

ALICE-USA Tier 2 Developments

ALICE Offline Week

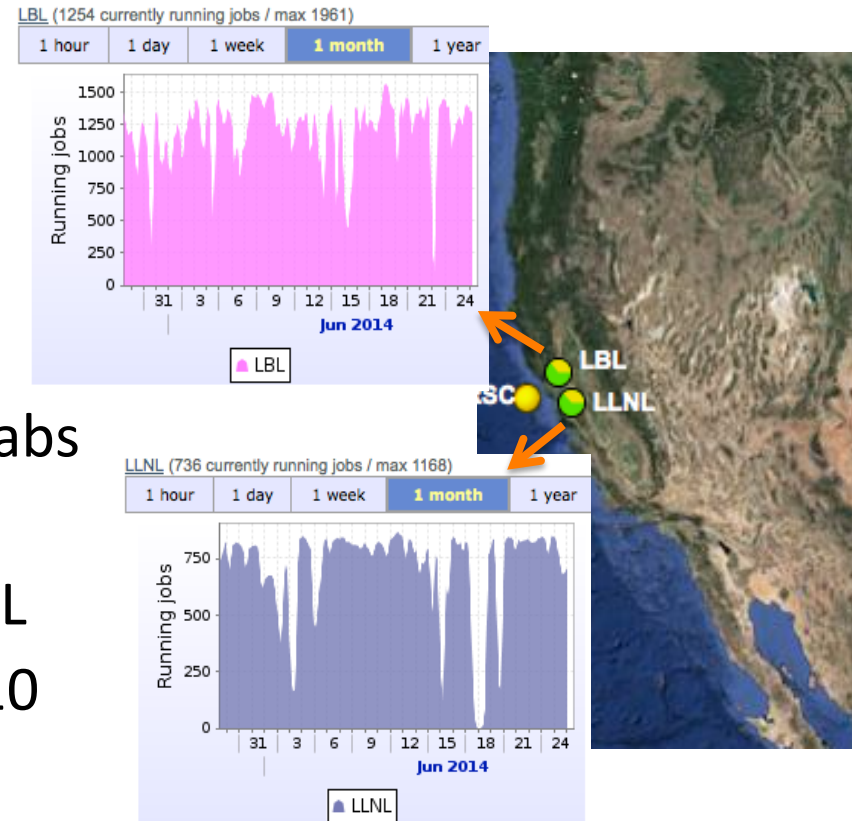
26 June 2014

ALICE-USA Computing Project

- Exists to fulfill MoU-based DOE-US computing obligations to ALICE

- Original Project Proposal
 - Operate facilities @ 2 US DOE Labs
 - NERSC Center at LBNL
 - Livermore Computing at LLNL
 - Approved & funded in early 2010

- Facilities operational since Aug 2010



Project Adjustment in 2015/2016

- LLNL will depart from ALICE for reasons of group size & logistics
 - ALICE Tier-2 site @ LLNL to be decommissioned after October 2015

- ALICE-USA Coordinator formed a small team to evaluate options & make recommendations for continuing the computing project
 - What is the level of computing resources (CPU, storage, bandwidth, expertise, etc.) that will be lost when LLNL ceases to be part of ALICE-USA?
 - What are the options for replacing this capacity?
 - Describe the potential sites and assess the pro and cons of each site, including issues like local expertise, cost effectiveness, migration paths to the O² time period, access, time required to be fully operational, etc.
 - Describe the pros and cons of 1, 2, or 3 site solutions.
 - If possible, provide a prioritization of the most realistic options.

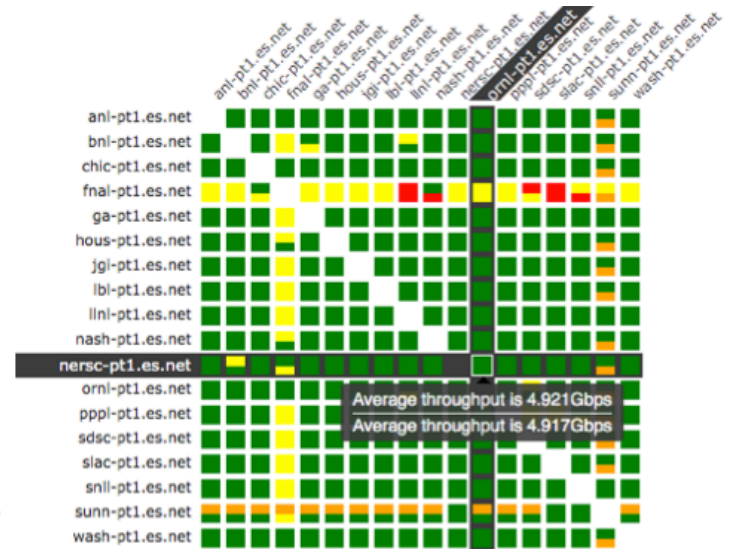
Review of Options

- Single site option at LBNL NERSC
 - Pro: Low administrative overhead
 - Cons:
 - Less flexibility for US operations
 - Strongly disfavored by ALICE
- 3+ sites ruled out as not cost effective
- Two site option (with LBNL NERSC)
 - Pros:
 - Can mitigate impact of outages at one of the T2s
 - Better geographical span in terms of data latency
 - Diversify procurement options
 - Cons:
 - Modest increase in admin overhead
- 2nd site selection criteria:
 - Strong institutional presence in ALICE-USA
 - interest in supporting ALICE-USA computing
 - Good network connection to rest of US and world
 - Established technical capabilities, embedded within larger facilities
 - Reasonable procurement and operational costs structure



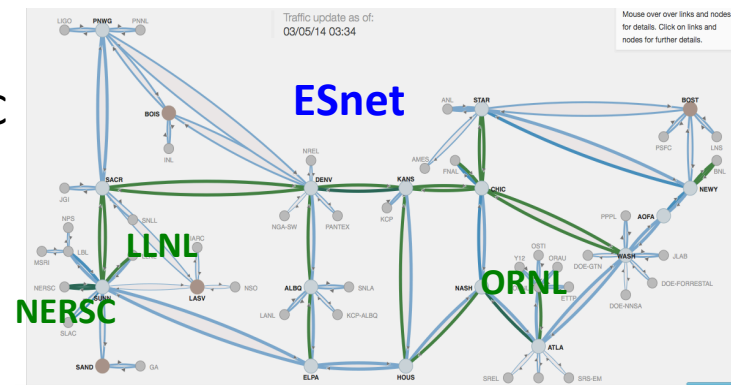
ALICE-USA Project Plan: New ALICE T2 facility at Oak Ridge National Laboratory

- ORNL CADES Facility:
 - Compute And Data Environment for Science
- LBNL NERSC & ORNL CADES
 - Both with Scientific Computing strength
 - High-bandwidth, monitored uplink to ESnet
 - Proximity to DOE HPC Resources with strategic alignment to O² project
 - Oak Ridge Leadership Computing Facility (OLCF)
 - NERSC DOE SC Flagship facility

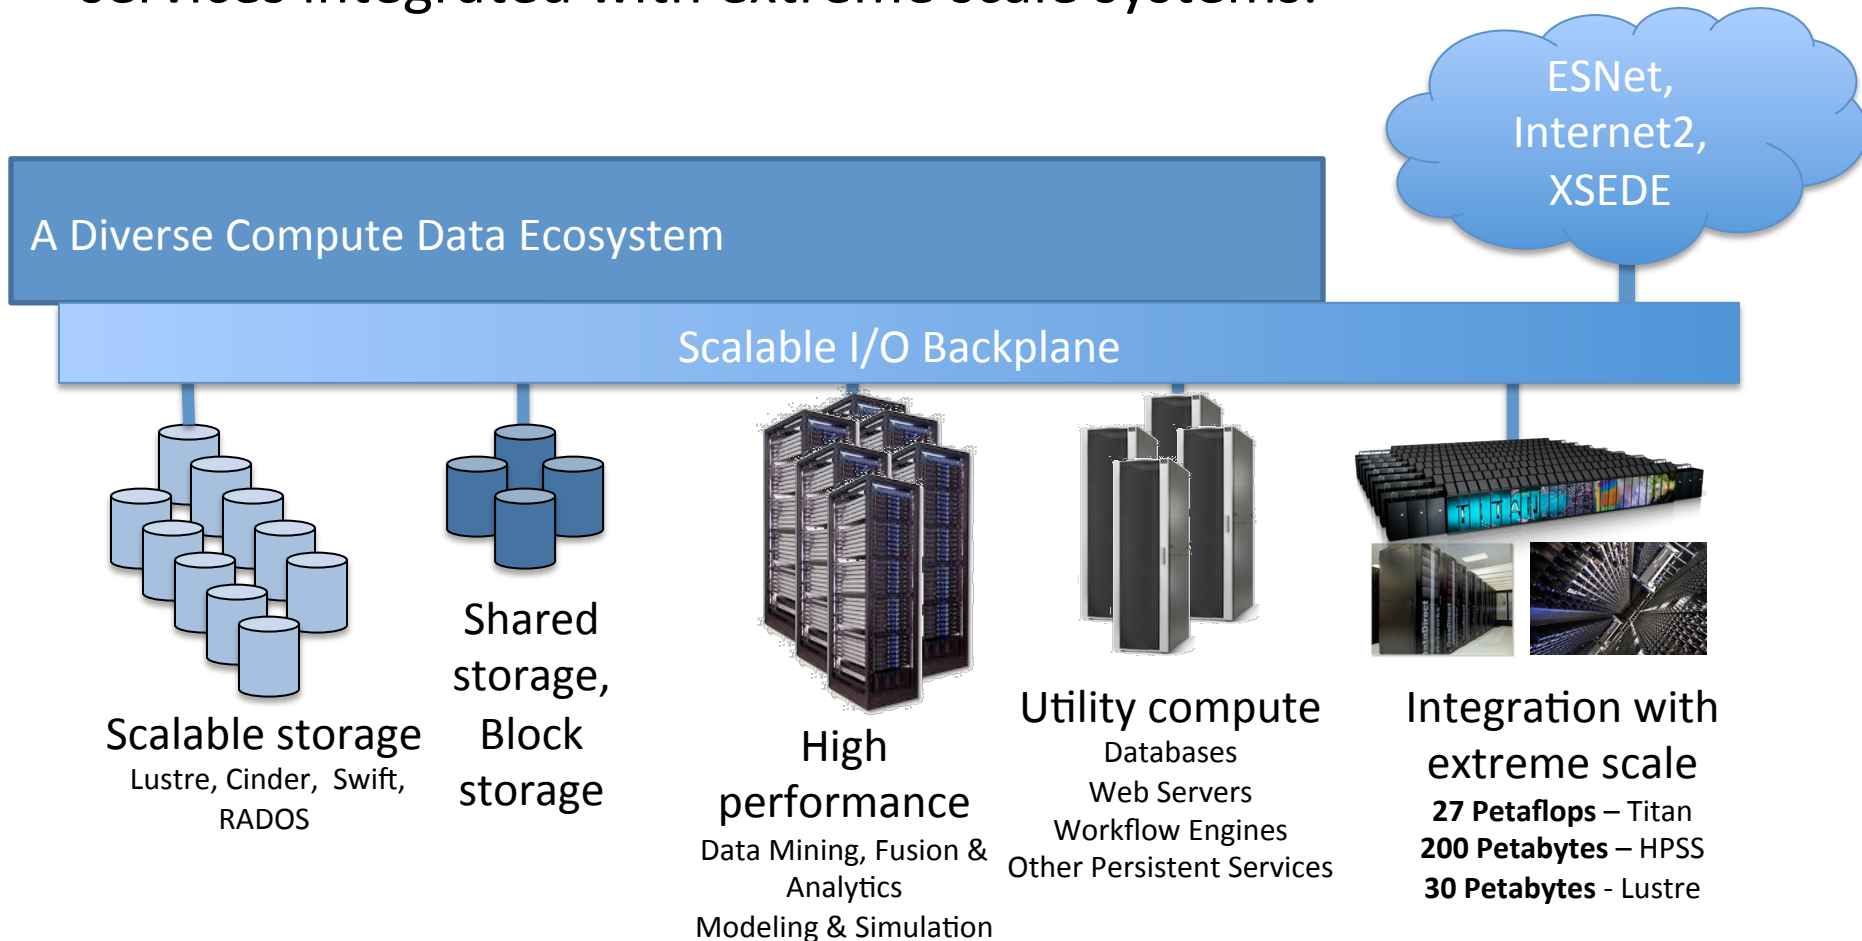


ESnet - ESnet Hub to Small DOE Site Border Throughput Testing

- 2014 Project Proposal:
 - Establish new distributed ALICE T2 facility
 - Continue deployment & operation at LBNL NERSC
 - Establish new T2 site early 2015 at **ORNL CADES**
 - Transition operations from LLNL to ORNL in 2015
- Project Proposal Reviewed, June 2014
 - Proposal endorsed with modest recommendations



- A diverse ecosystem of compute and data infrastructure and services integrated with extreme scale systems.



CADES Tier 2 Site


- Allows ease of expansion over time.
- Provides potential for ALICE workloads to elastically scale to other CADES compute and storage resources.





CADES new home shown in magenta. ALICE T2 system will be two racks (or slightly more) in highlighted area.

CADES Tier 2 Site

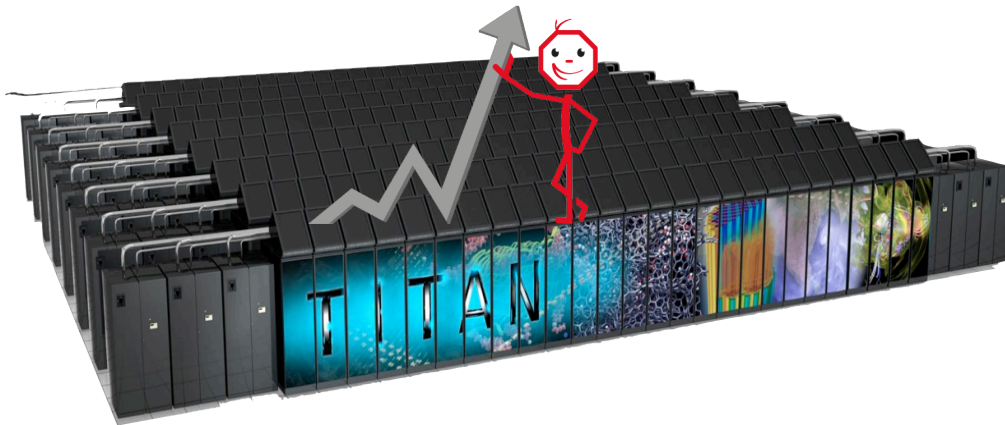
ORNL CACS - Cold Aisle Containment System



T2 Rack layout



PanDA Pilot



- **P**roduction **A**Nd **D**ata **A**nalysis (PanDA) workload management system used by multiple experiments for optimized GRID/cloud job production and data availability. Big PanDA is a DOE ASCR-supported project.
- PanDA-capable projects (i.e., ATLAS and ALICE) recently awarded 10 M hours on Titan over next 12 months.
- ALICE executables and libraries provided via CVMFS do work properly on Titan, with a tailored environment.

ALICE Tests PanDA Pilot

- Now possible to manually launch an ALICE simulation on Titan from elsewhere via a PanDA Development Server on Amazon EC2 (Elastic Cloud Compute) communicating with a PanDA Pilot process on Titan.
- Have launched ALICE simulations on Titan from the NERSC T2 VO box and the ORNL CADES VO box *manually* using PYTHON scripts.
- Work underway this summer:
 - Connect job submission and monitoring to AliEN. (Jeff, others, advice from Latchezar).
 - Optimize/test ALICE simulation production job on Titan. (OLCF postdoc, Ken, advice from Latchezar, others).
 - QA checks (OLCF postdoc, Ken, others).
 - Register output simulation files with AliEN and make available to ALICE (i.e., at NERSC and CADES T2s). (Jeff, others).

