What's new in FIFE

FIFE WORKSHOP
JUNE 16, 2014
MICHAEL KIRBY

Outline

- Changes to FIFE support
- Updates to documentation
- Cloud utilization
- Changes to FermiGrid
  - Improvements to job workflows and resource utilization
- New services

FIFE Support Group

- new group has been formed within service now
- can be used as for SNOW tickets where uncertain which service
- triage the issues and direct them to appropriate experts
- Act as Open Science Grid stakeholder for FIFE experiments
  - Interface with OSG team
  - Gather requirements and submit request to GOC
  - Track those request and deploy new services/features

FIFE Support Group

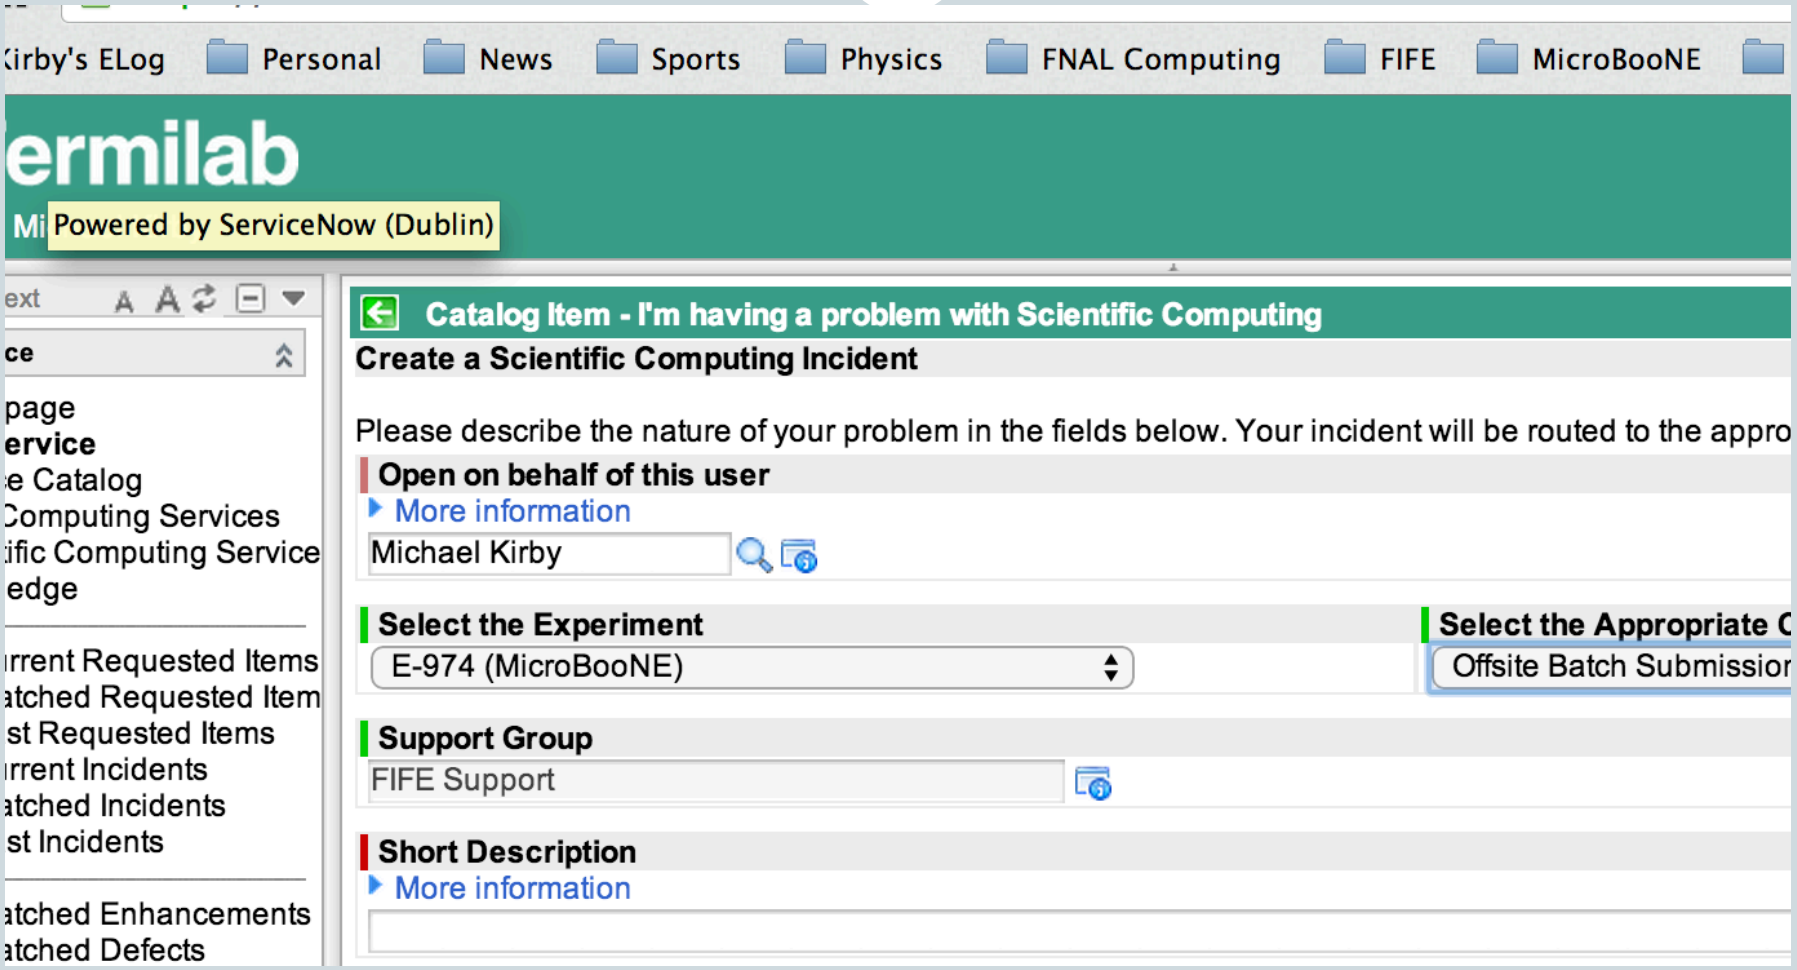

FIFE Documentation

- Katherine Lato written extensive FIFE intro/docs
- https://cdcvs.fnal.gov/redmine/projects/fife/wiki
- Covers introdcution as well as detailed links to services and projects within FIFE
- Also transitioning away from IFRONT page
  - Please let us know if you have old links to IFRONT so we can help you update those links
- UPS Documentation has been updated
  - http://www.fnal.gov/docs/products/ups/ReferenceManual/html/upsv4toc.html
  - https://cdcvs.fnal.gov/redmine/projects/ups/wiki/Documentation
  - https://cdcvs.fnal.gov/redmine/projects/ups/wiki/Getting_Started_Using_UPS
- More projects on the way (SAM, jobsub, etc)
- Better documentation for integrating services into experiment operations

FermiCloud On-Demand Services

- Our Goals:
  - Enabling users to access cloud resources both on- and off-site
  - Dealing with unusual/temporary resource needs not usually my by normal grid computing
  - Providing burst computing capacity
  - Providing capacity to make data movement, compute, and web caching services scalable on demand

Steve Timm

GlideinWMS – Grid and Cloud Bursting

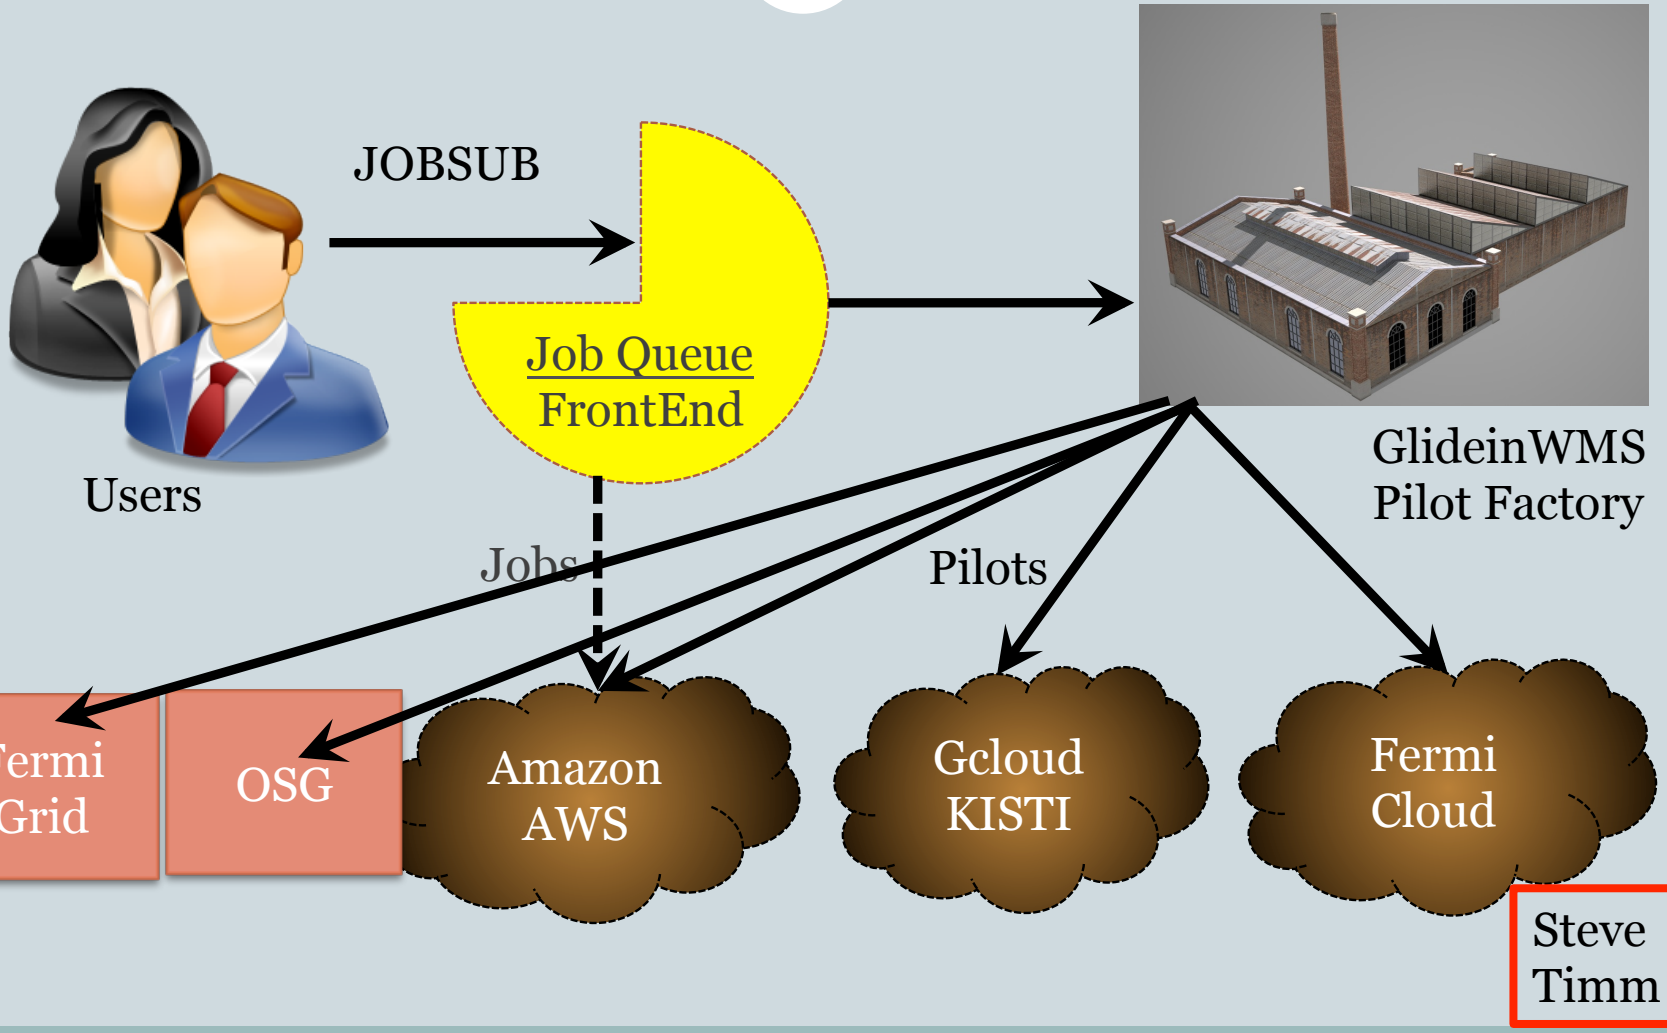

On-demand Grid-bursting – NOvA MC on AWS

Number of jobs 1088:

- Successes 1047
  - Relies on extensions to GlideinWMS
- Failures 41
  - File upload (ifdh) 1
  - File download (SAM) 1
  - nova executable 33
  - Art non-0 exit code 1
  - Geant4 failure 4
  - Hang job (100% cpu): killed

Total hours: 1135 (56 min/job)

- overhead: 6 min/job

Output data: 326 (incl. 4GB tests)

- data: 152GB (BA via BeStMan)
- log files: 170 GB (on fifebatch1)

Input data: 72 GB (free)

Cosmics Background in the Near Detector

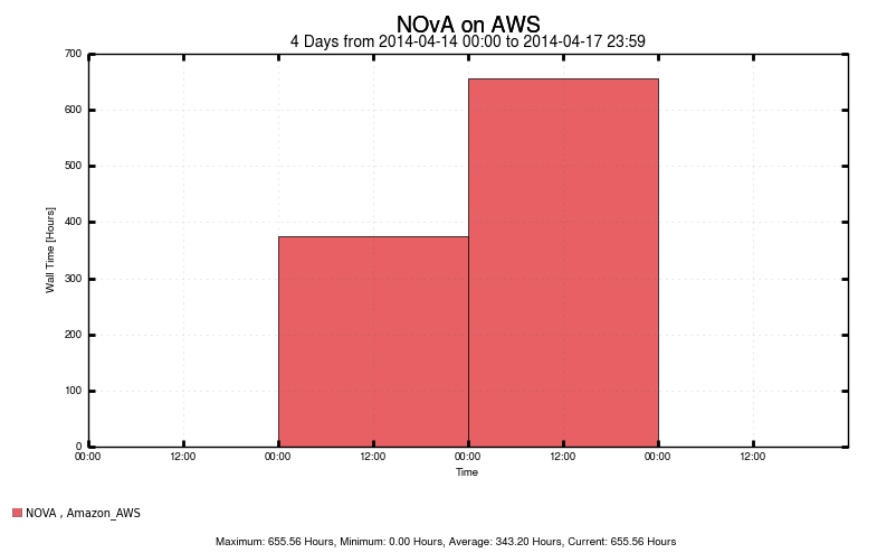

m1.medium spotpricing bid $0.07 →
got "blended" price $0.066 , instead: need to investigate

Total cost: $125

Data transfer: $39

CPU: $86

Steve Timm

Jobsub High Availability

- What is it? Suite of tools to manage batch/grid submission
  - Fault Tolerant: Minimal dependency between components
  - Scalable: High Availability (HA)
- Advantages?
  - Define common interfaces for experiments
  - Simplify the job submission process by:
    - Integrating complex grid tools in a sensible manner (credential management)
    - Automates mundane tasks interfacing with storage
      - Start SAM project
      - Facilitates data movement using ifdh to protect shared resources from overload
      - Stop SAM project
  - Jobsub along with GlideinWMS shields the user from complexity of running complex workflows on Grids & Clouds

Gerard Bernabeu

Partitionable Slots

- What is it?
  - One ‘partitionable’ slot per WorkerNode (server)
  - From which ‘dynamic’ slots are made
  - When dynamic slot exit, merged back into ‘partitionable’
  - Split happens at claim time (JDL/classadd)
- What does it buy us?
  - Partionable slots split on CPU, Disk, Memory. You need all to get a matching slot => will have to monitor all resources
  - When you’re out of one you’re out of slots

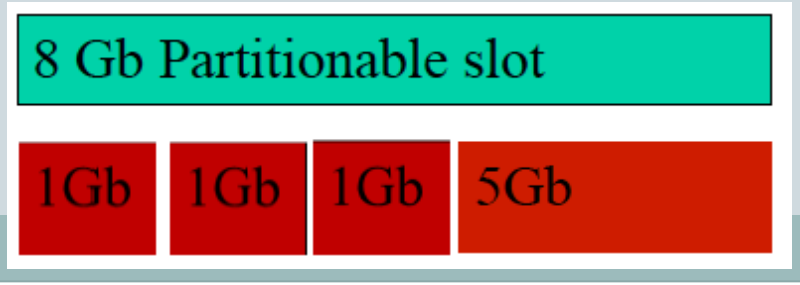

Gerard Bernabeu

Hierarchical quotas

- What is it?
  - Hard quotas
  - in a hierarchical way

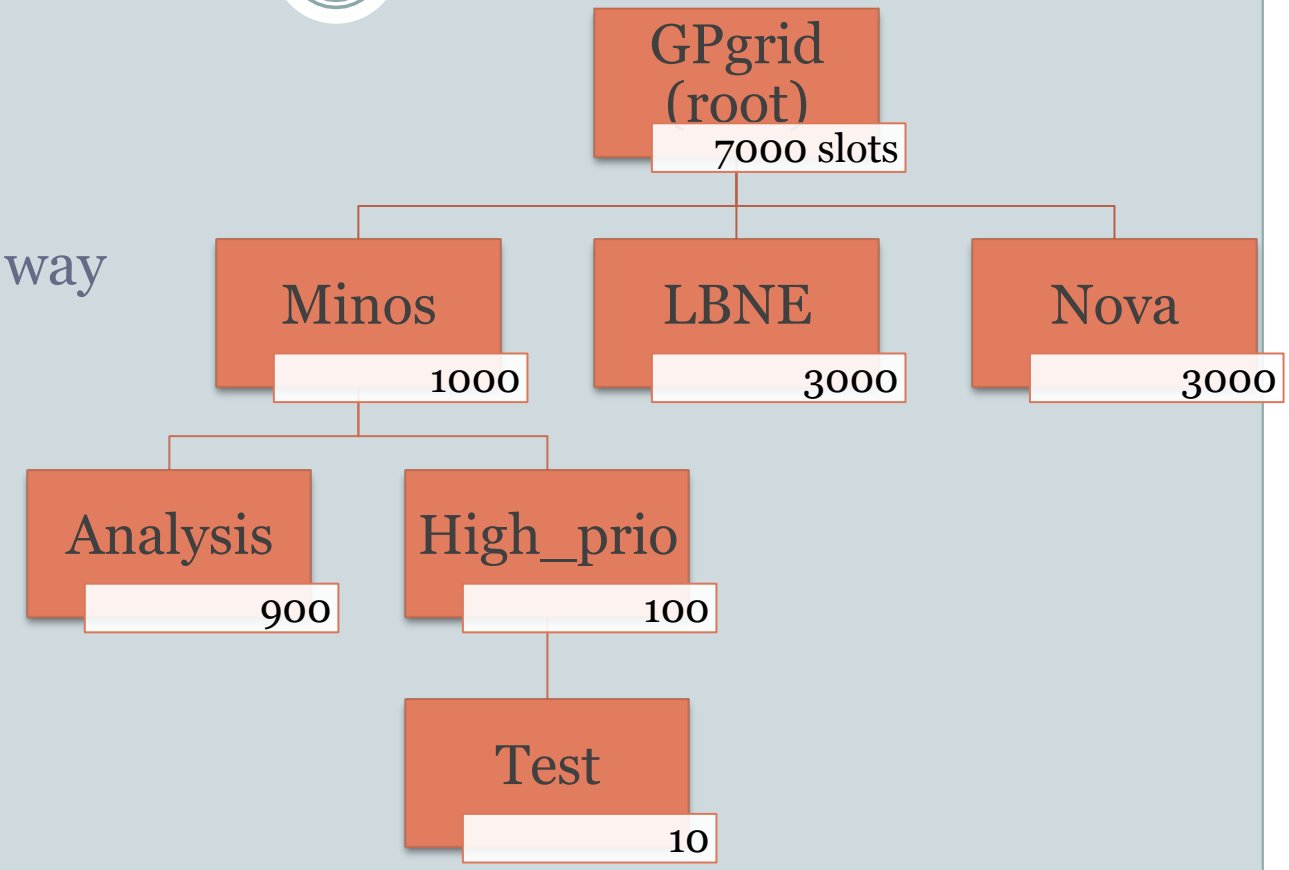

- Advantages?
  - Different 'queues' within the experiment

Gerard Bernabeu

Hierarchical quotas with surplus

- What is it?
  - Soft quotas (can go over if available resources)
  - Surplus follows the hierarchy (can be limited by a higher tier)
  - We need to agree on a maximum running time (in real hours, not cpu time) Eg: jobs must be shorter than 48h
    - Can have special 'queues' (AccountingGroups) with larger max time, but no surplus enabled for those
    - Time length is an open questions for experiments
- Advantages?
  - Allows Fermilab-based jobs to more easily grow to take up the whole GP Grid cluster if there is the demand.
    - Simplifies current jobsub structure; users will no longer have to pick opportunistic or non-opportunistic
    - Removes preemption for local resources (Gpgrid)
  - Tends to ensure desired quotas (if the farm is full, after the max running time)
  - Allows implementation of high priority accounting groups (ie queues)

Hierarchical quotas with surplus

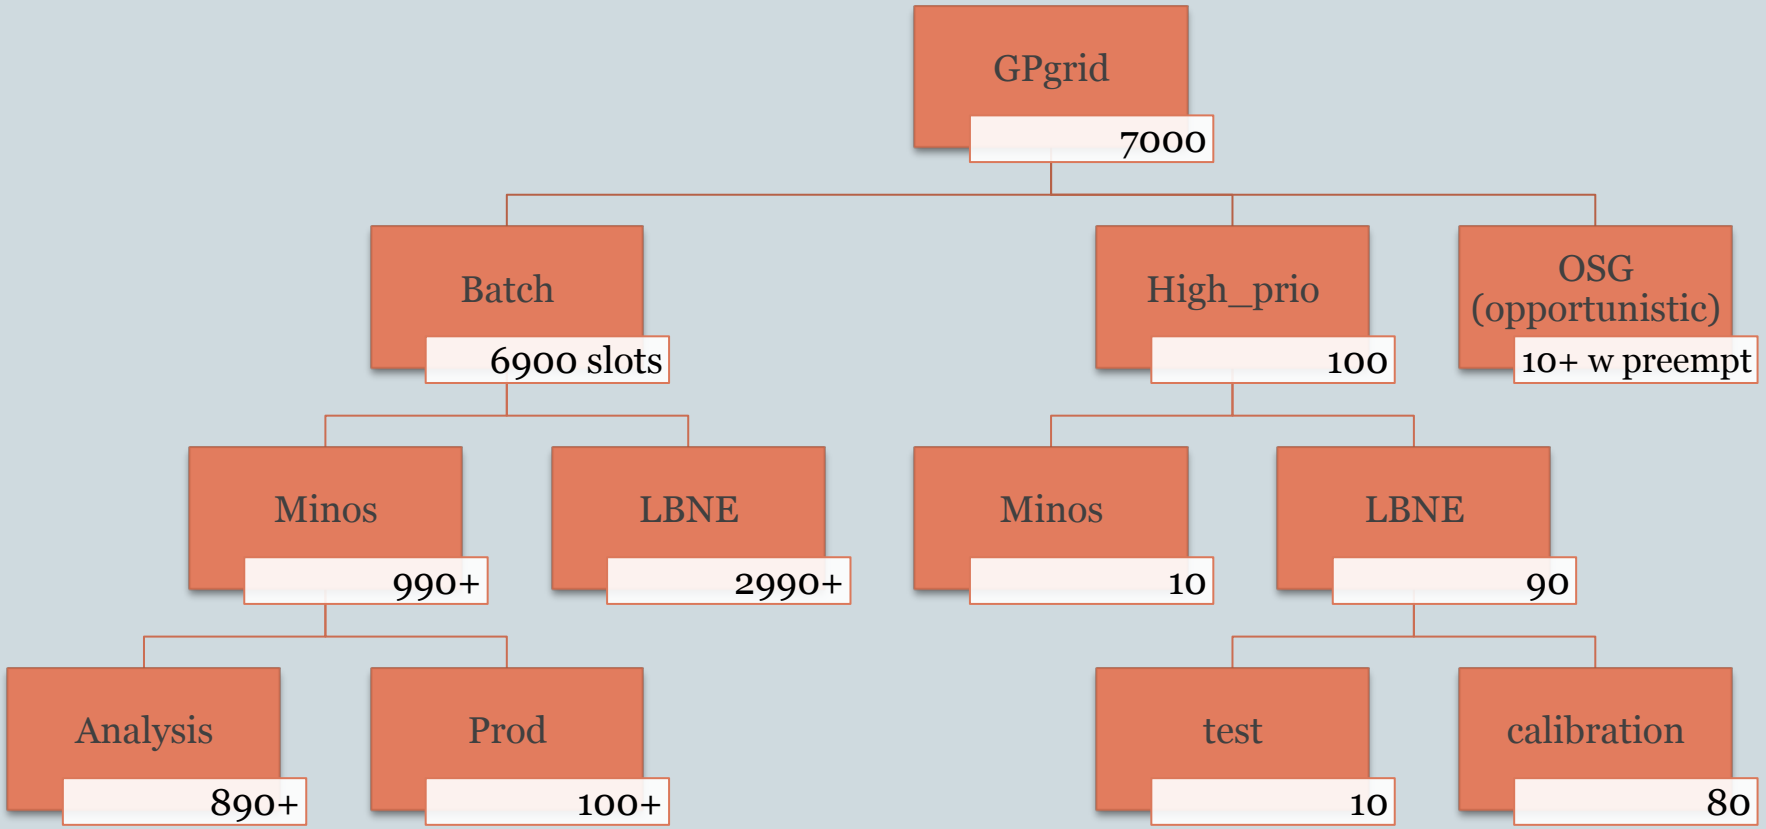

Gerard Bernabeu

GPgrid upgrade to SL6

- GPgrid WorkerNodes are being reinstalled in a phased way.
- Process will be completed by Fall 2014, meaning that there will be no more SL5 slots available in GPgrid.
- In order to obtain a reduced amount of SL5 WorkerNode slots jobs can be directed to FermiCloud SL5 on-demand WorkerNode (without NFS mount points?).

End of month	SL5 slots	SL6 slots
July'14	3232	4456
September'14	0	7688

Gerard Bernabeu

Other new services

- Continuous Integration
- Updates to CVMFS
- Potential larger memory allocation on OSG
- Data Handling
- Jobsub client
- Many new projects to help you get your science done

Hopefully not this kind of computing…

Thanks to everyone for their contributions and participating in the workshop!

Thanks Katherine for the image!