



ALICE



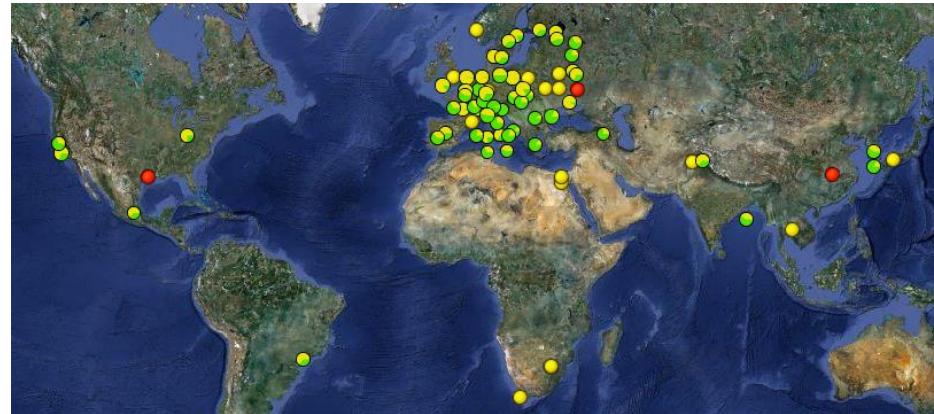
ALICE USA Computing Project Review

ALICE-USA Resource Review Meeting

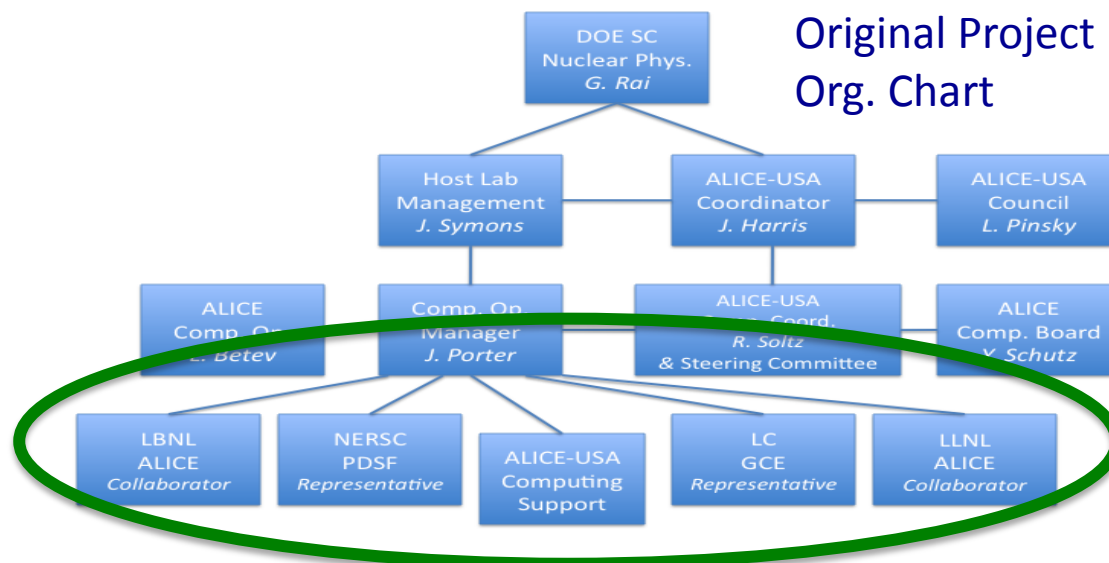
Apr 7-8, 2014

- ALICE-US Computing Project Overview
 - Project Operations
 - Resource planning
 - Facility snapshots
- Resource utilization & performance
 - History of project
 - Current year
 - Summary of project evolution
- Big Change on the way
 - LLNL Status & beyond

- **Goal: supply cost-effective Grid-enabled computing resources to ALICE**
 - Fulfill MoU-based ALICE USA obligations for computing & storage resources to ALICE
 - Based on ALICE USA participation at about 7-8% of ALICE
- **2009 Project Proposal**
 - Operate facilities at two DOE labs
 - NERSC/PDSF at LBNL
 - Livermore Computing (LC) at LLNL
 - 3-year procurement plan
 - LBNL as the host lab
 - Fully operational since Summer 2010
- **Operations defined in “Project Execution & Acquisition Plan”: PEAP**
 - Organization structure
 - Procurement strategy
 - Deliverables & Milestones



- Project personnel on steering/operations committee
 - Jeff Porter – project manager & ALICE Grid Manager for NERSC
 - Ron Soltz - Former Computing Coordinator & LLNL ALICE Rep.
 - Jeff Cunningham – LLNL System Admin and ALICE Grid Manager for LC-glcc
 - Iwona Sakrejda – PDSF project lead
 - Lisa Gerhardt – PDSF user support
 - Bjorn Nilsen – ALICE-USA contributor until 12/2013. Currently no replacement



- **Project Communications:**

- Local wiki: http://rnc.lbl.gov/Alice/wiki/index.php/ALICE-US_Computing
 - Document repository
 - Meeting Agenda & minutes
- Email list
- Monthly Meetings

- **Connection to ALICE Grid**

- Alice-grid-task-force email list
- Annual T1/T2 workshops
 - 2012 @ KIT Germany: Iwona & Jeff C
 - 2013 Lyon, Fr: Jeff C. & Jeff P.
 - 2014 Tsukuba, Jp: Jeff C. & Jeff P.
 - 2015 ?
- AliEn Developers Workshops
 - No longer held

ALICE-US Computing

Contents [hide]

- 1 ALICE-USA Computing
- 2 Documents
- 3 PDSF
- 4 LC

ALICE-USA Computing

- [Computer Steering Committee Meetings](#)

Documents

- [ComputingDocs](#)
- [\[Email Archives\]](#)

PDSF

- [PDSF home page is here](#)

LC

- [Livermore Computing](#)
- [Green Data Oasis](#)

- ALICE annual computing requirements
 - from ALICE Offline group
 - # events, event size, real/MC processing times & samples, data duplication ...

- Requirements vetted by WLCG
 - Final requirements in WLCG DB
 - 6 months before they take effect

- ALICE-USA Obligations:

- Fraction of total requirements defined by proportion of ALICE-USA/ALICE

Year	FY12	FY13
ALICE Requirements		
CPU (kHEPSPEC06)	336	290
Disk (PB)	22.0	30.3
ALICE-USA Participation		
ALICE Total Ph.D. (Total-CERN)	538	528
ALICE-USA Ph.D.	40	43
ALICE-USA/ALICE (%)	7.4	8.1
ALICE-USA Contributions		
CPU (kHEPSPEC06)	24.9	23.2
Disk (PB)	1.65	2.4

Table 1: Computing requirements from ALICE and planned ALICE-USA contribution.



Facility Snapshot: LLNL/LC



- **Livermore Computing**
 - Large & diverse institutional-based High Performance Computing Center
 - Supports Lab Science and Engineering activities
 - Lab interest in developing external collaborations
- **Cost effective procurement and operations model**
 - Able to buy into routine very very large purchases of scalable units
 - In-house managed OS (CHAOS) & other software (e.g. SLURM)
- **ALICE Deployment model @ LLNL/LC**
 - Separate single-use Grid facility
 - 100% ALICE
 - Grid only use → no user logins
 - Large HW purchase, refreshed every 4 years

ALICE-USA 2010 HW Purchases

13kHS06 @ \$25/HS06
680 TB @ \$210/TB

- 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 (100s k cores)
 - Special Clusters: PDSF, Visualization,...
 - Large archival storage (HPSS)
 - Data Transfer, Gateway & OSG/Grid Services
 - Evaluation Systems: GPU & Cloud Services
- Extensive user support services
- 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
 - Annual HW purchases to adjust to changing ALICE requirements

ALICE-USA Operations Activities


2009 OSG/AliEn Interface
2010 Fix OSG/AliEn scaling issues
2011 AliEn High Memory Killer
2012 OSG access with AliTorrent
2013 US MonaLisa Collector



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 project leverages OSG capabilities

- OSG Registration Authority
 - ALICE-USA user certificates
 - PDSF & LLNL machine certificates
- Resource reports sent to WLCG
 - ~~Availability and Reliability~~
 - ~~Critical services scans~~
 - Accounting Reports
 - Gratia site service
 - OSG central repository → WLCG

OSG Reports



Q1FY14 Report to DOE Oct 1 – Dec 31, 2013

➤ Funding agency monitors that we are not oversized for our mission

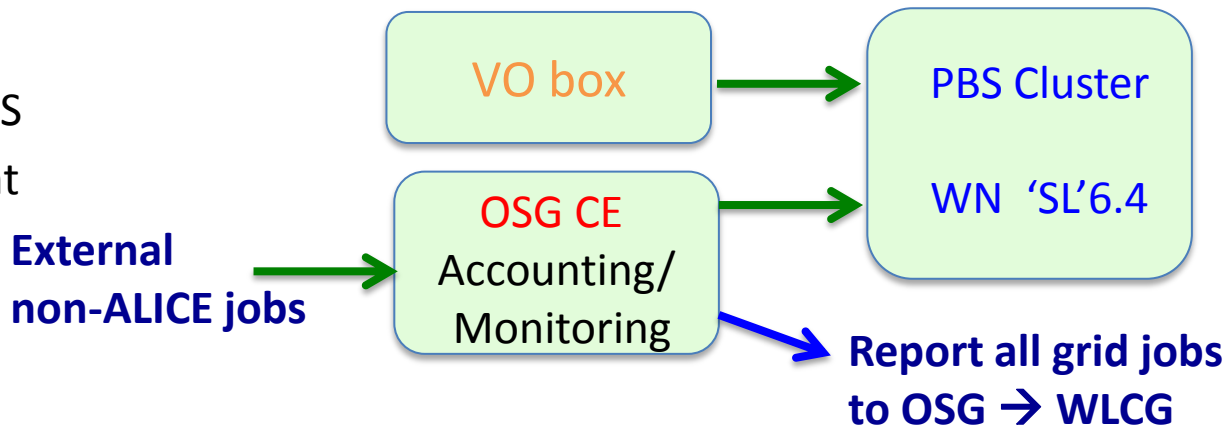


Site Configurations



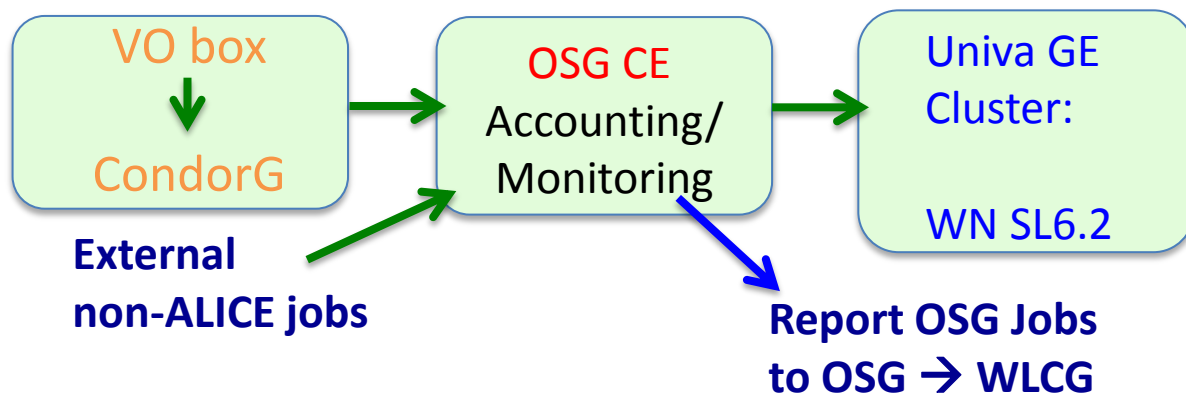
LLNL/LC

- Vobox Submits to PBS
- OSG-CE: independent external interface
- OSG Accounting



LBNL/NERSC/PDSF

- Submits to CondorG service on VO box
- CondorG submits to OSG-CE at PDSF
- OSG Accounting



➤ NERSC evaluating SLURM

- Target date ~ Sept 2014



Part II



- Annual Performance Summary

% used

NERSC/PDSF	716.1 TB	365.2 TB	351 TB	50.99%
LLNL/LC	687.8 TB	327.9 TB	359.9 TB	47.68%

Q1FY14 Report to DOE Dec 31, 2013

- Pledged Obligations

- LLNL/LC

- 650 TB since 2010

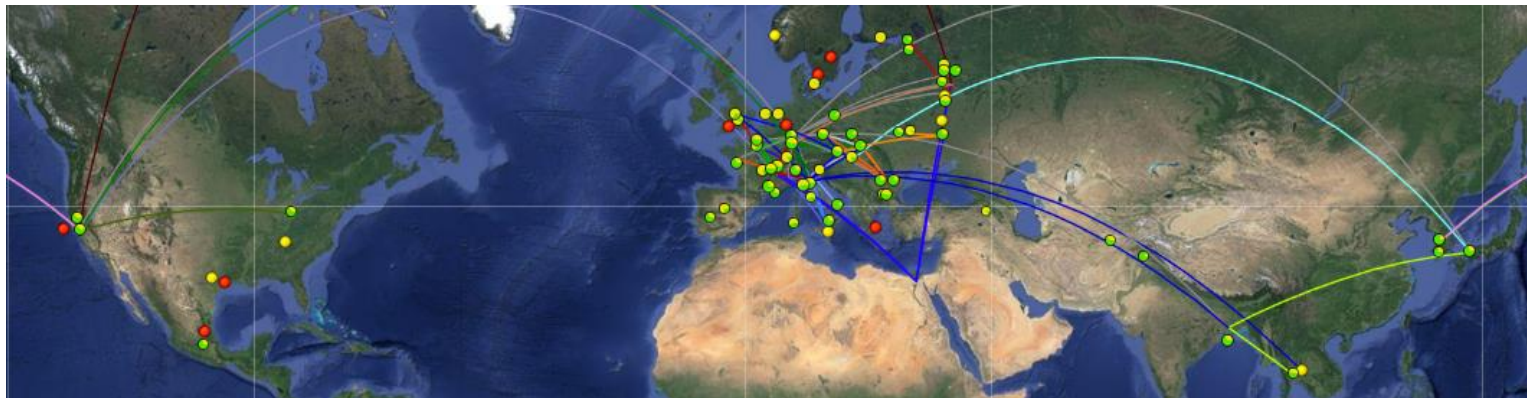
- NERSC/PDSF

- Steady ramp plan: 300TB → 740 TB → 1,020 TB → 1,200 TB → ...



- Installed Capacity

- LLNL/LC = 685 TB since Aug '10

- NERSC/PDSF = 720 TB since Oct '11



- High availability in AliEn SE tests
 - LLNL::SE → 96%
 - LBL::SE → 98%

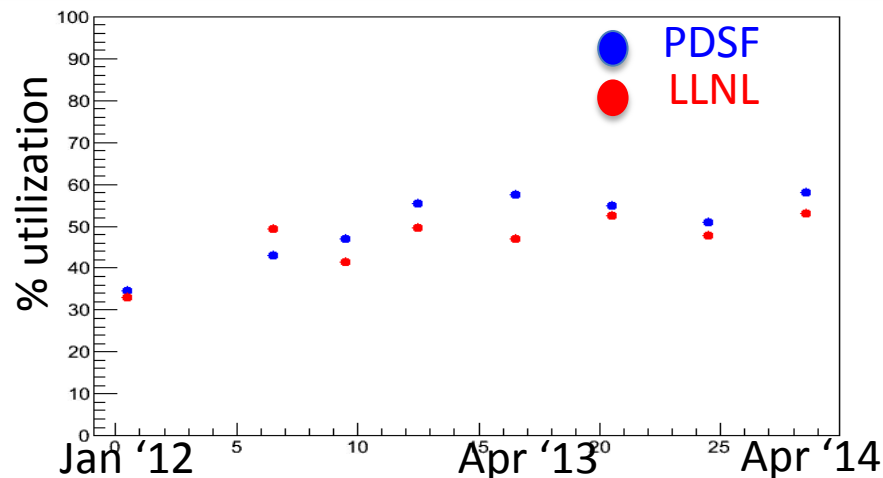


Statistics						
Link name	Data		Individual results of writing tests			Overall
	Starts	Ends	Successful	Failed	Success ratio	Availability
LBL::SE	31 Mar 2013 06:05	07 Apr 2014 04:19	4383	75	98.32%	98.47%
LLNL::SE	31 Mar 2013 06:04	07 Apr 2014 04:20	4268	190	95.74%	96.01%

June reconstructon @ KISTI

- Stable Utilization for 2+ years

- LLNL::SE → 35-50%
- LBL::SE → 35-55%



- “Nearby” SE effect

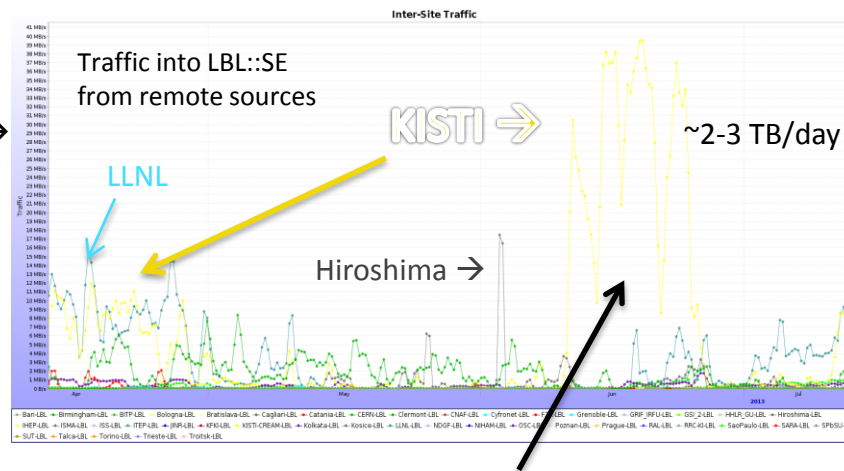
- Small nominal rate → LBL::SE

- Typically ~5MBps
- LLNL is largest writer, ~5-10 MBps

- Larger rates

- during reco @ KISTI
- Periodically from Hiroshima

30 MB/s →




June reconstruction @ KISTI



ALICE

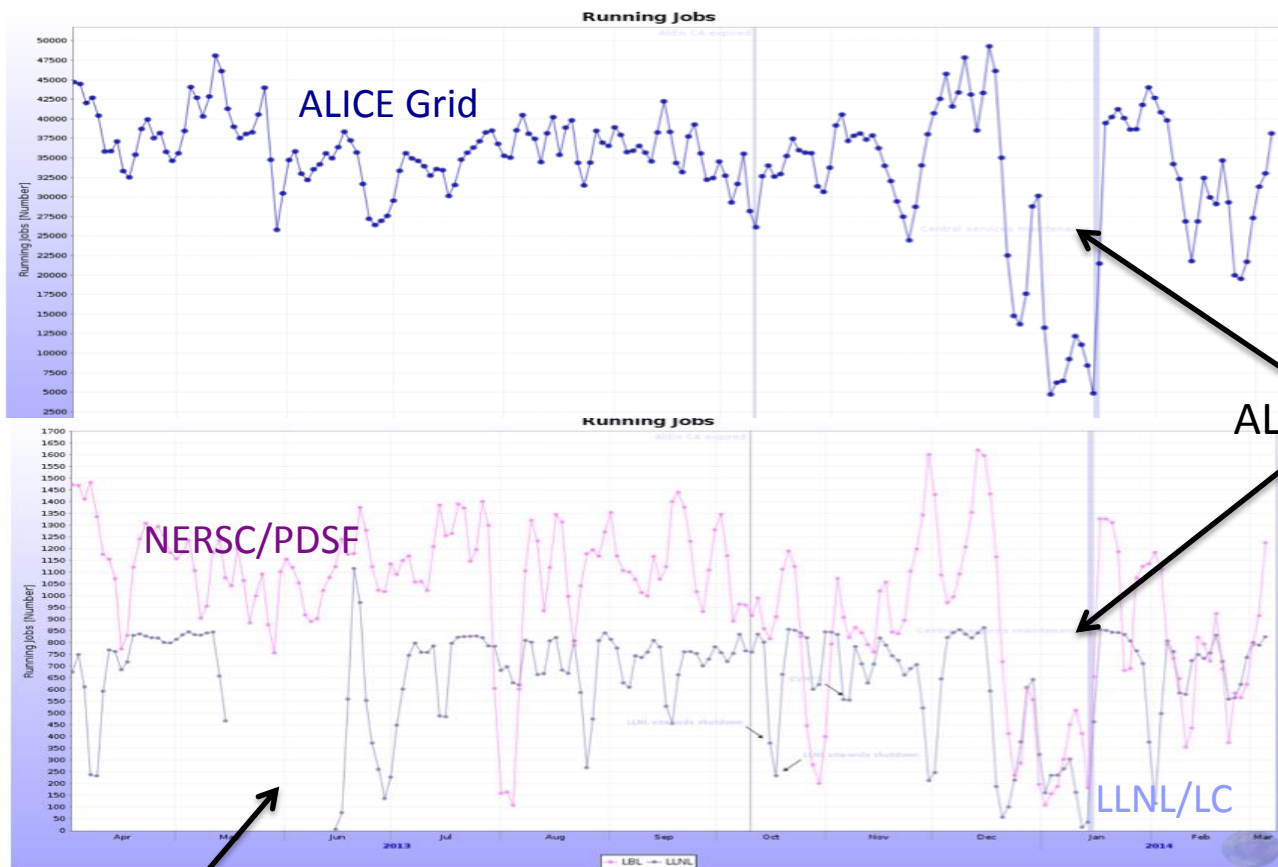
Project History: Job & core count





ALICE

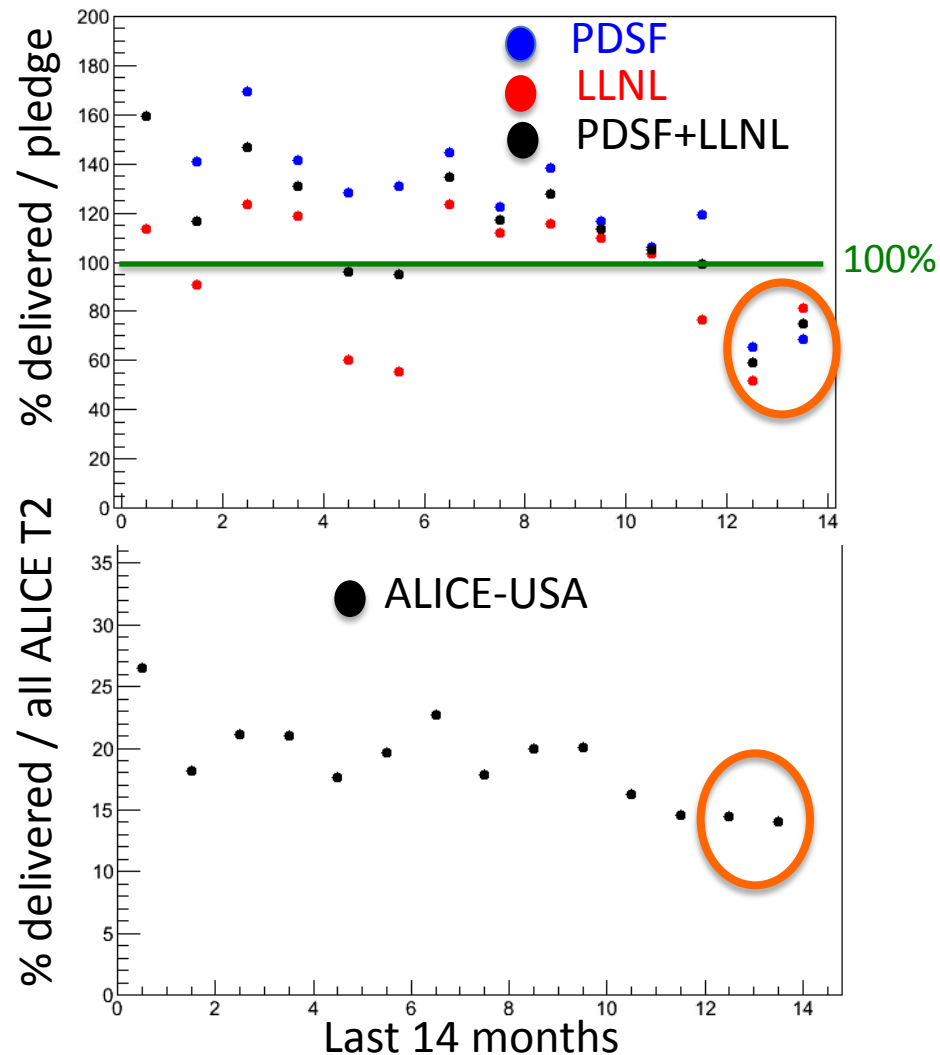
Job & core count: 2013 – Feb 2014



LLNL upgrade:

- Both Sites often over 100%
 - PDSF opportunistic cycles
 - Made up LLNL/LC upgrade loss
 - Illustrates value of 2 sites

- Ave >~ 100%
 - Except last two months (Jan+Feb)
 - Lower Overall ALICE Grid Use





CPU Utilization 2013 RRB Year 10 Months: 3/13 – 01/14

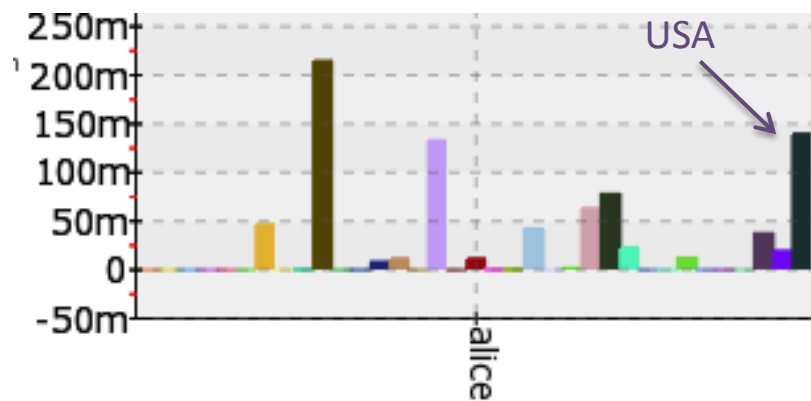


- CPU Utilization eff: CPU/wall
 - LLNL/LC & NERSC/PDSF ~ 65-70% efficiency
 - 70% Allowed by WLCG Accounting for T2



Efficiency factor for Tier-2 sites - utilisation 70% of pledge as specified in TDR

EGI Portal Country T2



- Utilization relative to pledges
 - LLNL/LC
 - Pledge : 11,500 HS x 24 x 300 x 0.7 (allowed eff.) = 58.0 MHS-hrs
 - Delivered: 51.6 MHS-hrs → **89%**
 - NERSC/PDSF
 - Pledge: 12,900 HS x 24 x 300 x 0.7 = 65.0 MHS-hrs
 - Delivered: = 79.9 MHS-hrs → **123%**
 - Combined
 - Pledge = 123.0 MHS-hrs , Delivered = 131.5 MHS-hrs → **107%**

- **Project Funded in Jan. 2010**
 - Both Tier 2 sites fully operational by Sept 2010
- **WLCG MoU**
 - LBNL signed, 4/2011
 - LLNL signed Lol, 2012
- **External Project Review, 2012**
 - Executive Summary: “As of today both participating sites have demonstrated their outstanding ability to reliably contribute to ALICE’s managed production and user analysis activities at excellent performance.”
- **Project plan updates based on new ALICE requirements**
 - FY13 update, July 2012
 - FY14 update, July 2013 ←included description for refresh all LLNL HW in 2014
- **External project review was scheduled for Feb 20-21, 2014.**



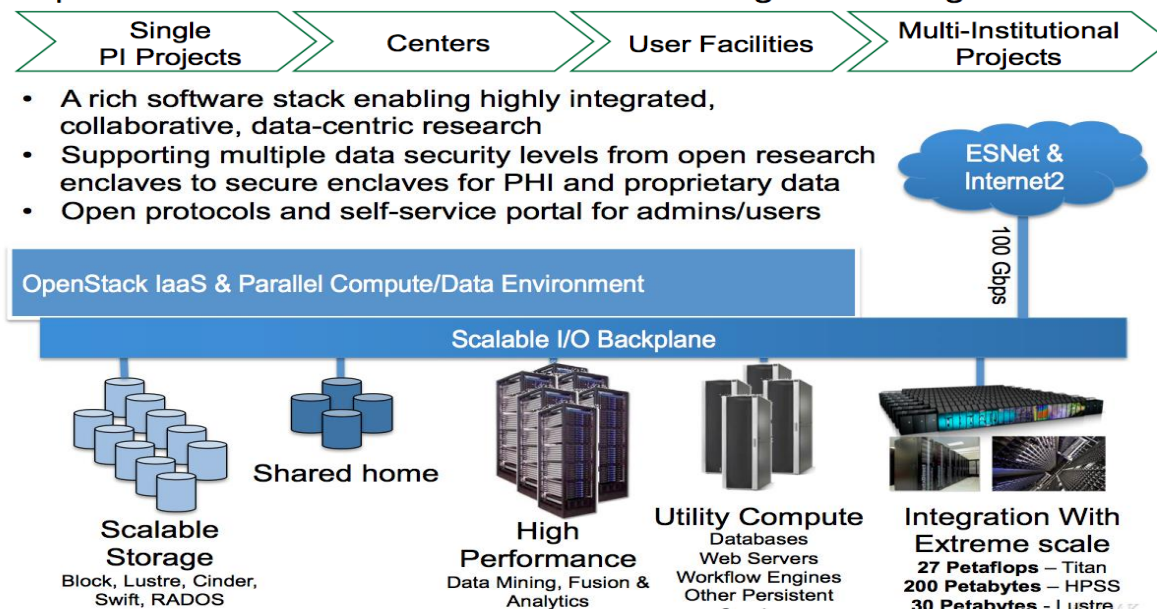
Part III



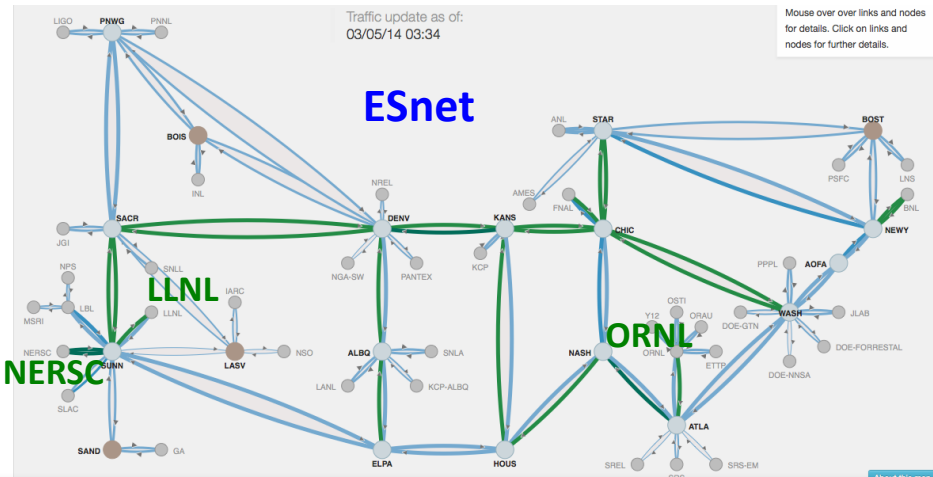
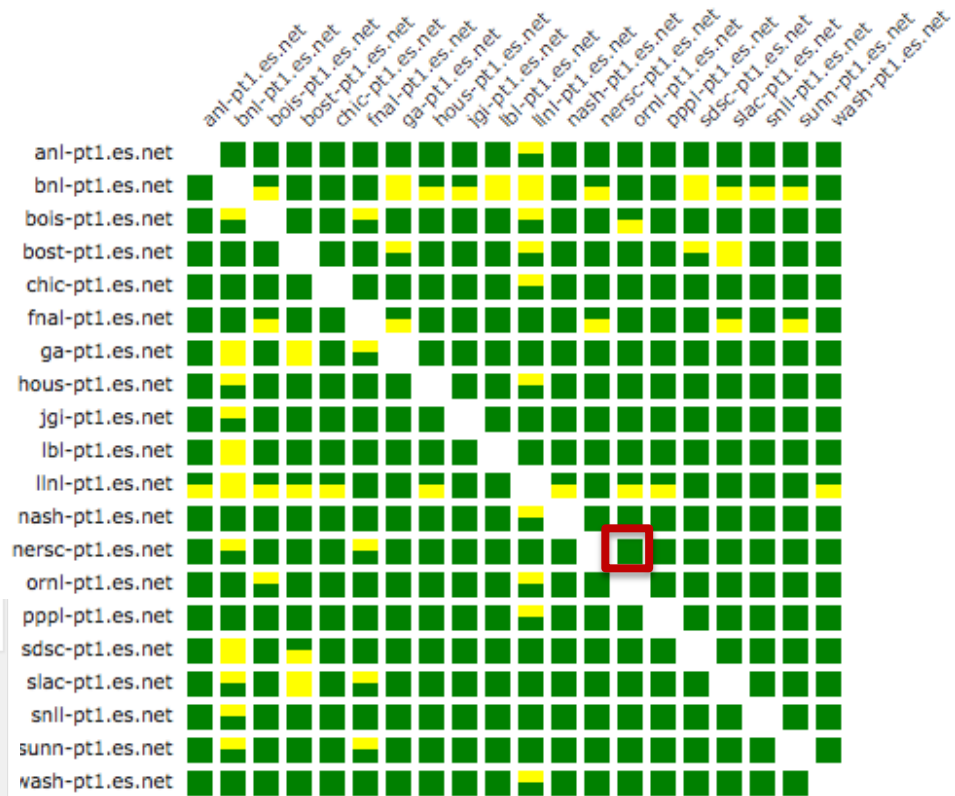
- Big Changes Coming

- Group at LLNL has decided to shift efforts from ALICE-USA
 - ALICE Tier-2 site @ LLNL to be decommissioned
 - Target dates → Oct 2014 (extending to Oct 2015)
- ALICE is losing an important partner:
 - Extremely cost-effective procurements
 - Experience built up over the past 4+ years.
- ALICE-USA Evaluation for replacement has recommended ORNL CADES

Compute and Data Environment for Science targets full range of needs



- NERSC/PDSF + ORNL/CADES
 - Scientific Computing strength
 - High-bandwidth connection: ESNet
 - Favorable cost structure
 - Proximity to HPC Resources
 - Oak Ridge Leadership Class Facility
 - NERSC Flagship facility
 - Strategic alignment with O² project

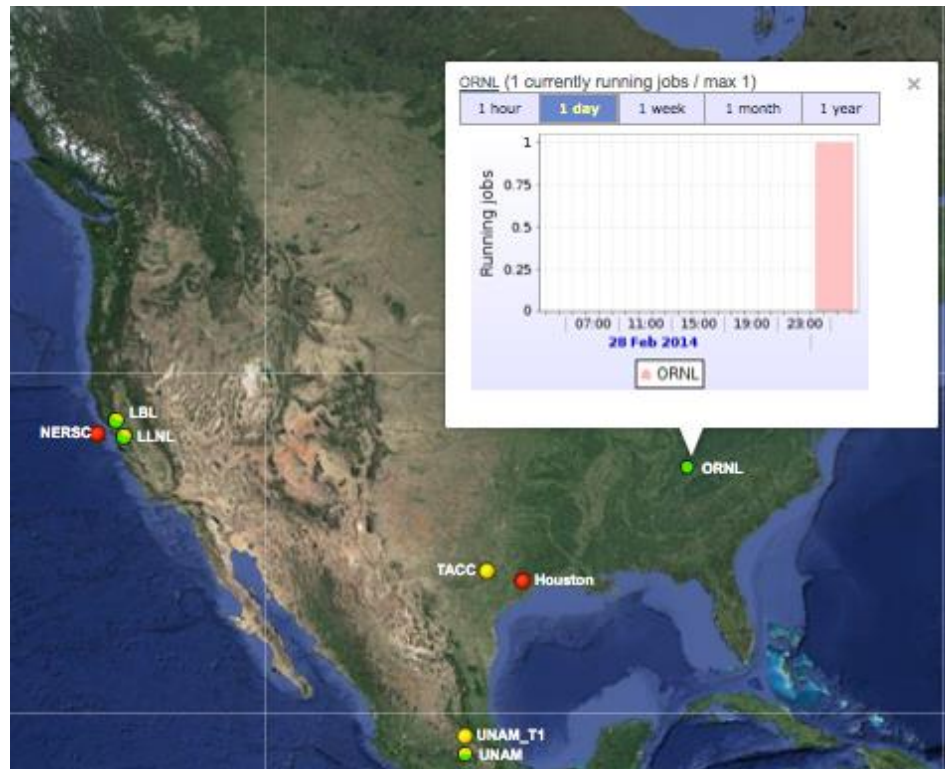


ESnet - ESnet Hub to Large DOE Site Border Throughput Testing

■ Throughput >= 2000Mbps
 ■ Throughput < 2000Mbps
 ■ Throughput <= 500Mbps

ESnet Monitoring

- Set up a demo ALICE grid site at CADES
 - VO box with AliEn services, local and CVMFS installs
 - Demonstrate network capability with AliEn monitoring
 - Demonstrate basic job processing



Excellent network
connection to US
& worldwide sites

- **ALICE-USA Computing project**
 - Two US based ALICE Tier-2 facilities, LLNL/LC & LBNL/NERSC
 - Currently satisfies all US pledged resources
 - ALICE used >100% of pledged CPU resources
 - 80% of pledged storage resources installed, due to lower utilization
- **The project will undergo a shift leading up to (during) Run 2**
 - Two US based ALICE Tier-2 facilities, ORNL/CADES & LBNL/NERSC
 - Replace LLNL/LC resources with new resources at ORNL/CADES
 - Rebuild expertise lost at LLNL
- **Challenge over the next 12-18 months**
 - Essential rebuild of our facilities with limited loss of service