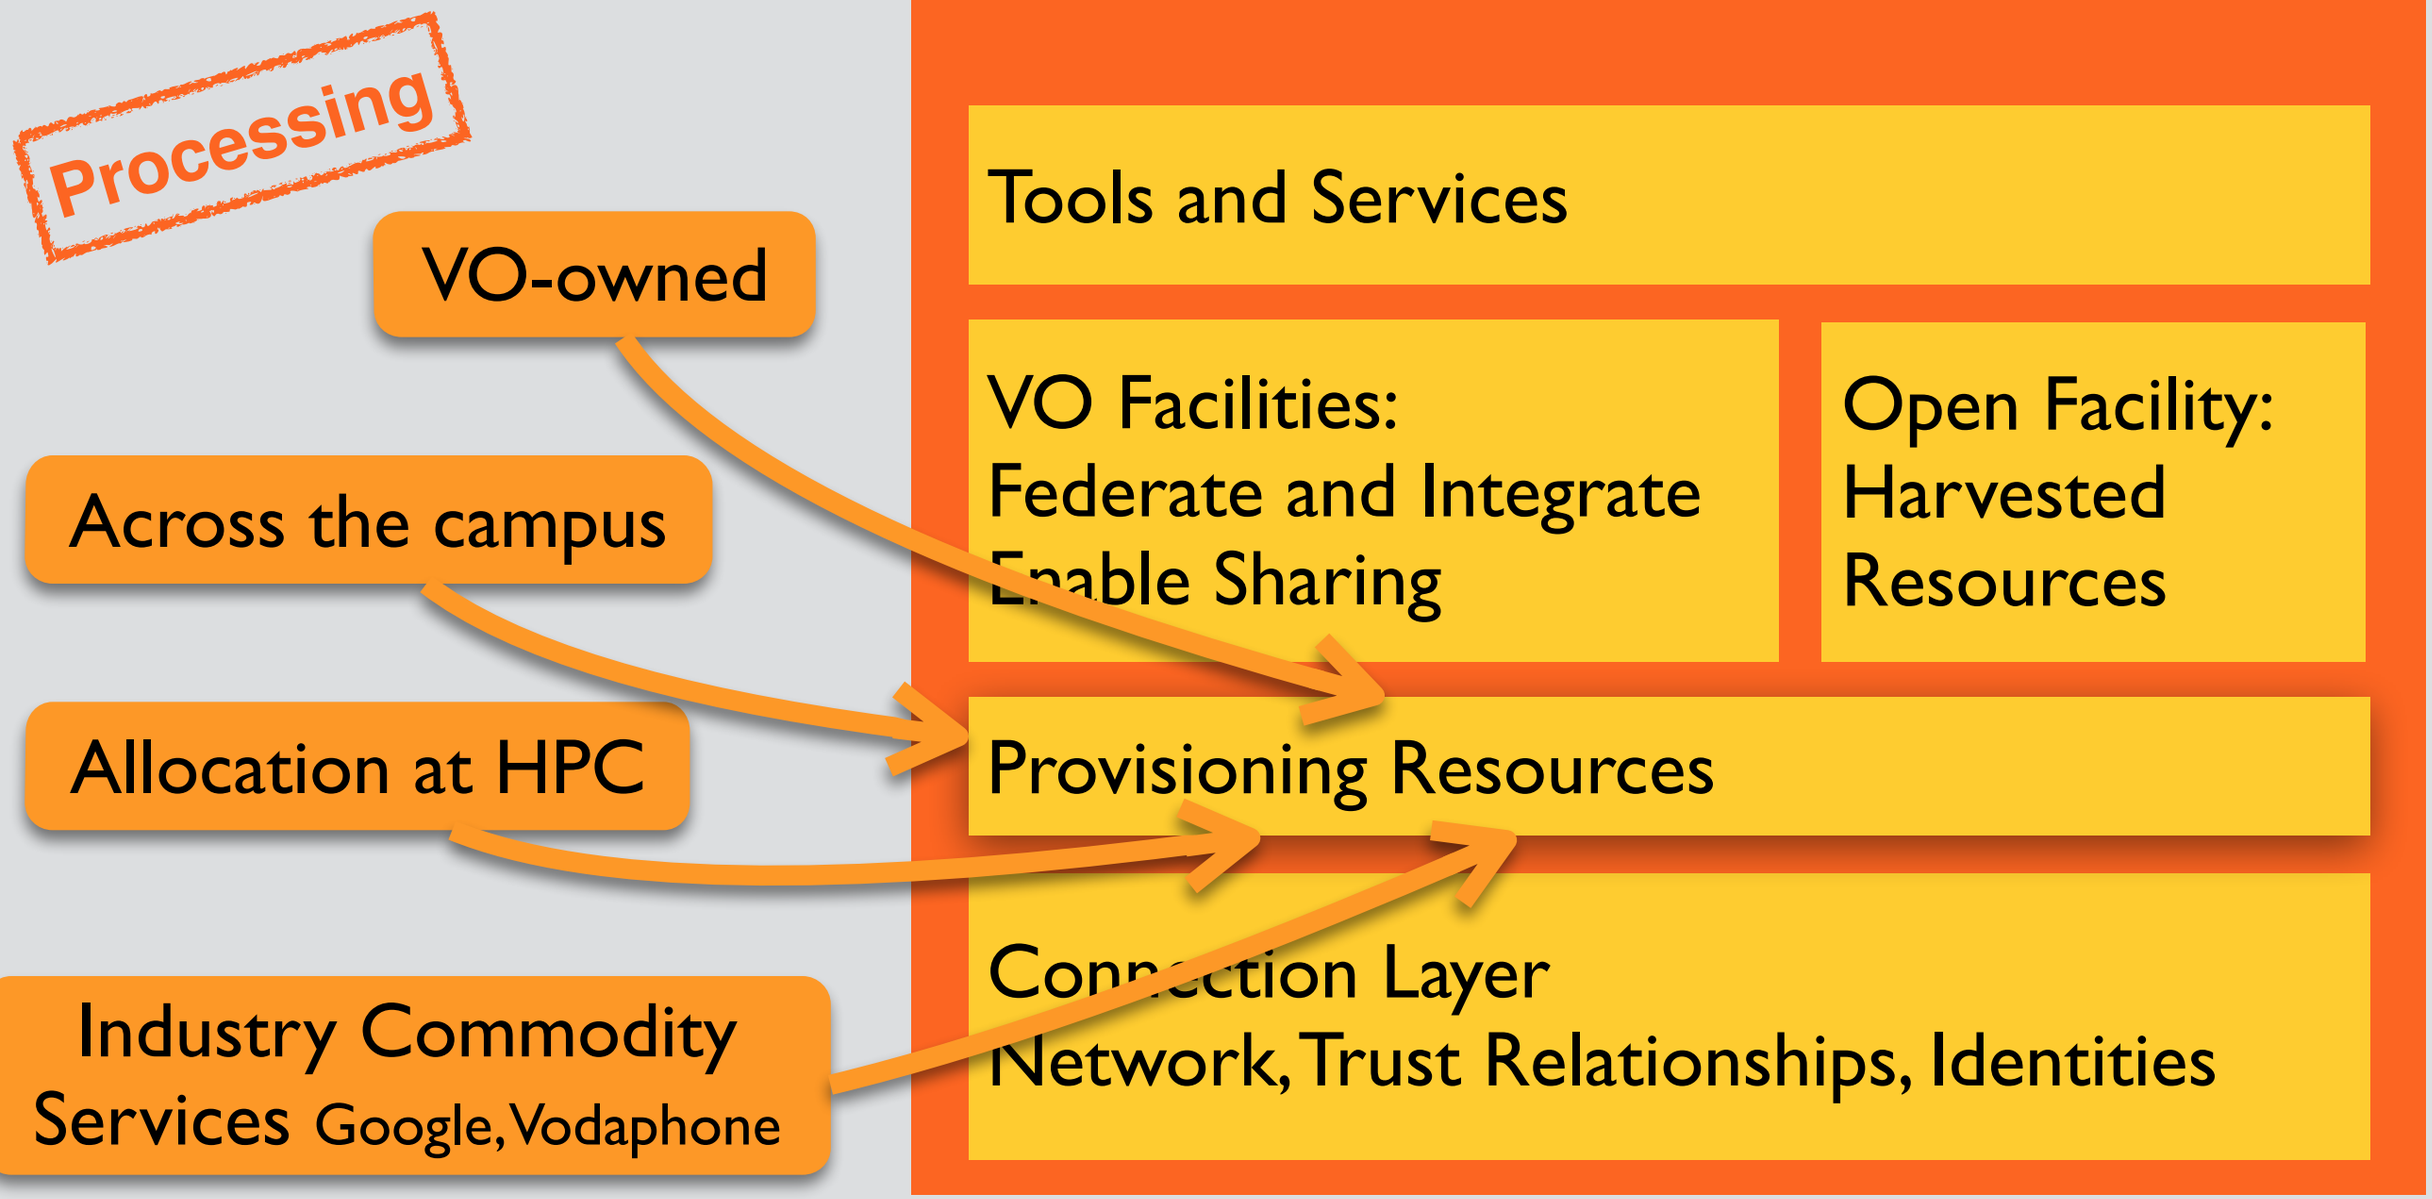

- Focus on **Dynamic Resource Provisioning**
  - ★ Statically federated resources need to be integrated with dynamically allocated resources causing new challenges for resource planning, acquisition, provisioning

Resource Provisioning to Campus Users

Open Science Grid

- OSG does not yet have a flexible “provisioning” capability
  - ★ although it has been discussed as a natural extension addressing a number of shortcomings
  - ★ raises new questions of liability, accountability, audit-ability for the “money” spent?
  - ★ can this model be extended to “enabling acquisition” of commercial cycles o behalf of an OSG member
  - ★ maybe even be extended to the provisioning of opportunistic resources
- Many Open Questions that OSG will need to address in the 2nd half
  - ★ how will OSG assist VOs provision resources with different costs?
    - clouds are a “business model” that Resource Providers use to provide resource
    - for user, “costs” for using a specific compute resources are not necessarily money
      - other examples are CPU power, network locality, allocation
  - ★ how can OSG act as an intermediary, between VOs and resources?
    - in this model, instead of OSG going to NERSC to ask for an allocation for it’s users, e.g. Atlas has the allocation and OSG helps to manage and provision it to the Atlas workflows

Open Science Grid

OSG providing support to Data Infrastructure

- for the Large VOs (LHC exp) the OSG eco-system offers its users sophisticated data infrastructures
  - dynamic data placements and high throughput robust data transfers
  - data discovery and metadata tools
- The real killer feature: Connectivity and Network Throughput
  - including high-speed networks to resources, distributed IDM, etc
  - enabling remote data access, federating storage systems and allowing **global access to locally managed data**
- OSG has not yet succeeded (nor yet really tried) to make a full-fledged data service available to the “long tail”
  - important use case of common input data cache, currently based on iRods, is being further developed
  - OSG Connect has a rudimentary “stash” service
  - a simple general data archiving service (tapes) was started at Fermilab
- We will do more work in this area during years 3-5

Success of OSG Campus Program is about minimizing obstacles for users

- ★ once users adapt they don't look back (Rob Gardner)

- ✦ So OSG needs to be at its best game to convince users to adapt
  - ★ while keeping a focus on the LHC and the start of Run2
- ✦ We have a great toolset and an emerging understanding of what works and what does not work so well:
  - ★ OSG Connect, OSG-VO, BOSCO and seamless access to the overlay job execution environments, we have established OSG as a XSEDE Service Provider, remote data access technology allowing local data management of globally accessible data, we have a program of going away from user certs, etc
- ✦ These all are great ways to minimize impedance and get campus users to "take the gateway drug and get high" with their science on DHTC
- ✦ We'd like to strengthen partnerships: Campuses, ACI-REF, XD/XSEDE, satellite projects, and re-invigorate relationships with EU projects
  - ★ discussion about EU-T0 "framework for collaboration" at Barcelona

Questions/Discussions?

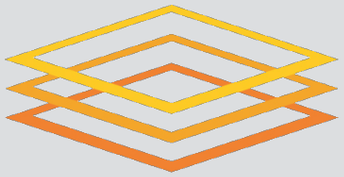

OSG for the "Long Tail of Science"

- "long tail" of researchers that do not have access to sufficient in-house computing resources and/or skills, or that profit from “more”
  - “OSG will be considered successful if we not only keep up with the growing LHC needs but also expand to enhance the computing throughput to a broad spectrum of scientists at a variety of scales, from individual users at a single campus to multi-institutional experiments”
- We see two extremes of such challenges
  - On one extreme, there are the large collaborations in HEP that need innovation to help the transition from petabytes to exabytes of data volumes, and from 10,000 cores to 100M cores within the global DHTC systems over the next ten years
  - On the other extreme, there are an ever increasing number and diversity of small to mid-size collaborations of domain scientists that struggle with crossing order of magnitude boundaries in their computing — transitioning from gigabytes to terabytes of data volumes at scales of 1-10 million CPU hours
  - Over the next 10 years, we expect the number and diversity of groups at this end of the spectrum to increase as data and compute intensive science becomes the norm rather than the exception

Survey of Who Provides Resources that Researchers Utilize

Open Science Grid

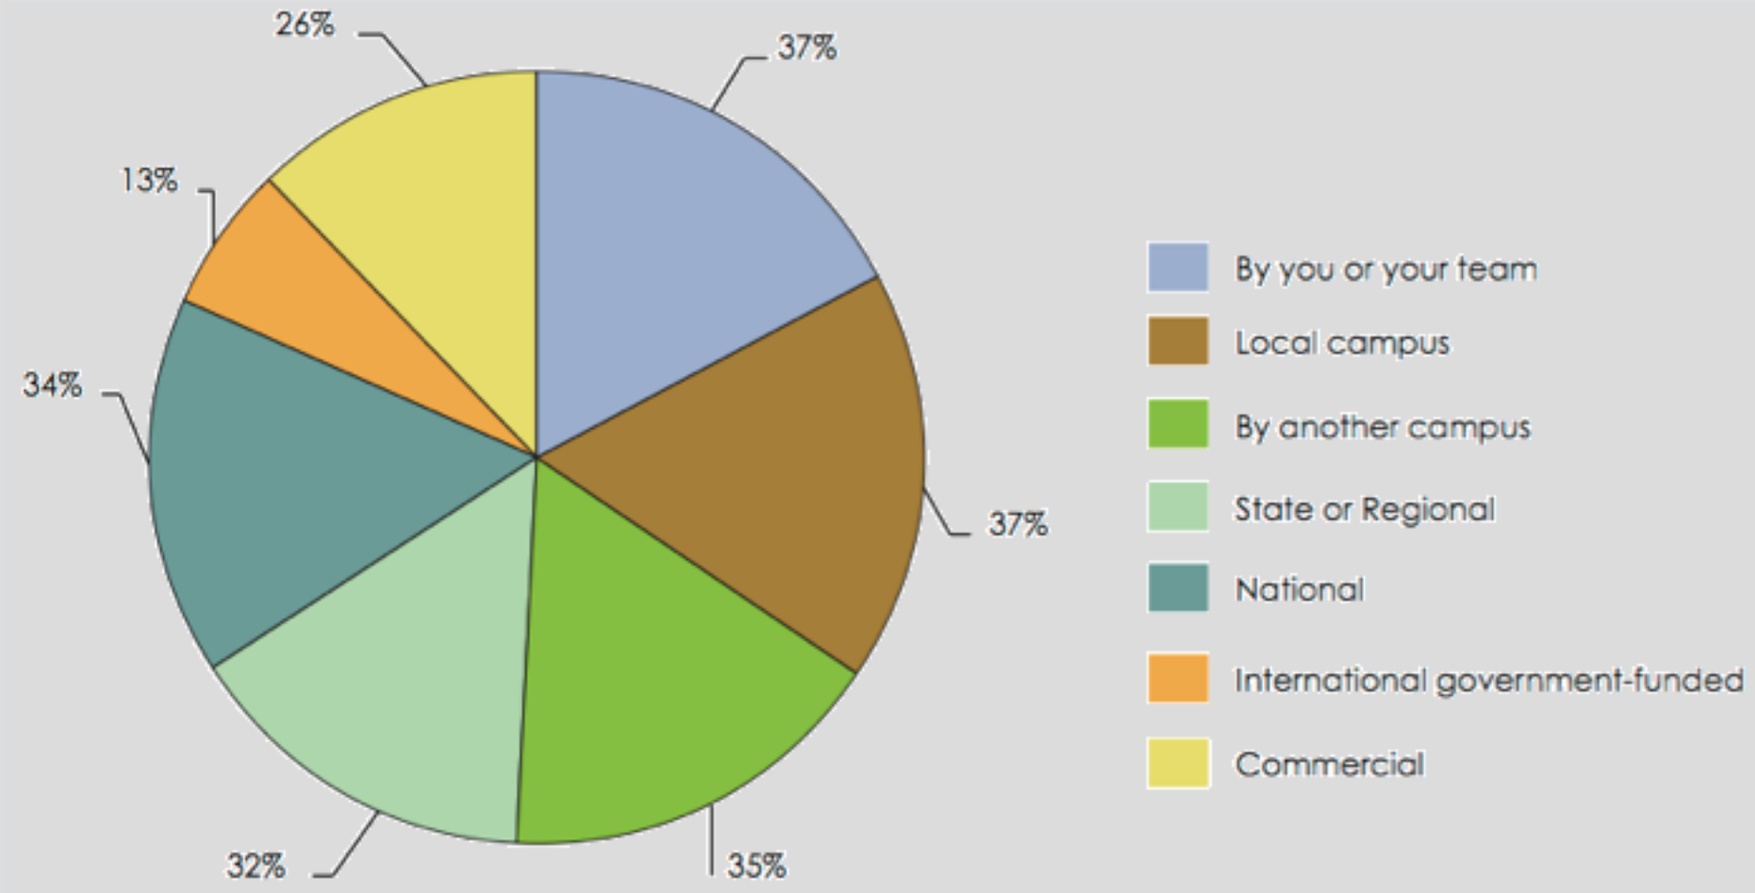

Figure and data from Campus Bridging: Software and Service Issues Workshop Report, McFee et al, http://hdl.handle.net/2022/13070

★ Responses to asking if researchers had sufficient access to cyber infrastructure resources

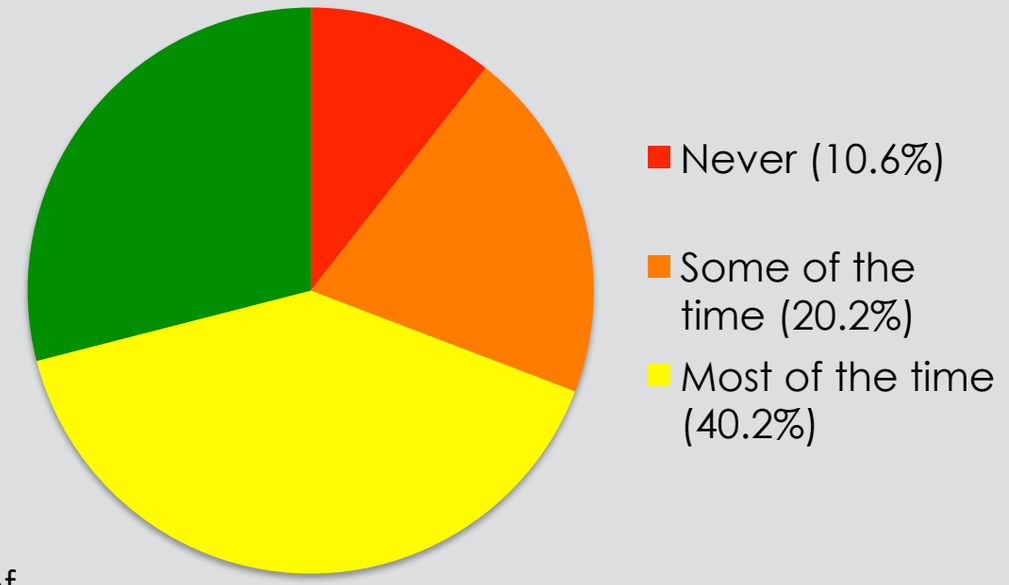

C.A. Stewart et al: Survey of cyberinfrastructure needs and interests of NSF-funded principal investigators. 2011. hdl.handle.net/2022/9917

The Gap

- In our assessment, it is crucial to prevent the capability gap between these two extremes from growing. While “exascale” problems need to be solved for the large collaborations, it is equally necessary to ensure solutions are available for the many scientists challenged at the terascale and petascale on the passage to the exascale.

- To help our growing spectrum of users, we need to keep up with increasingly dynamic and heterogeneous environments while ensuring that domain scientists with limited computing expertise can use those environments

- We believe that the Campuses will have to play a role in filling this gap and supporting scientists and groups to cross over to DHTC

OSG Services and Facility Ecosystem

Open Science Grid

VO-centric Model

Campus User Model