ALICE USA Computing Project for developing AliEn-PanDA interface

Ideas for a Plan

PanDA Workshop

Why are we…. ?

- Interested in PanDA?
  - AliEn originally planned as a light weight interface to leverage other middleware
  - Prototype common pilot job framework: CMS, ATLAS & CERN IT
  - ALICE Team is small wrt supporting AliEn

- Looking to ALICE-USA to initiate AliEn-PanDA investigation?
  - US Project on LCFs → an opportunistic motivator
  - Experience in AliEn development → extensions for US operations
  - Available disk storage capacity in the US
  - Interest in other US-HPC resources
    - As well as evolution of current ALICE-USA resources

ALICE-USA Computing Project

- Goal: high-performance, cost-efficient ALICE computing facility
  - Fulfills MoU-based ALICE USA obligations for computing & storage resources to ALICE
  - ALICE USA participation is about 7-8% of ALICE

- 2009 Project Proposal to DOE
  - Operate facilities at two DOE labs
    - NERSC/PDSF at LBNL
    - Livermore Computing (LC) at LLNL
  - In operation since Summer 2010

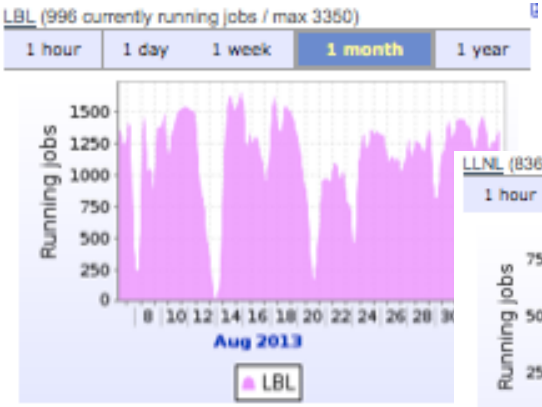

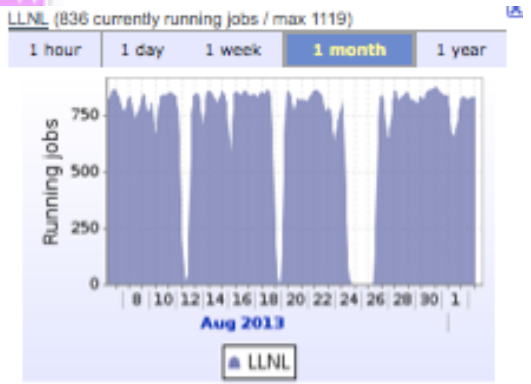

- Project personnel on steering/operations committee
  - Jeff Porter – project manager & ALICE Grid Manager for NERSC
  - Ron Soltz - Former ALICE-USA Computing Coordinator & LLNL ALICE Rep.
  - Jeff Cunningham – LLNL System Admin and ALICE Grid Manager for LC-glcc
  - Iwona Sakrejda – PDSF project lead
  - Bjorn Nilsen – ALICE-USA contributor

Facility Snapshot: LLNL/LC

- Livermore Computing
  - Large & diverse institutional-based High Performance Computing Center
  - Supports Lab Science and Engineering activities
- Cost effective procurement and operations model
  - Able to buy into routine very large purchases of scalable units
  - In-house managed OS (CHAOS) & other software (e.g. SLURM)

- ALICE Deployment model @ LLNL/LC
  - Separate single-use facility
    - 100% ALICE
    - Grid only use → no user logins
  - Large HW purchases, refreshed every ~4 years

Facility Snapshot: NERSC

- NERSC: US Department of Energy (DOE) Office of Science Flagship High Performance Scientific Computing Center
  - Available to all DOE Office of Science sponsored research
- Computing for Scientific Research
  - Large HPC Systems (~200k cores)
  - Special Clusters: PDSF, GPU, Visualization,…
  - Large archival storage (HPSS)
  - Data Transfer, Gateway & OSG/Grid Services
- Broad user support & engagement services
  - Including Grid & Science Gateways

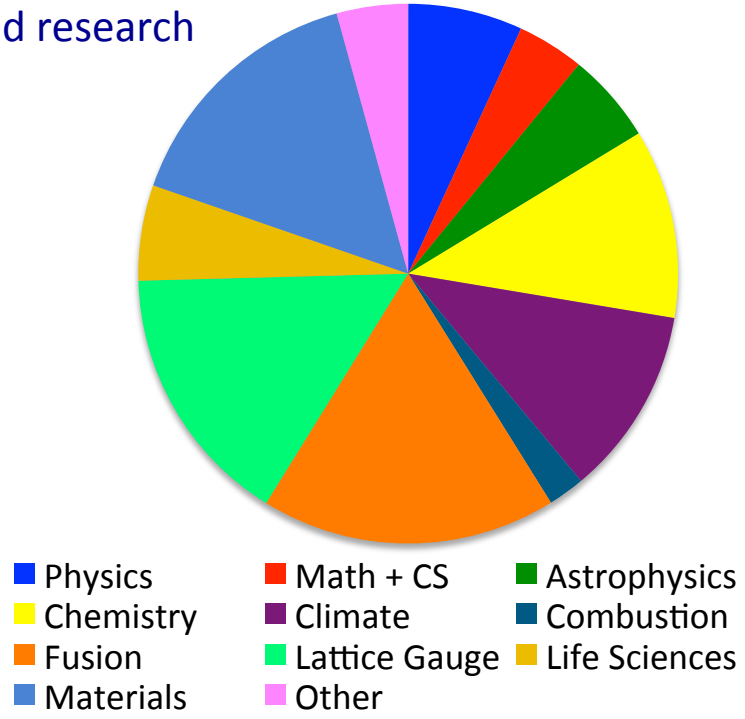

- ALICE Deployment Model @ NERSC
  - Project resources deployed on PDSF for ALICE Grid (see next slide)
  - Users can have login access with ALICE client tools available
  - Smaller Annual HW purchases to adjust to changing ALICE requirements

Facility Snapshot: PDSF

- Multi-group facility for Nuclear & High Energy Physics experiments
  - Allocations as "share" of resources
  - Fair share done in SGE (UGE)
- Share calculation includes
  - HW investment
  - FTE contribution
- Nuclear Science shares
  - ALICE 40%
  - STAR 30%
- Physics Div. shares
  - ATLAS T3 15%
  - Dayabay 10%

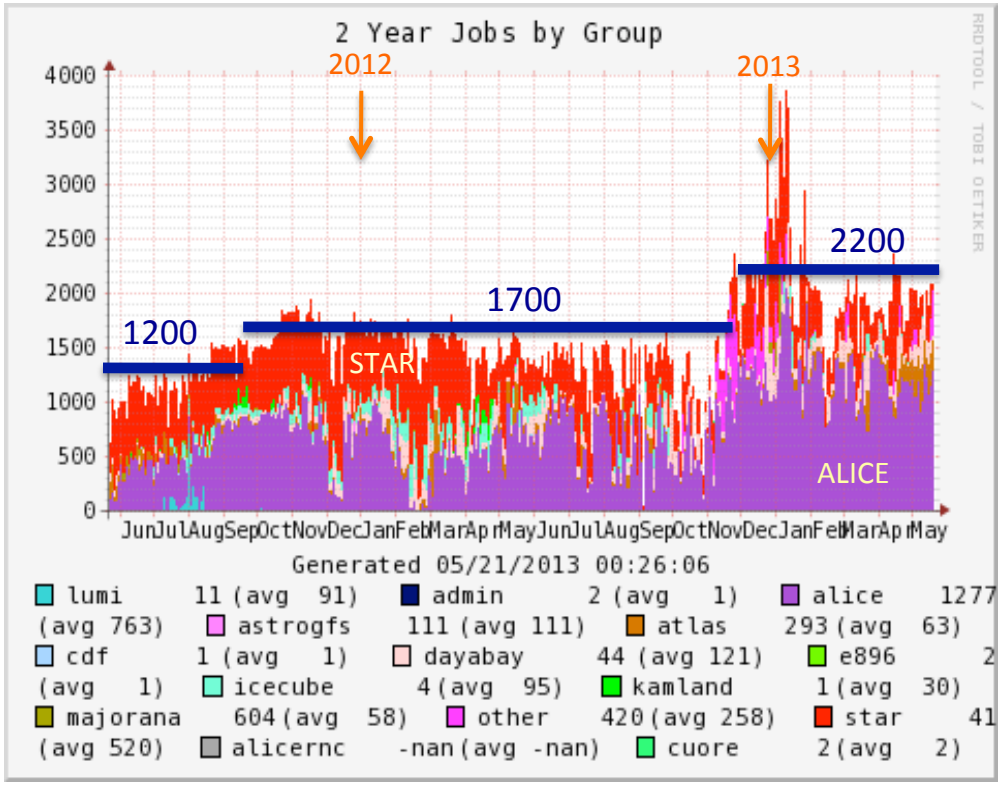

Running jobs

ALICE-USA & Open Science Grid

- ALICE-USA project leverages OSG capabilities
  - OSG Registration Authority
  - Resource reports sent to WLCG
    - Availability and Reliability Rep.
      - Critical services scans
    - Accounting Reports
      - Gratia site service
      - OSG central repository
  - Disposition of unused cpu @LLNL
    - OSG CE for DOE/NP funded groups

OSG Reports

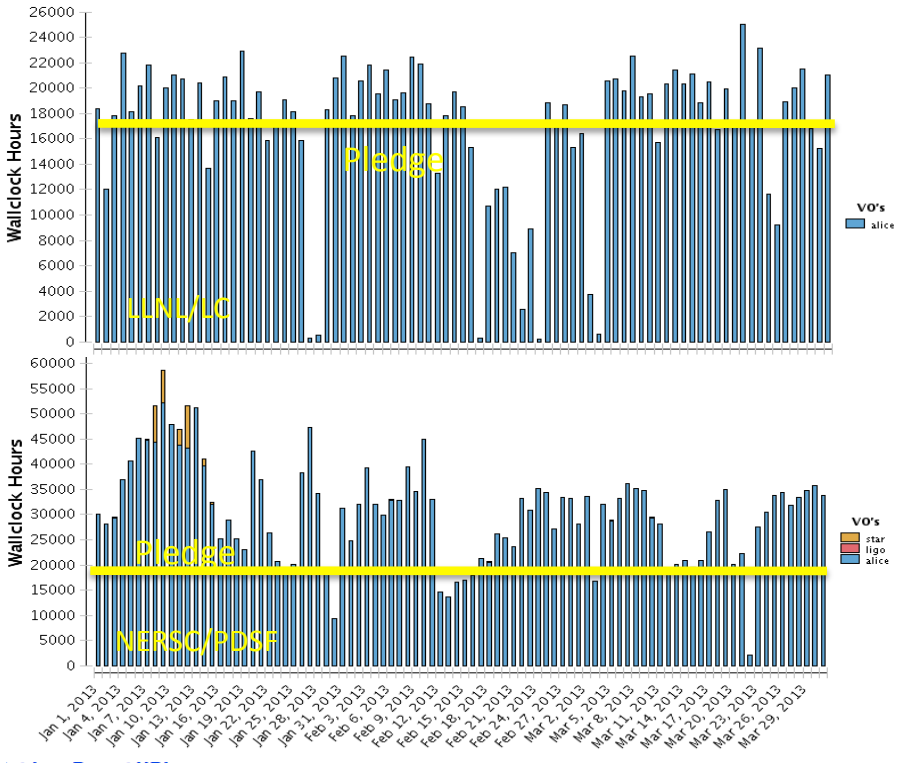

January 1 – March 31, 2013

AliEn & the ALICE Grid

- ALICE Grid Facility
  - In production since ~2004
  - more than 80 sites
  - Small, central operations team
    - + site admins around the world
- AliEn: 'Alice Environment'
  - Central Services at CERN
    - Task Queue, Job & File Catalog
    - LDAP registry of sites & configurations
    - Software management (build & deployment
  - VO box site-specific operations
    - JobAgent (pilot job) submission
    - ~~Software deployment~~
    - Resource monitoring
  - Grid Monitoring with MonALisa
  - Data Management
    - AliEn FileCatalog
    - Grid-Enabled XRootD SEs

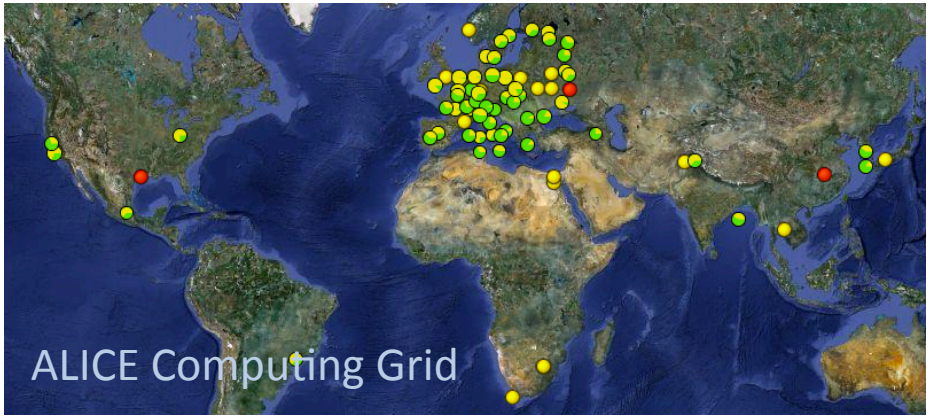

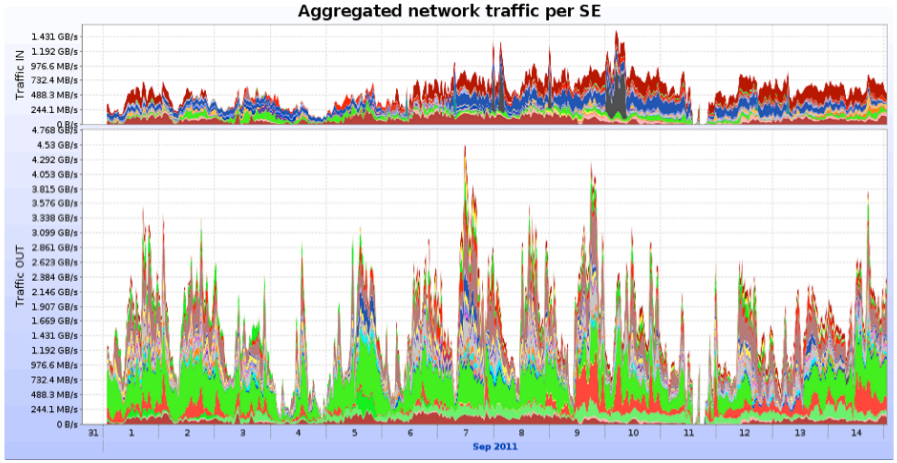

AliEn-OSG Interface

- Developed in 2009, collaboration of ALICE, OSG & NERSC
  - AliEn-CE submits to Condor-G (on VO box)
  - Condor-G submits to OSG-CE
  - VO Box operations otherwise standard

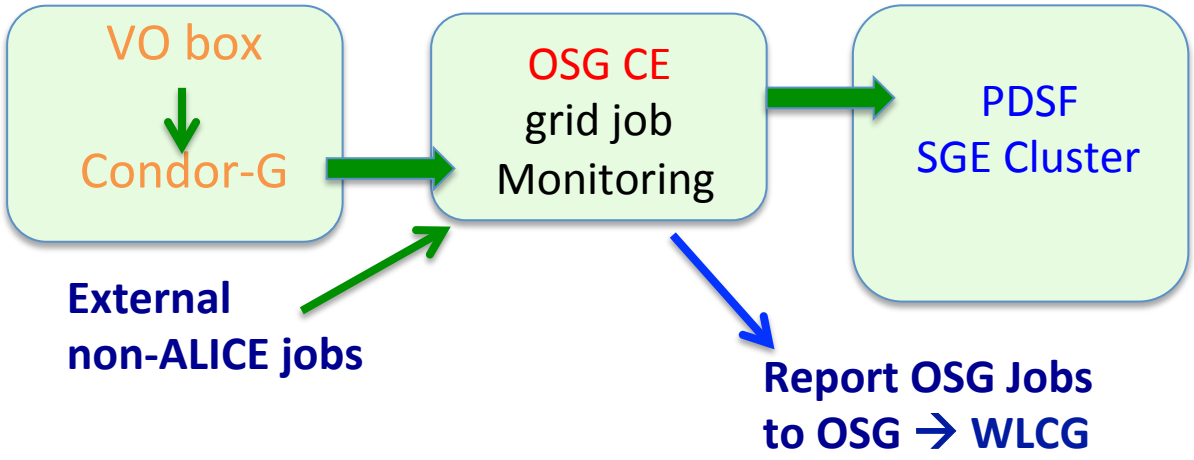

- Model extended in 2012 for opportunistic use of OSG resources

Opportunistic use of OSG Sites

- AliEn adaptor to access remote OSG sites
  - VO Box testbed at NERSC/PDSF
  - OSG-CE Endpoints
    - Managed as a dynamic list
    - Endpoint:MaxJobs:RSL
  - Test on selected sites
    - OSG-ITB sites + LBNL & LLNL
  - Automatic pairing to SE
    - Close::SE → 700 TB @ PDSF

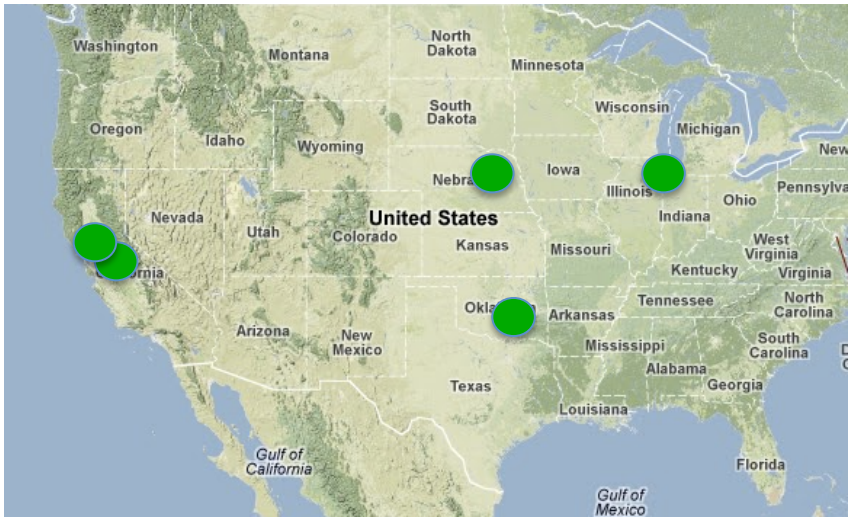

  - Point of the exercise was to exploit new AliTorrent software distribution

➢ Challenge became dynamic throttling of site utilization, without hooks on real-time resource capacity for opportunistic use.

Storage Element Deployments

- Grid Enabled XRootD SE
  - Site-level Redirector
    - LBL::SE
    - LLNL::SE

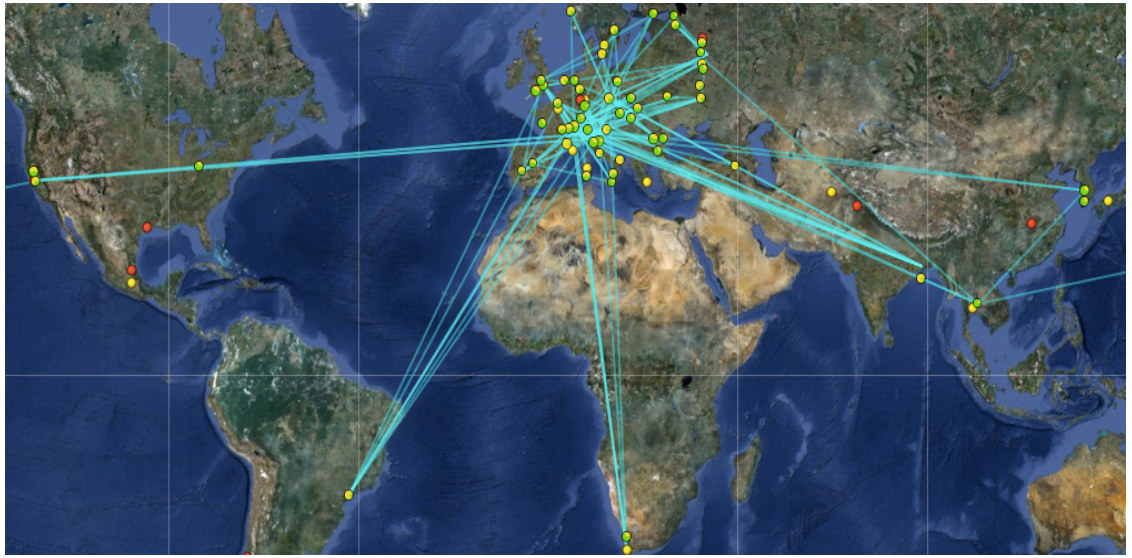

- Installed Capacity
  - LLNL/LC = 685 TB since Aug '10
  - PDSF = 720 TB since Oct '11

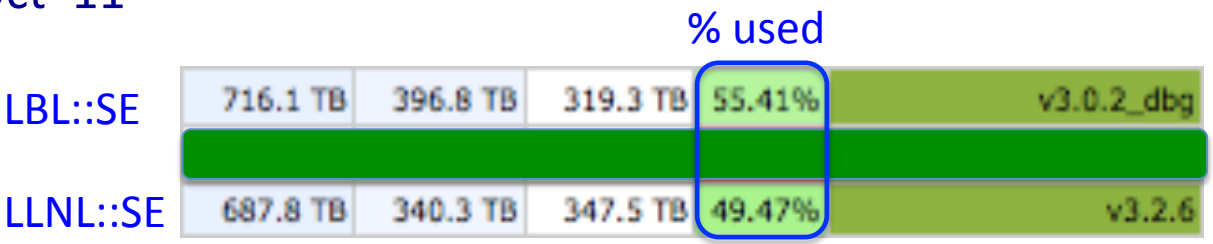

AliEn-PanDA Interface: How to Proceed?

- Work out a more detailed plan
  - Discussion at CERN in October …
- We will need access to resources / queue under PanDA server
  - Accessible via service certificate
- Modify AliEn-CE to submit to the PanDA server
  - Potential use of test VO box at NERSC
  - Retain AliEn workflow details
  - Target MC so no need to co-locate jobs with data
- Wrapper for PanDA to manage output data
  - Xrootd based mover with ALICE Authentication
  - ALICE data distribution model (Close::SE + replicas)

Potential Long-term Benefits to ALICE

- Simplify maintenance of thinner AliEn software
  - Replace (where practical) AliEn implementations with PanDA variants
- Leverage new features from common software effort
  - e.g. Event Service?
  - Reduced site management effort?
- Increased access to HPC resources at ALICE-USA sites
  - Backfill at NERSC: Hopper, Edison, …
    - Both ATLAS and ALICE have footprints at NERSC
  - LLNL/LC relatively new open campus extension
  - Both centers include large disk resources for ALICE
- Alter future HW procurements at ALICE-USA sites
  - Both centers can leverage costs for HPC-like resources if we can use them

Additional Resources

- NERSC Initiatives:
  - NISE: NERSC Initiative for Scientific Explorations
    - http://www.nersc.gov/users/accounts/allocations/nise/
    - 2012 Awards ranged from 100k – 13M hours
  - Data Intensive Pilot Project:
    - http://www.nersc.gov/users/accounts/allocations/data-intensive/
    - CPU paired with large &/or new storage technologies