

# Present and Future Usage of the OSG

Brian Bockelman  
ISGC 2009

# The Open Science Grid

- The OSG provides organization and infrastructure for a cohesive national grid in the United States.
- The OSG focuses on scalability, distributed infrastructure, and
- Enabling the computational programs of its core users - USLHC, LIGO, FNAL Run II, and others.

# The OSG

- One of the strengths of the OSG is its focus on decentralized software - it's possible to use the same software at a smaller scale to deploy a campus or regional grid.
- But this decentralized nature makes it difficult to collect information about sites!

# OSG Growth

- What questions can we ask about the grid?
  - How is it used?
  - How “big” is the OSG?
  - How will it grow in the future?


# OSG Usage

- The usage is the easiest question to answer.
- The next few slides show the OSG's growth over the last 12 months.

# VO-centric View

Monthly Wall Hours per VO

52 Weeks from Week 13 of 2008 to Week 13 of 2009




1,000,000 hrs  
= 114 years

- cms
- cdf
- usatlas
- dzero
- nysgrid
- engage
- glow
- ligo
- theory
- dosar
- mipp
- miniboone
- star
- accelerator
- ilc
- minos
- fermilab
- hypercp
- cdms
- Other

Maximum: 18,733,363 , Minimum: 0.00 , Average: 12,641,874 , Current: 17,145,578

Note the top 4: USLHC and FNAL Run II.

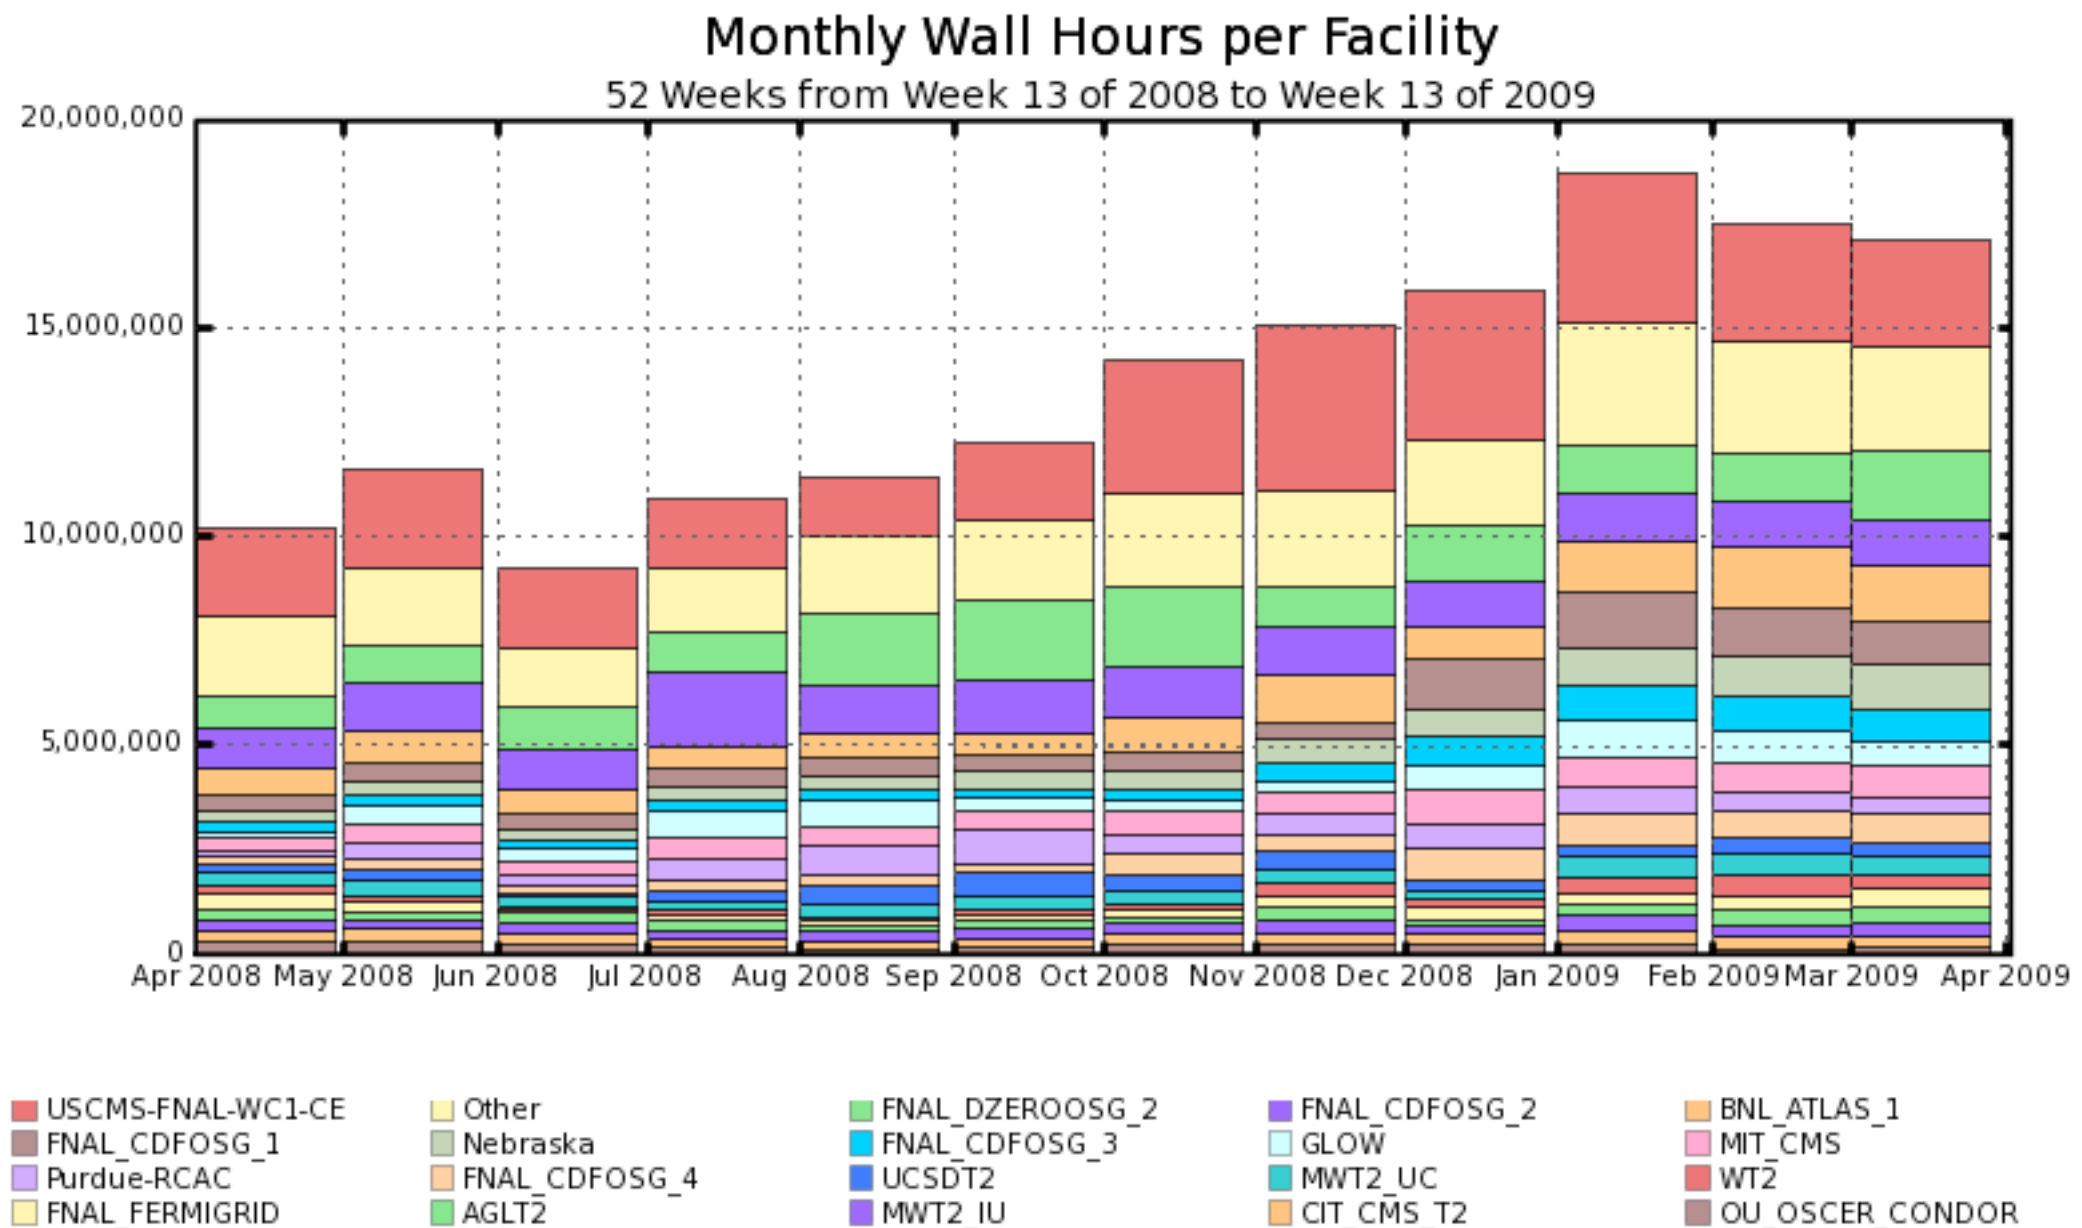
# VO-centric (Job Count)



Maximum: 11,027,137 , Minimum: 0.00 , Average: 6,805,969 , Current: 323,032

Notice ATLAS prefers a large # of short jobs - OSG infrastructure is able to handle this

# Resource-centric



Maximum: 18,733,363 , Minimum: 0.00 , Average: 12,641,963 , Current: 17,146,733


The “Other” category is the aggregate of all sites too small to list individually: the OSG has a “long tail.”



# Who are the users of Resources?

- Sum of non-HEP usage is about 1/3 of the size of USATLAS.
- OSG is actively trying to engage new communities, but recognize the stakeholders are the priority.


OSG Wall Years by VO Classification from 2008-04-01 to 2009-04-01 (Total: 18,764)




# Who are the providers of Resources?

- Compare this to the previous slide.
- Primary beneficiaries of opportunistic resources are FNAL Run II and non-HEP users.

Wall Years by Resource Owner from  
2008-04-01 to 2009-04-01



# Non-HEP Usage



Maximum: 1,403,146 Hours, Minimum: 0.00 Hours, Average: 785,479 Hours, Current: 449,515 Hours

Lots of these organizations don't have steady production like the HEP VOs -- usage is seasonal according to research activities

# How big is the OSG?

- This is a favorite question - how big/powerful is the OSG - but fairly tricky to answer.
- Let's catalogue the ways we've failed to do this.

# Failing to measure the OSG

- Survey sites: We don't "own" the sites, so there was low participation; sites misunderstood the questions (i.e., # of cores vs # of CPUs) and filled in the wrong value.
- Query the batch systems: It's impossible/difficult to understand all the possible batch system configurations and know what cores are grid-accessible.
- Have site admins fill in information during configure time: Sometimes site admins are lazy and just put in fake numbers. When resources are added, they don't always update the configuration.

# Stories of Failure

- Nebraska under-reported their cores in the site survey by 50% due to confusion in the survey questions.
- SLAC is a huge, huge computation farm (~9k cores) that only allows about 300 grid jobs to run at a time.
- UFlorida accidentally wrote a few too many zeros in their configuration and advertised 320,000 cores (whoops!).

# New Approach

- “Proven Capacity”: We only want to count a size as having a certain size if we can use that many cores.
- Site size is equal to the maximum observed utilization in Gratia accounting.
- Probably a lower bound - we implicitly assume that the site has been fully utilized at some point (but if we can't fully use it, we probably shouldn't include those cores in the OSG anyway!!).

# Proven Capacity


- So, if a site reports 168k wall hours in a week, then the proven capacity will be at least  $168k / 7 \text{ days} / 24 \text{ hrs} = 1,000$  cores.
- Pros:
  - No human input = no human error.
  - All sites report accounting numbers - very reliable.
  - Accounting numbers are examined daily for errors.



# Normalization

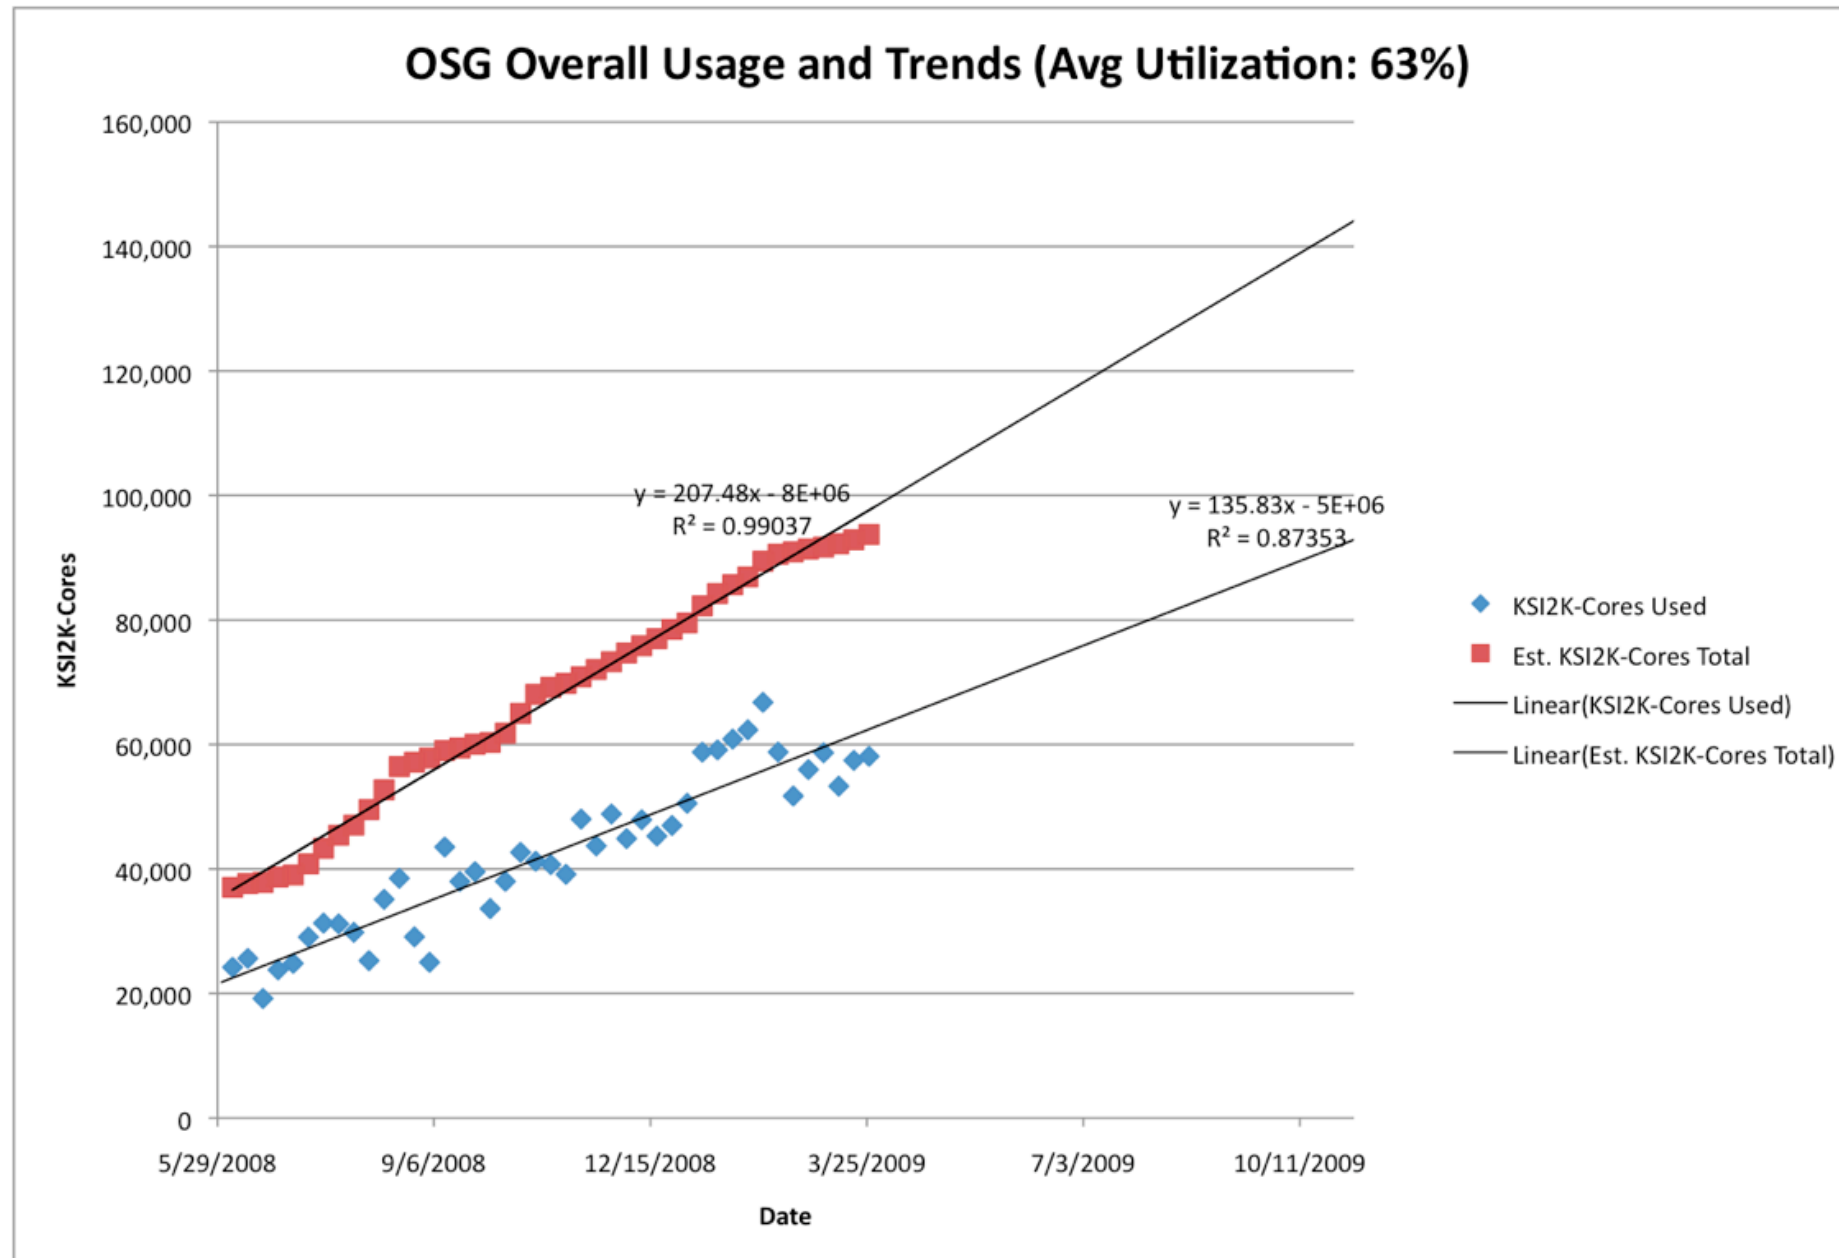
- We still do collect information about the CPU models at each site - not to determine size, but to guess average KSI2K score per CPU.
- Currently, average normalization constant is 2.3 KSI2K
- Remaining graphs don't show cores, but KSI2K-cores.

# Proven Capacity



Top number is all cores in all clusters - regardless if we've ever accessed them (or even could!)

# How will the OSG Grow?



Simple linear fit to data.

# Growth

- The proven capacity grows by about 207 KSI2K-cores / day; usage grows by 135 KSI2K-cores.
- At this rate, we'll have accomplished over 655 million KSI2K-hours during Year 3.
- This is assuming the current trend continues -- beware of predicting the future!

# Predicting Usage

- We know there's one major event - *LHC turn-on* - that will probably disrupt our usage patterns.
- In terms of growth, will this help or hinder?  
We don't know.
- Certainly, if the USLHC VOs utilize their owned resource, they will want to start using opportunistic resources too!
- Increased competition? Will this push out other VOs? These questions are unanswered.

# Conclusions

- OSG continues to grow as a grid!
- Main drivers are HEP, non-HEP usage is presently small but expected to grow.
- Measuring the size of a grid is non-trivial.
- Our favorite metrics are ones that avoid human input/error.
- If trends continue, OSG will accumulate around 1 billion KSI2K-hours by around October 2009.
- But we know our primary growth drivers (HEP community) are having a huge event in October 2009, so our crystal ball is a bit murky.