The Open Science Grid

HEPAP Meeting

July 13th, 2007

Frank Würthwein (UCSD)

As an aside: fkw's roles in OSG

- OSG Applications Coordinator (with Torre Wenaus (BNL))
- Co-lead of "CMS Computing Commissioning" (with Stefano Belforte (INFN))
- DISUN technical lead
  - DISUN = 4 sites on OSG plus infrastructure commissioning effort
- User of OSG as CDF physicist

My perspective on OSG is, of course, influenced by these roles.

Outline

- OSG Vision
- OSG facts
  - Organizational
  - Utility to HEP thus far
- Some of the Challenges
  - Commissioning LHC Computing
  - Cybersecurity
  - “Tier-3” program
  - Data Movement
  - Managing Change
  - Engaging new communities
- Summary & Conclusion

OSG Relevance to HENP

- Distributed Computing Infrastructure for the LHC.
- Significant resource for the Tevatron program …
- … and many others, ongoing as well as planned.
  - STAR, MiniBooNE, geant4, ILC, …

The LHC Problem

-- personal take --

- Hundreds of Institutions across many tens of countries across many continents collaborate on 4 experiments.
- Many tens of Millions of $$ worth of computing resources distributed (almost) as widely as the human resources.
- *To make this work requires new Technologies and organizational structures.*
- We are *turning "problem" into "virtue",* creating new institutions for the benefit of a broader scientific audience in ways that's never been done before.

Open Science Grid Vision

"Transform processing and **data** intensive science through a cross-domain self-managed **national distributed** cyber-infrastructure that **brings together** campus and community infrastructure and facilitating the needs of **Virtual Organizations** (VO) at all scales."

(Miron Livny, OSG-PI, SciDAC June 2007)

OSG - forming communities

Science Drivers

(HENP, LIGO, BioTech, NanoTech, …)

OSG enables community formation to solve compute and data intensive scientific problems.

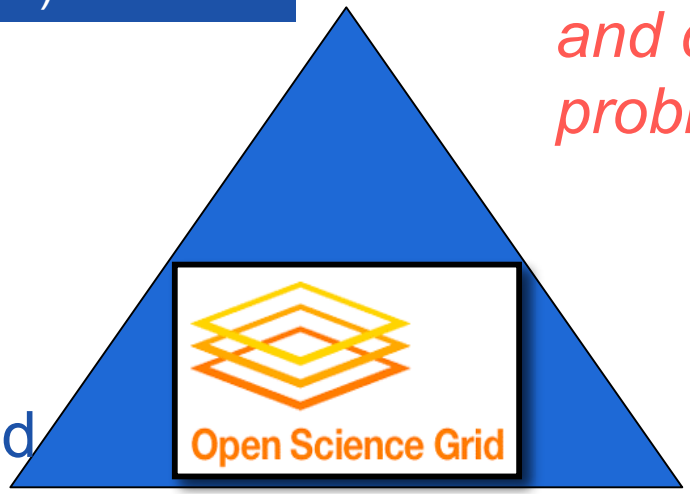

IT shops at Universities and National Labs

(e.g BNL, FNAL, LBNL, SLAC, LHC-T2s, DISUN, …)

Computer Science

(e.g Condor, Globus, SRM, …)

OSG - from local to global

National & International
CyberInfrastructure
for Science (e.g. Teragrid, EGEE, …)

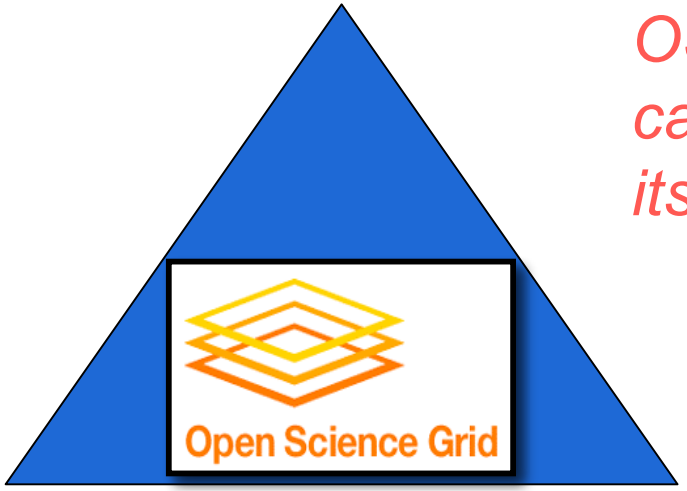

OSG harmonizes community, campus & national CI to enable its users to operate globally.

CS/IT
Campus
Grids

(e.g DOSAR, Fermigrid, GLOW, GPN, GROW, NYSGrid, NWICG, …)

Science
Community
Infrastructure

(e.g Atlas, CMS, LIGO…)

Open Science Grid Facts

OSG Evolution

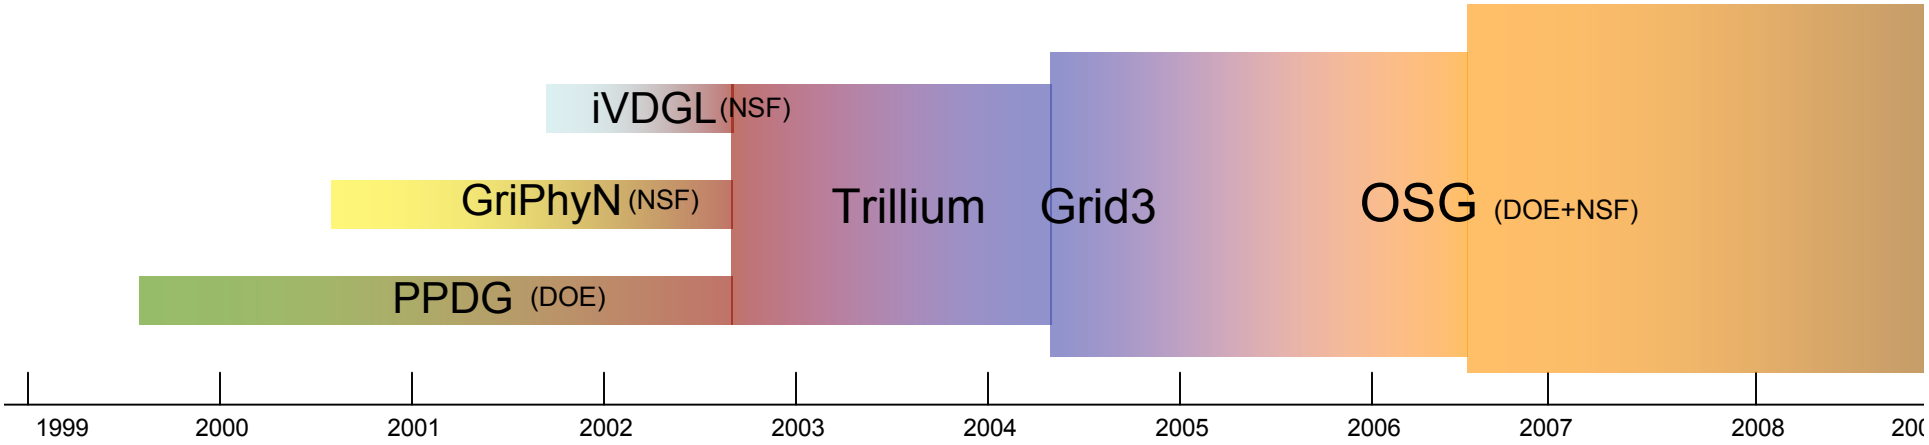

OSG *Consortium* since ~2005 → form collaborations
- >100 institutions
- ~30 user organizations

OSG *Project* since Fall 2006 → sustain infrastructure
- $30 M across 5 years
- ~17 Institutions
- ~34 FTE

OSG Consortium

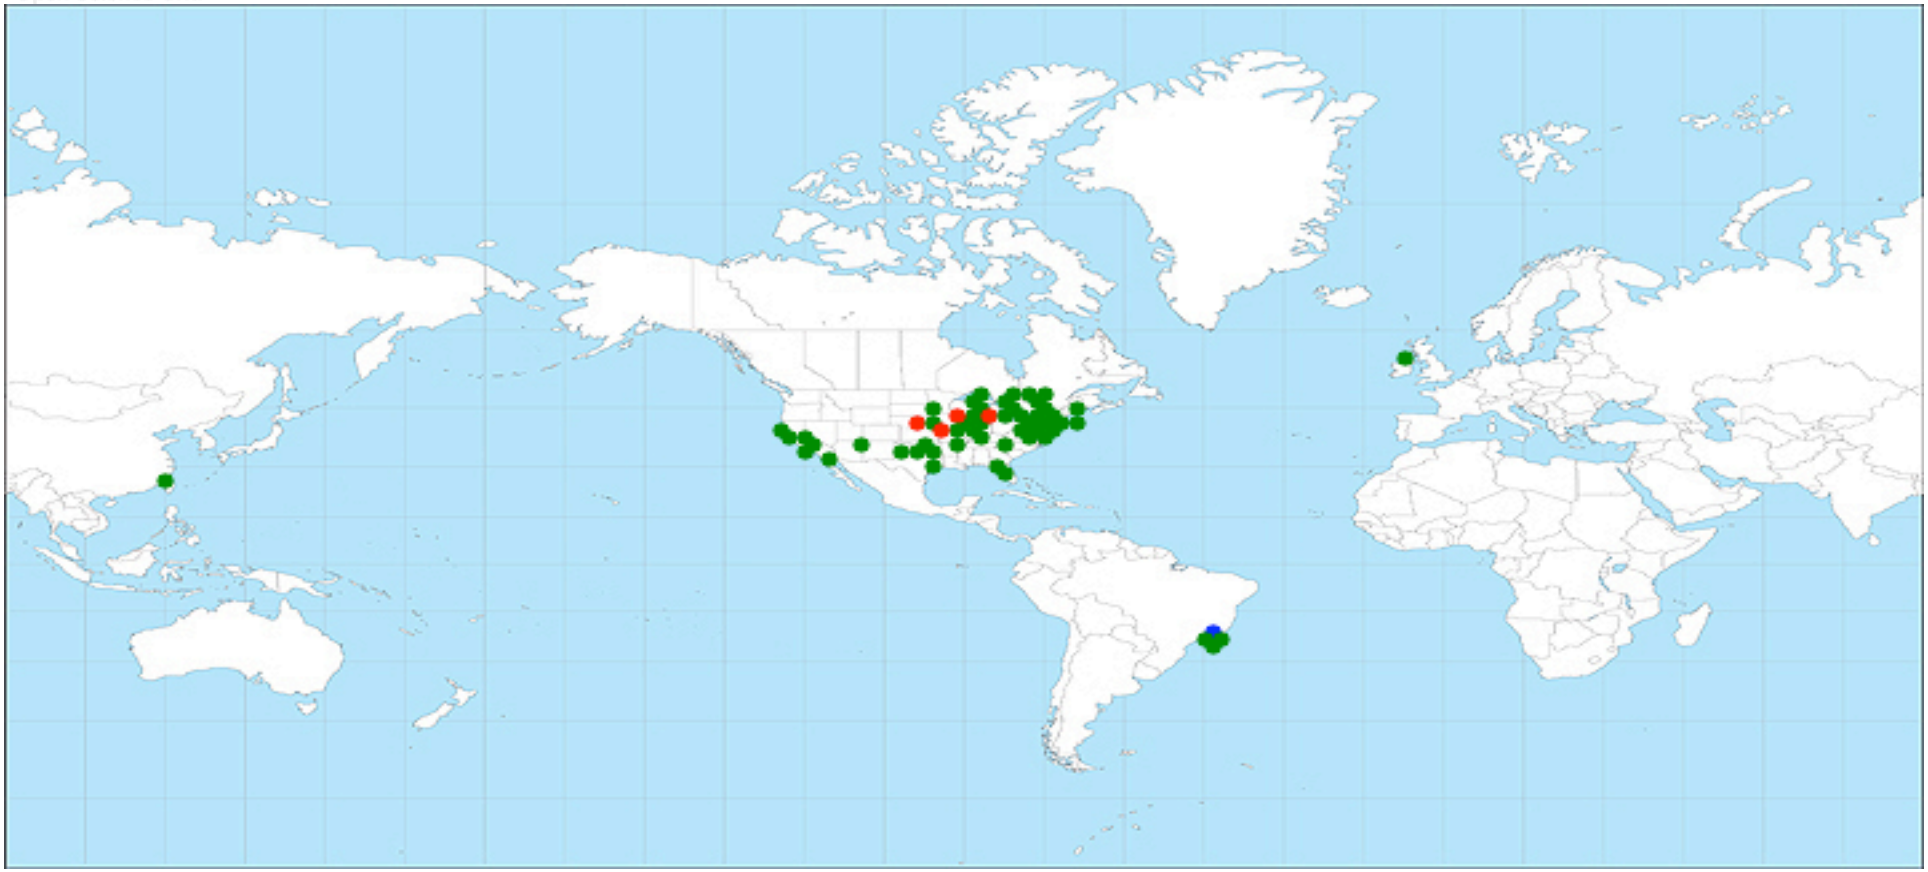

103 Resources registered on the grid.

30 user organization registered, among them:

7 HENP, 3 Astro, 9 Campus or regional, 5 bio/nano tech, …

OSG Project Effort

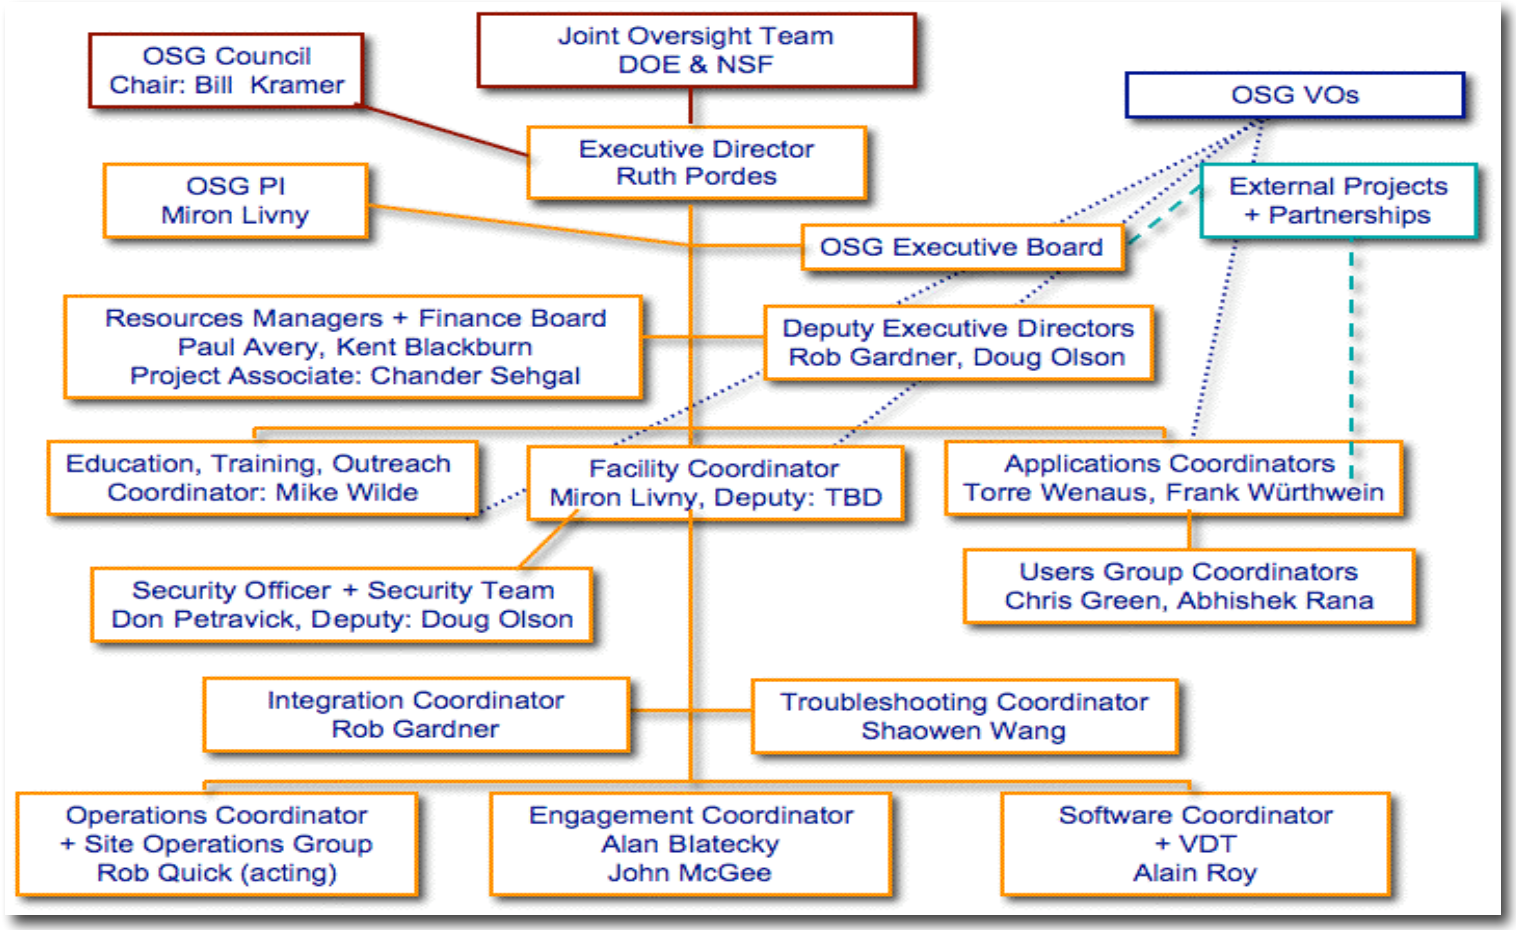

Roughly 2/3 of leadership positions filled from outside HEP !

Use of OSG last month

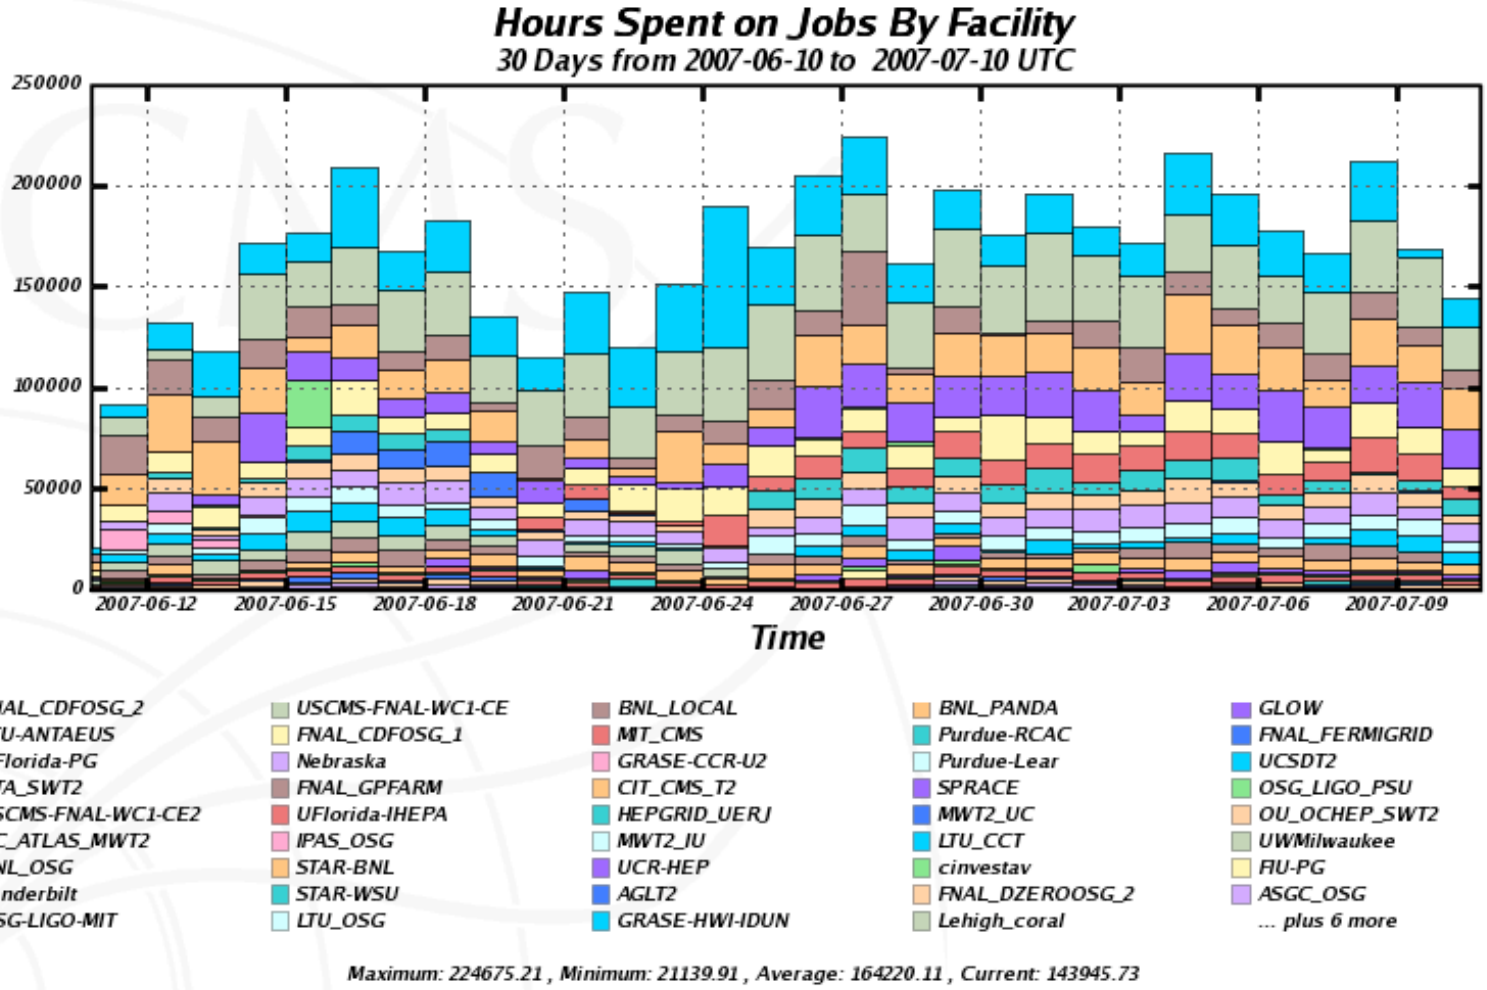

Routinely providing >200k hours of computing per day.

Utility to HEP last month

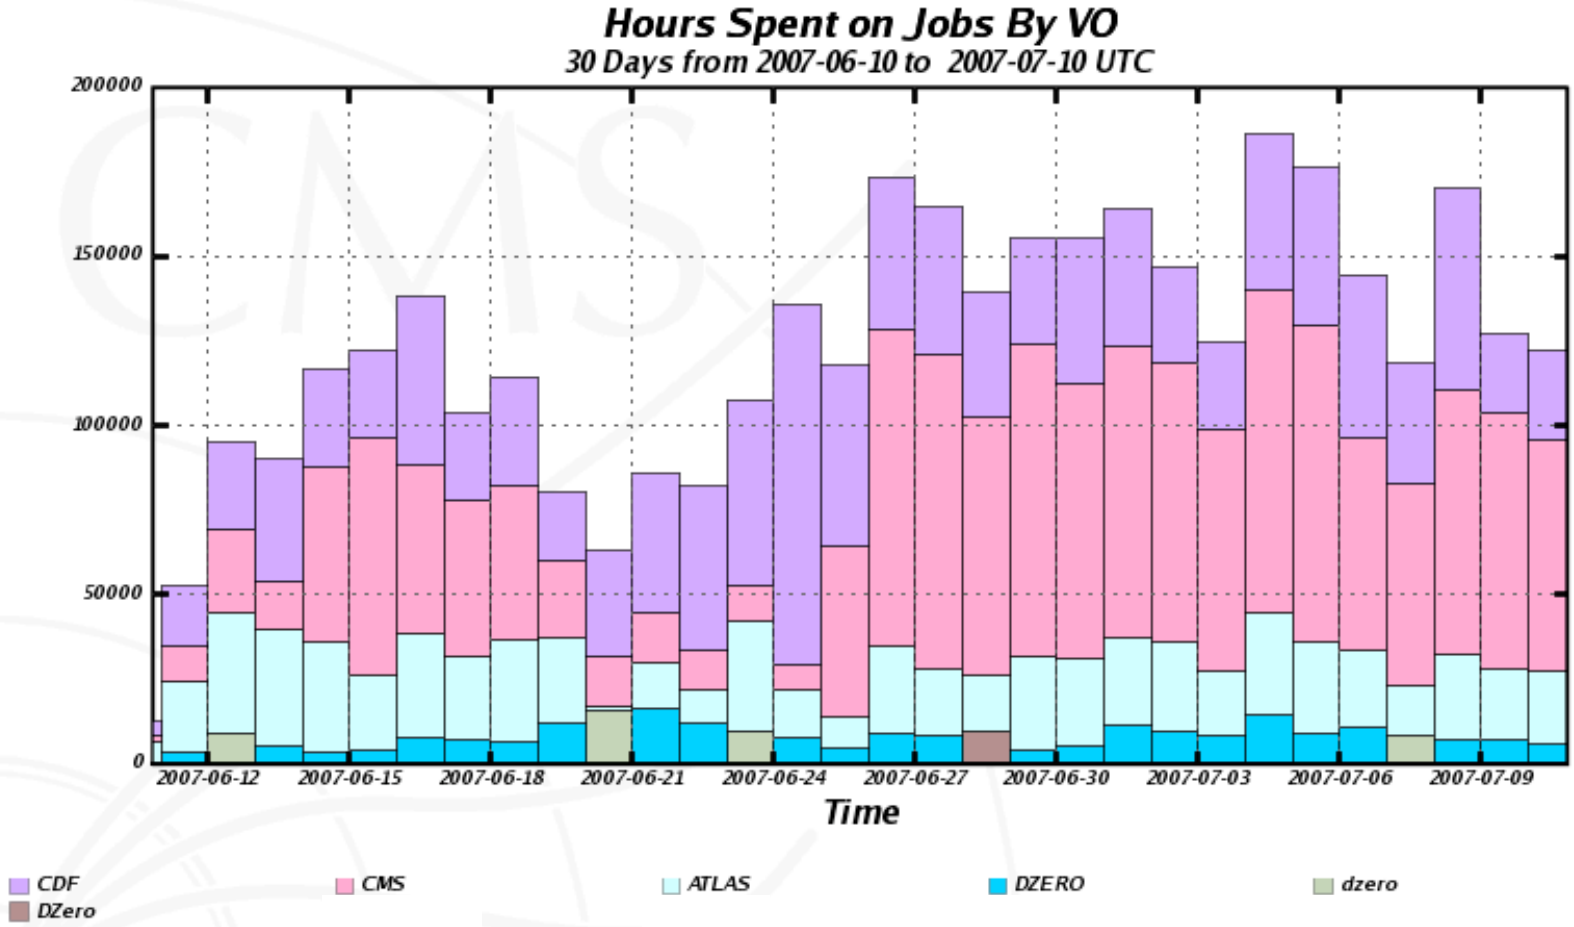

Routinely providing >150k hours of computing per day to HEP.

(Plot is misleading as not all Atlas consumption accounted yet in OSG accounting system!)

Utility to HEP so far

- Middleware stack for Atlas and CMS in US, Brazil, and a few other places.
  - Commissioning of the distributed computing systems for Atlas and CMS.
- Tevatron program
  - CDF for MC production
  - D0 for reprocessing
- MiniBooNE, …
  - Use of the FNAL campus grid

Example: D0 re-processing

In Nov '06 the D0 experiment asked the OSG to use 1.5-2K CPUs for 2-4 months for re-processing of a dataset (~500M events) for the summer conferences in July '07. D0's own resources were committed to the processing of newly acquired data and analysis of already processed datasets.

By the end of May, re-processing of 445M events was completed. OSG contribution to this effort was

- 286M events
- 286k jobs executed
- 2M CPU hours
- 48TB of input data
- 286M files of final results
- 22TB of output data

How did the D0 Reprocessing happen?

- The Executive Board estimated there were currently sufficient opportunistically available resources on OSG to meet the request; **We also looked into the local storage and I/O needs.**
- The Council members agreed to contribute resources (processing, data and FTEs) to meet this request.
- D0 had 2-3 months of smooth production using >1,000 CPUs.
- To achieve this
  - D0 testing of the integrated software stack took until February.
  - OSG and D0 staff then worked closely together to reach the needed throughput goals - facing and solving problems
    - sites - hardware, connectivity, software configurations
    - application software - performance, error recovery
    - scheduling of jobs to a changing mix of available resources.

D0 Re-Processing

More than 50% of reprocessing done on OSG …

Total Events

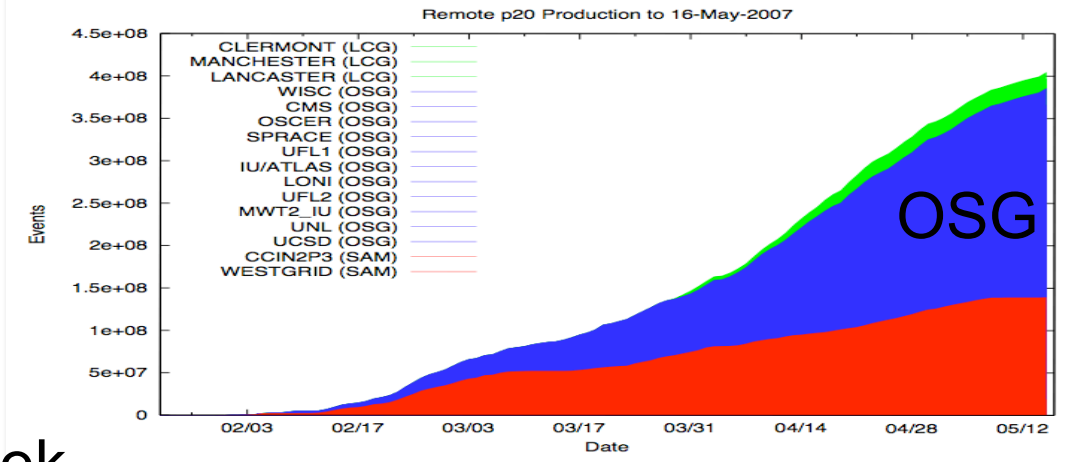

OSG CPUHours / Week

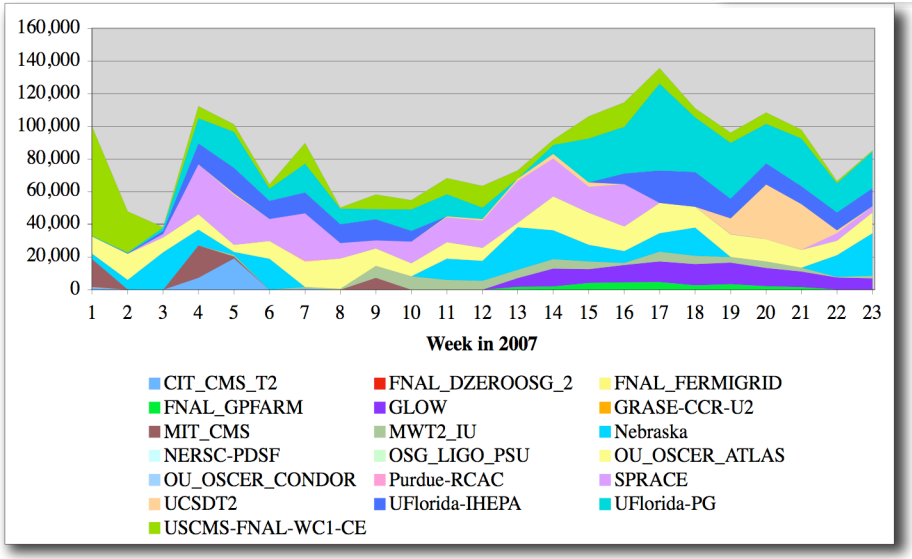

.. most of it on resources that would not have been accessible to D0 without OSG.

Obvious,
and not so obvious challenges

Commissioning LHC Computing

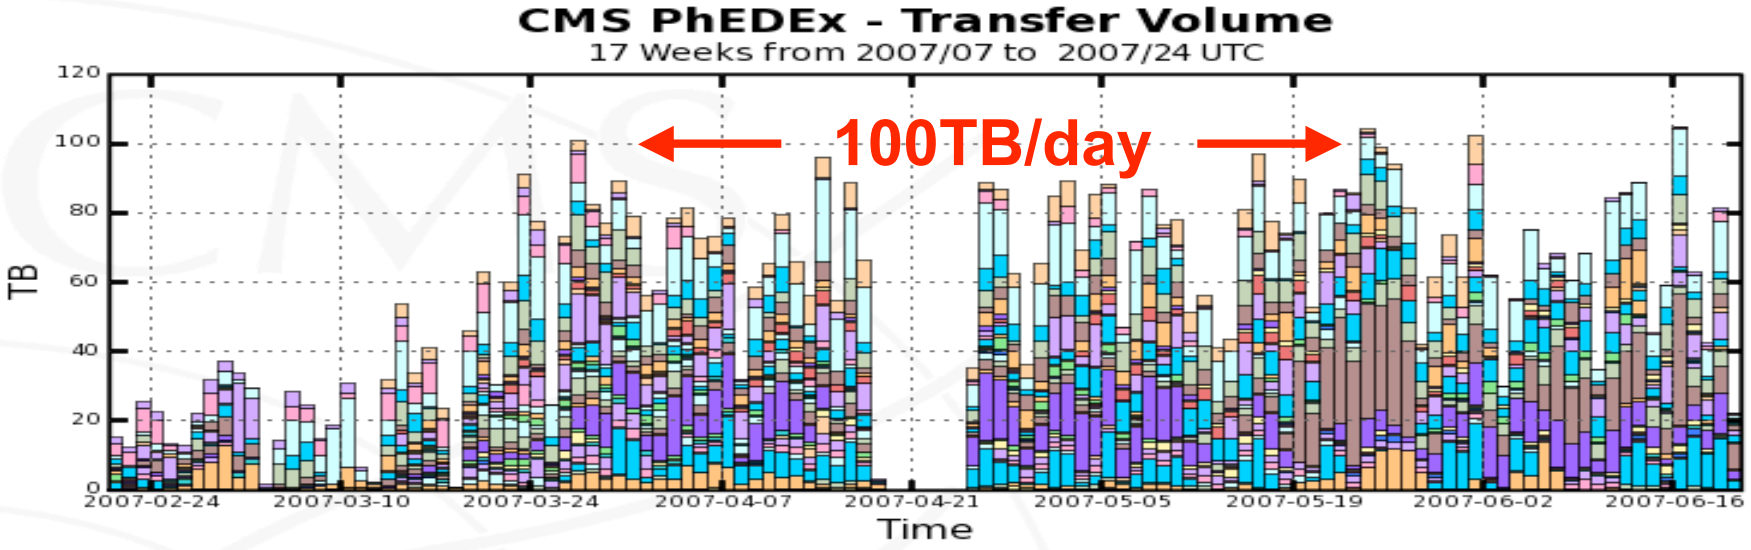

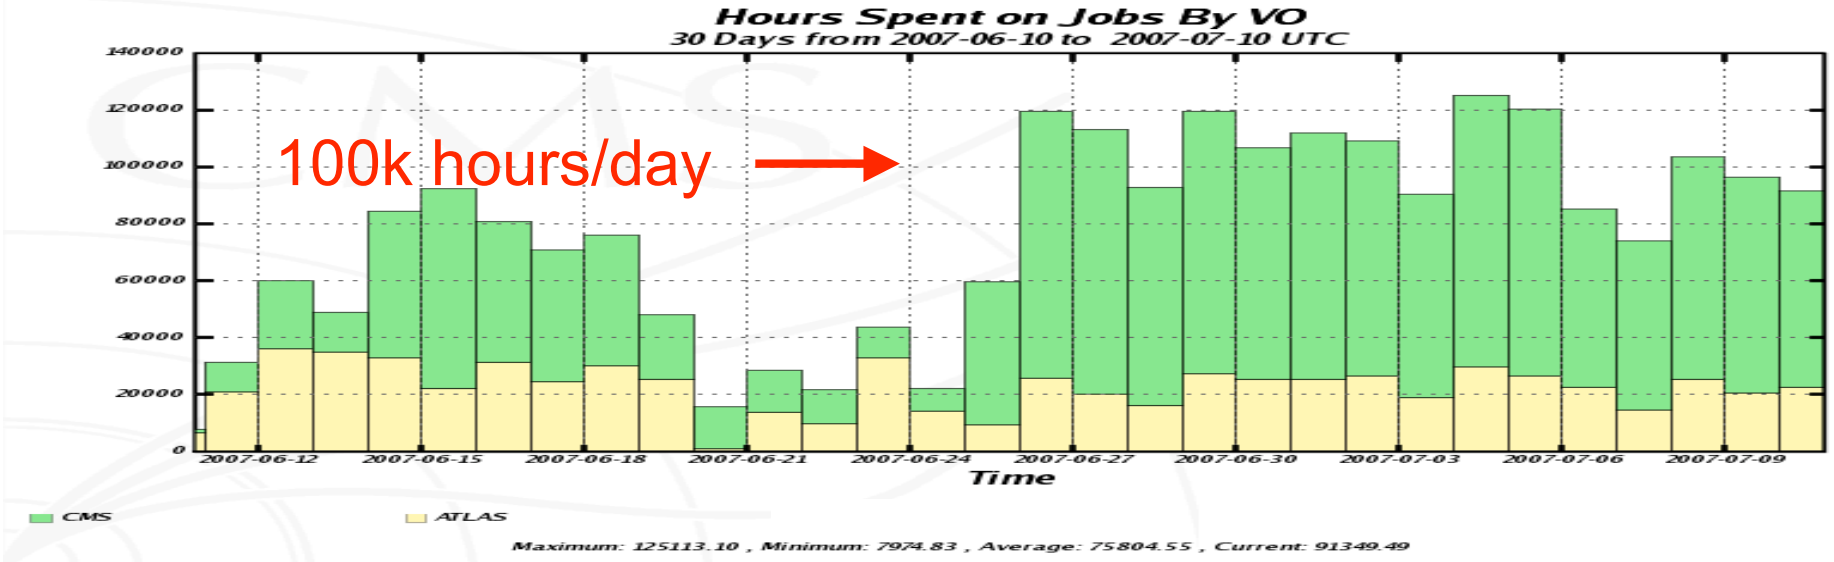

Challenges easily overlooked

- Cybersecurity
  - Day-to-day operations of cybersecurity
  - Global organization of cybersecurity
- The LHC “Tier-3 problem”
  - Significant resources owned by local groups at Universities.
  - How best to integrate into global LHC computing infrastructure?
- Managing Change
  - LHC program needs to be able to adjust to changes in technology without disruption to physics program.
- Engaging New Communities
  - Adapt legacy systems to common infrastructure

OSG is a vehicle to address all of these, in addition to the obvious challenges!

OSG Engagement

- Embedded assistance to bootstrap research groups
- Assist researchers by adapting existing job submission and management scripts to utilize OSG
  - Most researchers already have job mgmt scripts for local resources, the engagement team brings the OSG expertise into those scripts, and transfers knowledge during that process
- Develop and maintain a hosted infrastructure to enable scientists to ease into becoming full fledged OSG partners
- Example: Protein design research group at UNC-CH
  - Within two weeks, the research team was self sufficient and consuming >150k cpu hours for real science. Very low impact on existing processes. Successful engagement led to two new research groups seeking similar assistance

Summary & Conclusion

“don’t” and “do” of OSG

- Does not –
  - The OSG Facility does not “own” any compute (processing, storage and communication) resources
  - The OSG Facility does not “own” any middleware
  - The OSG Facility does not fund any site or VO administration/operation personal
- Does –
  - Help sites join the OSG facility and enable effective guaranteed and opportunistic usage of their resources (including data) by remote users
  - Help VOs join the OSG facility and enable effective guaranteed and opportunistic harnessing of remote resources (including data)
  - Maintains and supports an integrated software stack that meets the needs of the stakeholders of the OSG consortium
  - Reaches out to non-HEP communities to help them use the OSG
  - Train new users, administrators, and software developers

What Can the OSG Offer?

- Middleware that provides dependable “horizontal” capabilities on which “vertical” (end-to-end) solutions can be built, deployed and operated.
- Organizational support
  - The OSG Consortium (brings together the stakeholders)
  - The OSG Facility (brings together resources and users)
- Technical Support
  - Support with the “bleeding edge” distributed computing technologies.
- A “bridge” that forms an integrated national cyber-infrastructure by connecting desk-tops to campus cyber-infrastructure and to national and international facilities

Benefits to HEP thus far

- LHC
  - Middleware stack for the LHC distributed computing systems of ATLAS and CMS
  - Strong partner to negotiate technical and operational problems with EGEE and Nordugrid.
  - Framework for integrating “Tier-3” resources.
- Tevatron and other FNAL based HEP
  - CDF: MC production on OSG
  - D0: reprocessing on OSG
  - Other HEP benefit via FNAL campus grid
- Other HEP starting to show interest as well.

Backup

Sites for HEP on OSG

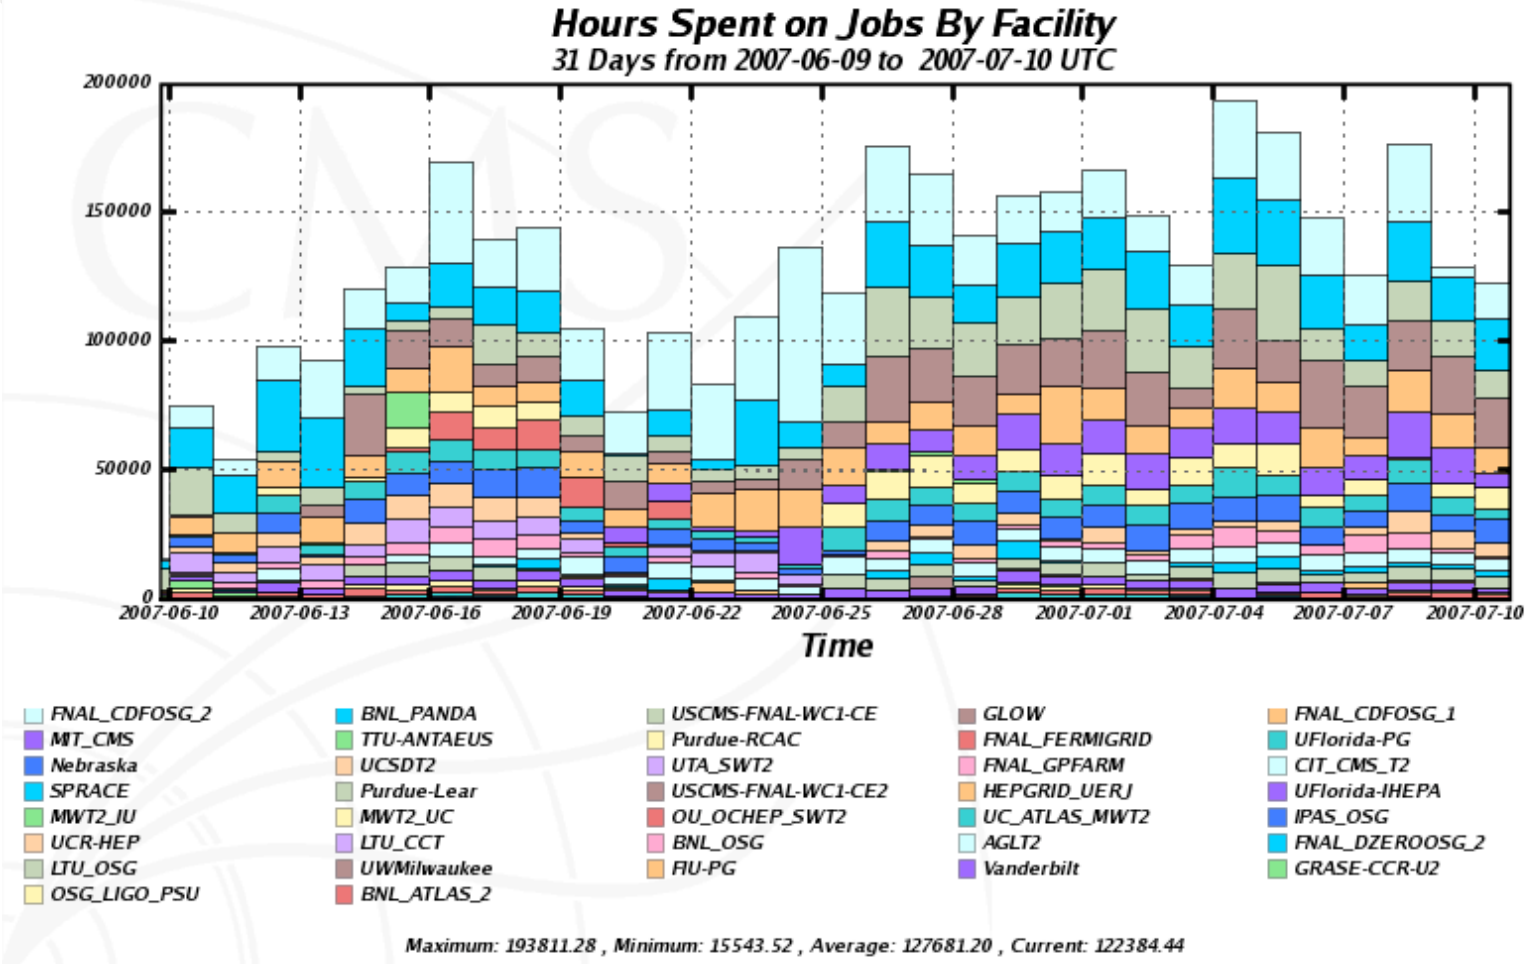

A large number of sites contribute, many of which are not LHC funded.