NERSC Operations Review and Plans

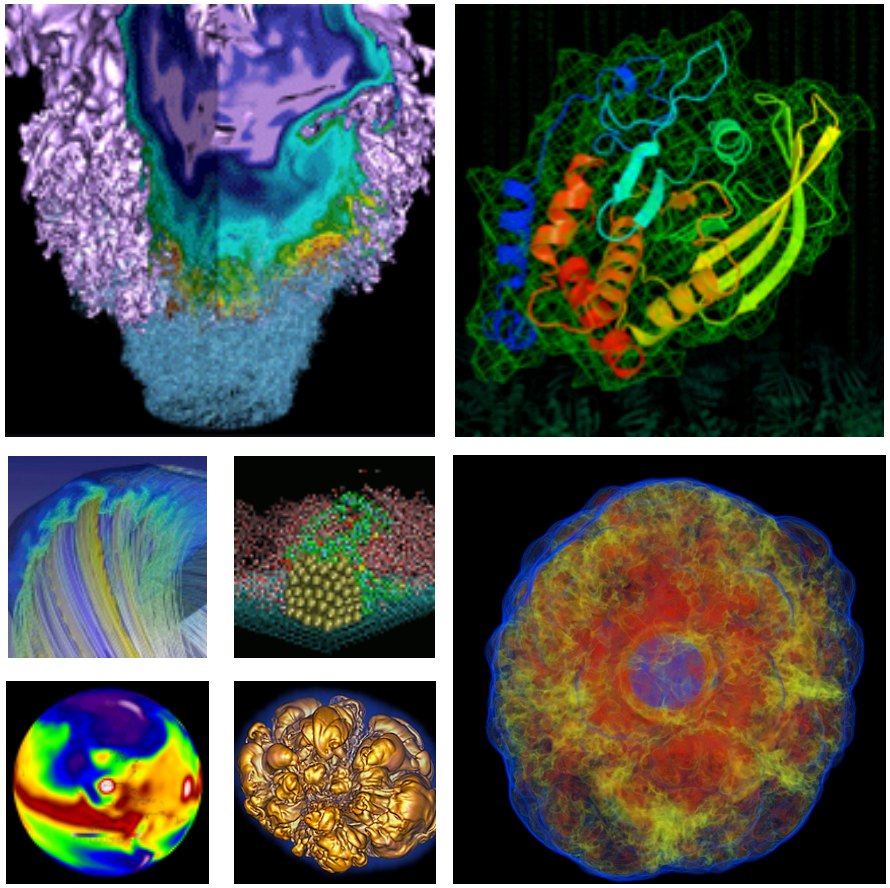

Iwona Sakrejda
Lisa Gerhardt
NERSC

April 8, 2014

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

NERSC Systems Today

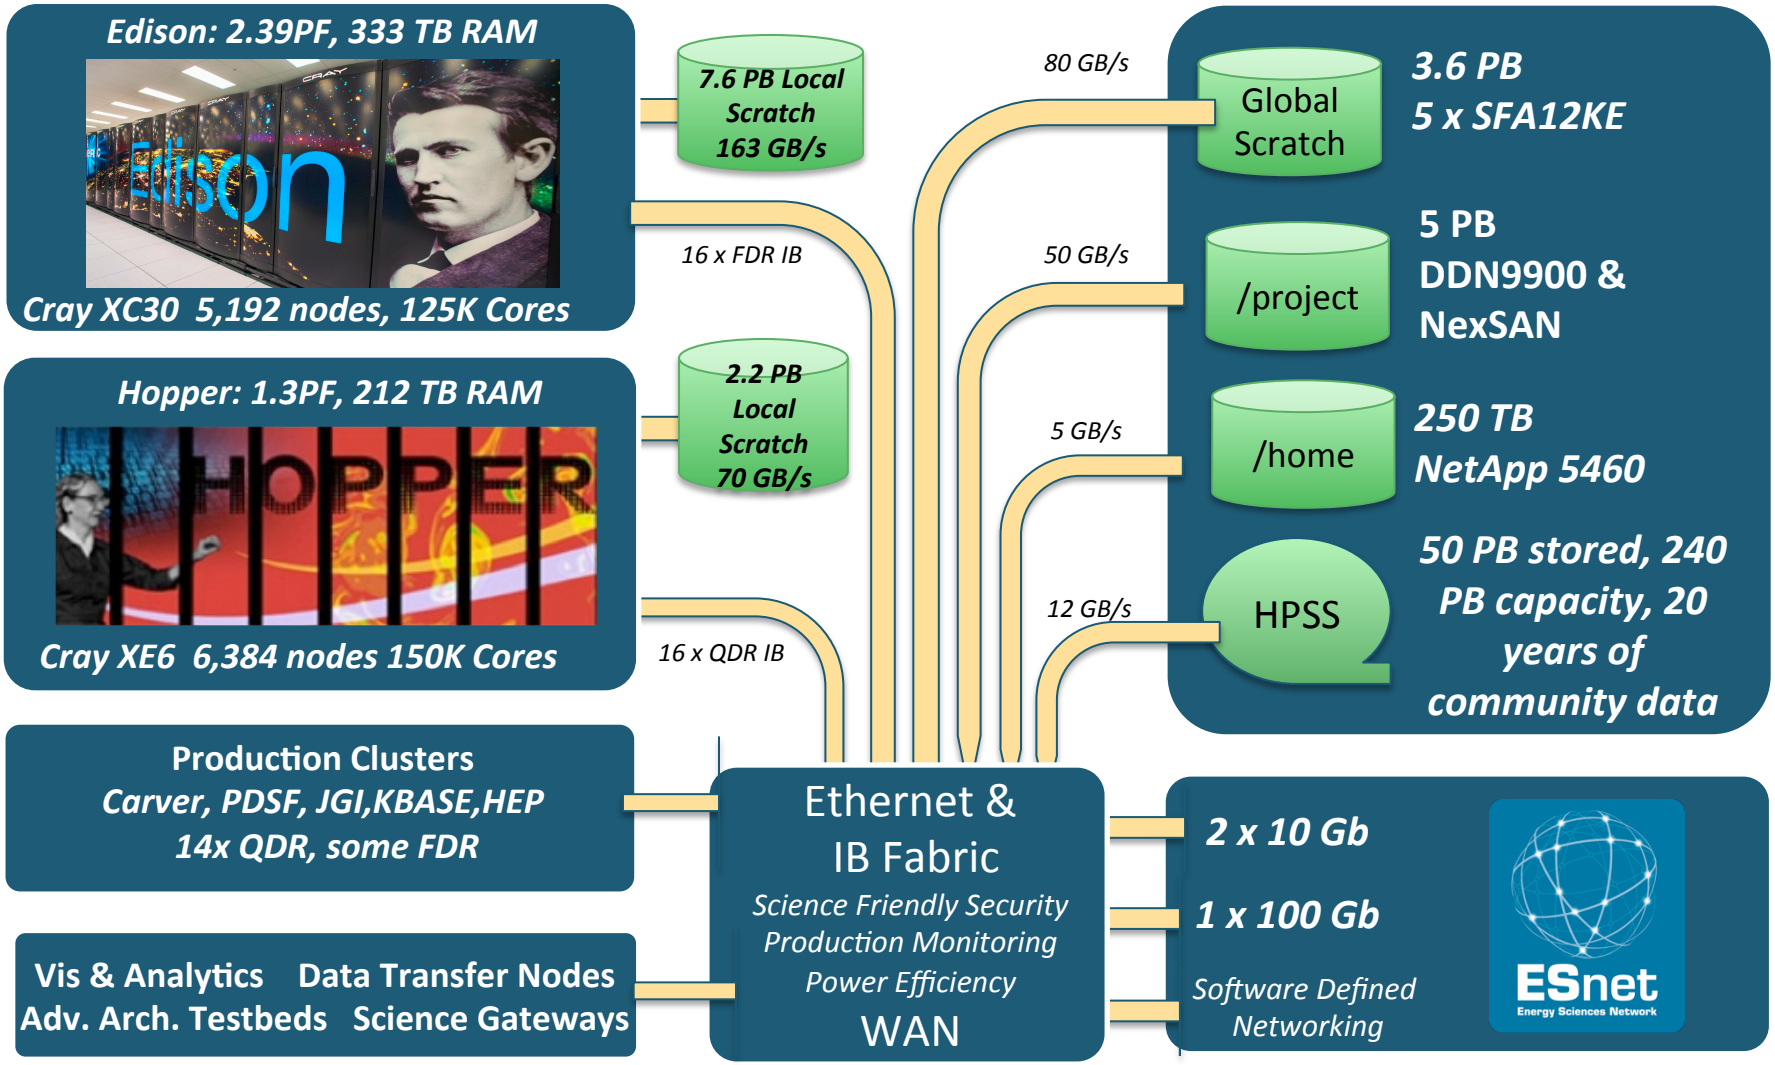

NERSC Systems

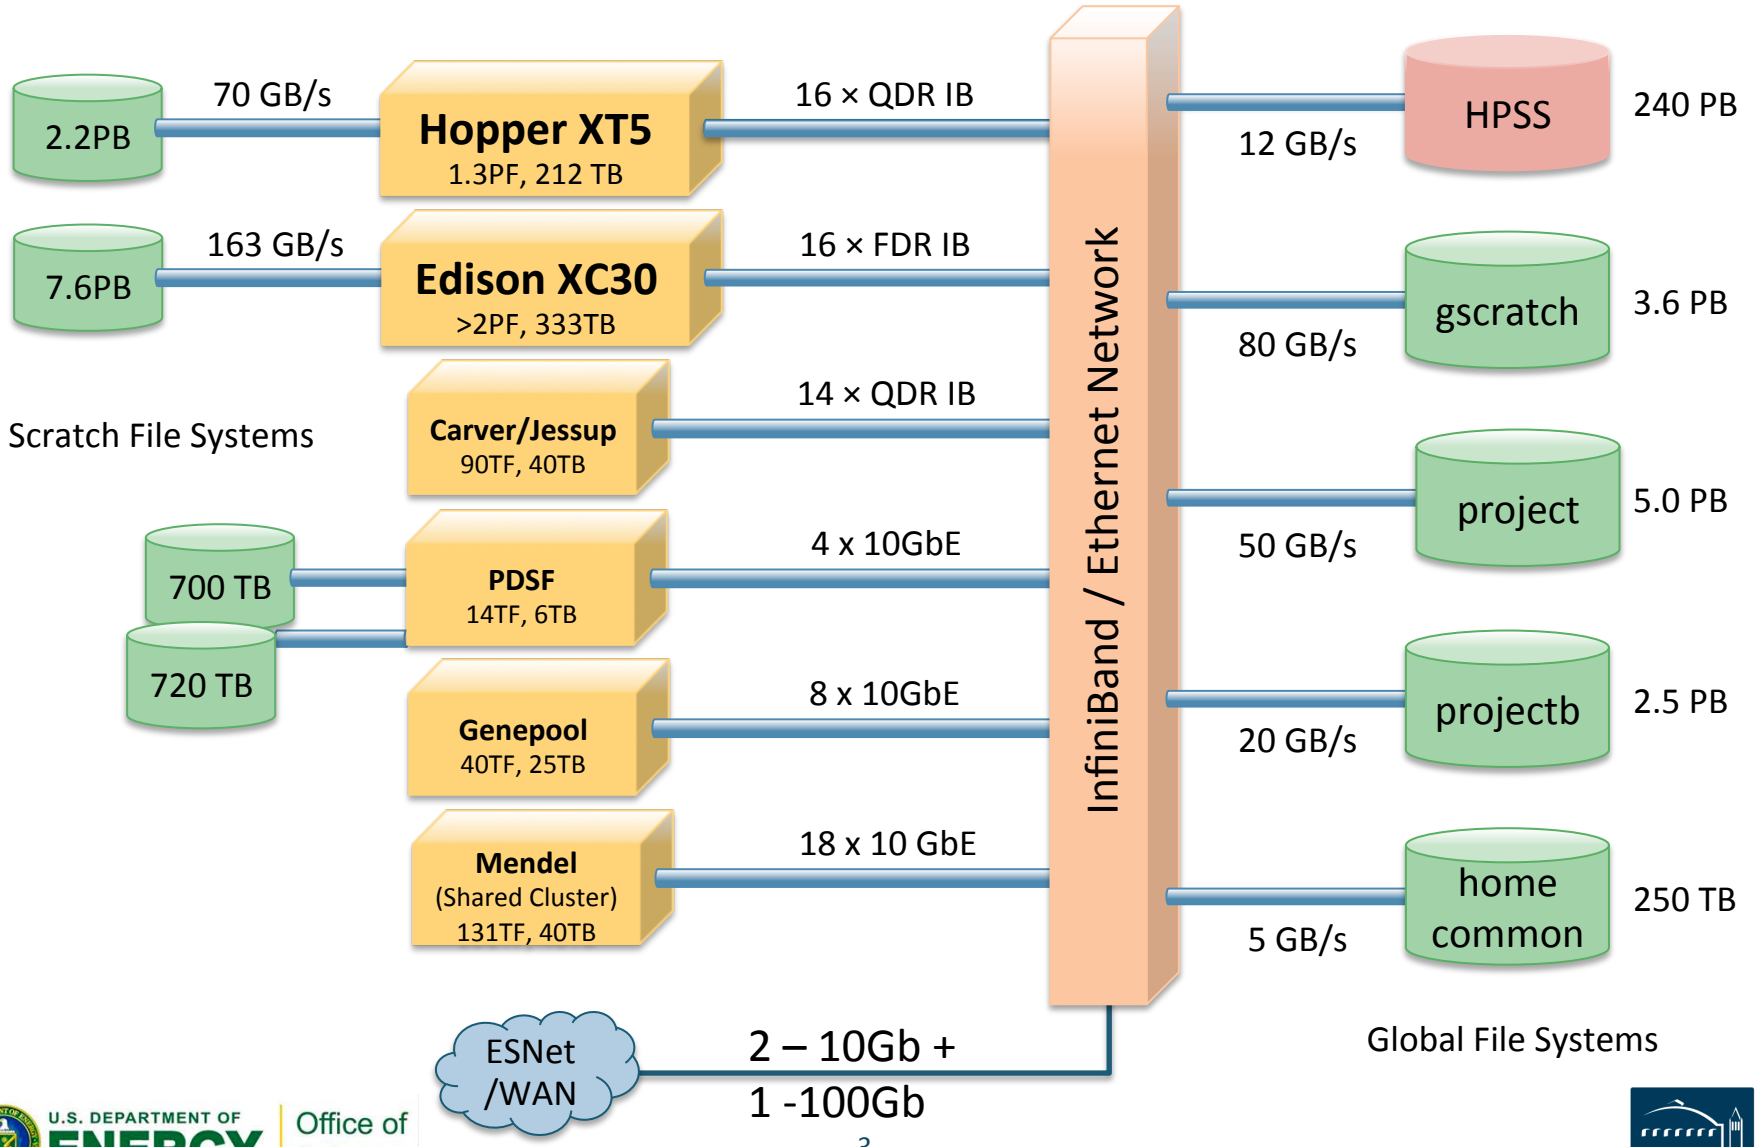

Midrange Systems at NERSC in 2012

- **Three x86_64 midrange computational systems:**
  - PDSF: ~200 node commodity cluster for High Energy Physics and Nuclear Physics; exclusively serial workload; SL 6.2, 6.4 and 5.3 environments; UGE
  - Carver: ~1000 node iDataPlex; mixed parallel and
    - serial workload; Scientific Linux (SL) 5.5;
    - TORQUE+Moab
  - Genepool: ~400 node commodity cluster
    - providing computational resources to the DOE JGI (Joint Genome Institute). Mixed parallel and serial workload; Debian 6; Univa Grid Engine (UGE)
- **All three systems needed expansion**

Mendel Midrange Systems Expansion

- **Each midrange system needed expanded computational capacity**
- **Instead of expanding each system individually, NERSC elected to deploy a single new hardware platform ("Mendel") to handle:**
  - Jobs from the "parent systems" (PDSF, Genepool, and Carver
  - Support services (NX and MongoDB)
- **Groups of Mendel nodes are assigned to a parent system**
  - These nodes run a batch execution daemon that integrates with the parent batch system
  - Expansion experience must be seamless to users:
    No required recompilation of code (recompilation can be recommended)
  - Design facilitates easy re-grouping of nodes between clusters.

The Layered Model

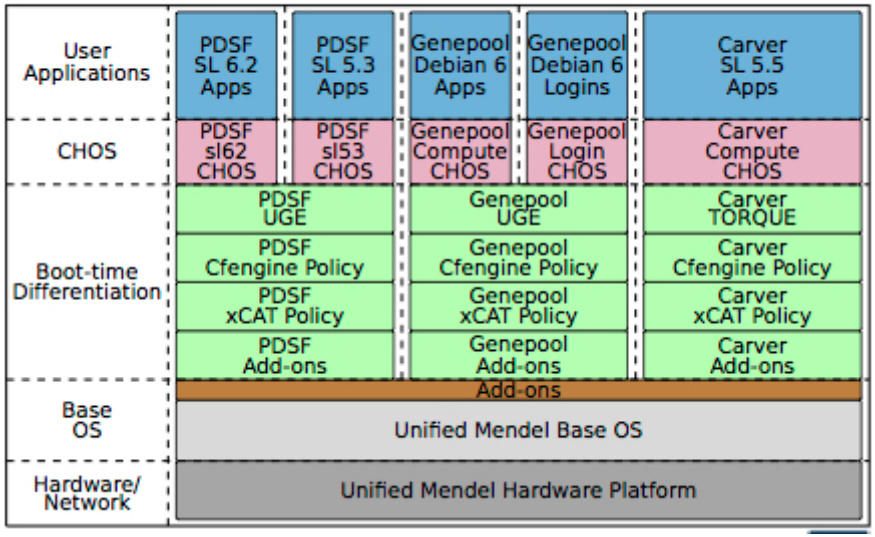

Mendel Hardware

- **Vendor: Cray Cluster Solutions (formerly Appro)**
  - Scalable Unit expansion model
- **FDR InfiniBand interconnect with Mellanox SX6518 and SX6036 switches**
- **Compute nodes are half-width Intel servers**
  - S2600JF or S2600WP boards with on-board FDR IB
  - Dual 8-core Sandy Bridge Xeon E5-2670
  - Multiple 3.5” SAS disk bays
- **Power and airflow: ~26kW and ~450 CFM per compute rack**
- **Dedicated 1GbE management network**
  - Provisioning and administration
  - Sideband IPMI (on separate tagged VLAN)

PDSF Hardware Layout

- **XROOTD Storage Cluster**
  - 10 Dell RPower 710 servers
  - MD1200 Dell JBODs direct SAS attached (4 per server)
  - Dell RPower 410 server (redirector)
  - 10Gb/s ethernet
- **PDSF Compute Cluster**
  - ~200 Dell RPower 410 Servers (8, 12 cores, memory mostly 4GB/core)
  - 68 Mendel Servers (16 cores, memory 4GB/core, FDR IB – 30Gb/s link between PDSF core router and Mendel)
  - 3 hosts behind load balancer for interactive access
  - 4 backup interactive nodes used for special services and development
  - Auxiliary servers (mostly Dell RPower 410)
    - 2 VO boxes with Condor-G
    - 2 CE gatekeepers with SGE job managers
    - 2 SGE servers (master, shadow) for reliability
    - 2 admin servers (managing deployment and configuration)
- **Networking**
  - Combination of Dell, HP and Cisco switches
  - Cisco core router
    - 2x10Gb/s connection to other NERSC systems/storage
    - 2x10Gb/s connection to the border router

PDSF Summary

- **Evergreen and condo**
- **2600 cores with 20 – 60 GB RAM, 230 nodes**
- **690 TB dedicated GPFS storage**
- **720 TB of storage dedicated to Alice in XROOTD SE**
- **Used by Nuclear Physics and High Energy Physics groups**
  - Simulation and analysis
  - Data mirrors

PDSF Shares 2014

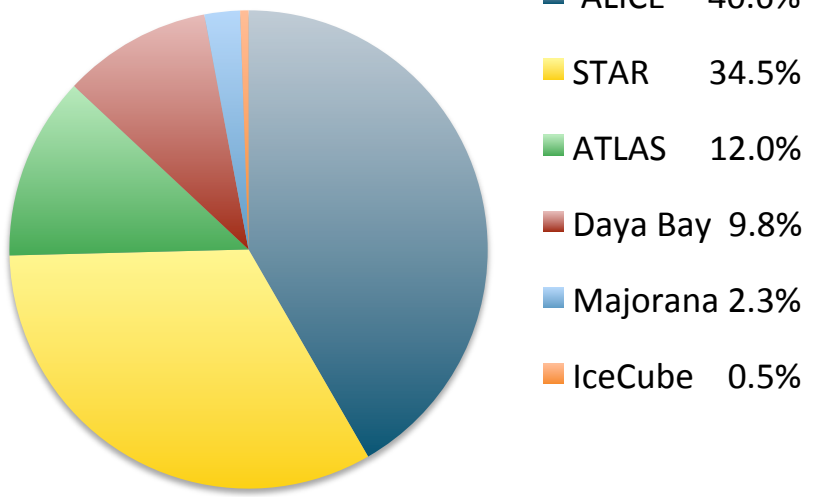

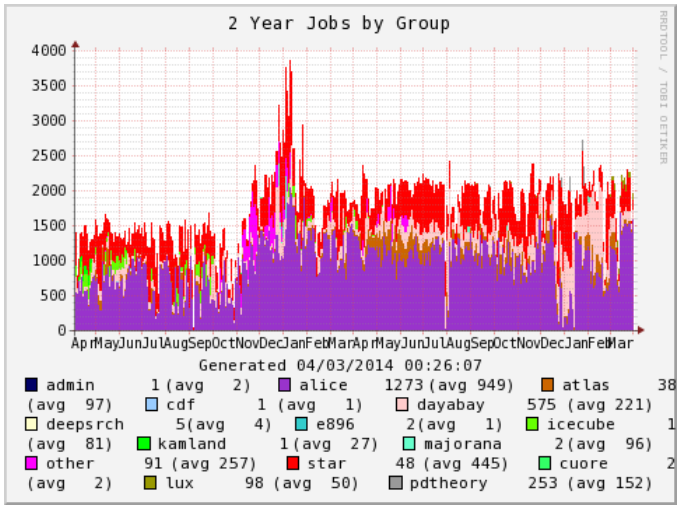

PDSF Utilization by Alice

- **Due to redundancy in configuration number of downtimes minimized**
  - 5/21, 7/30, 10/28 (rolling) , 2/11 Center-wide maintenance
  - 11/(12-16) cable replacement in Mendel
  - 3/29 slow submissions due to a hardware issue on vo box

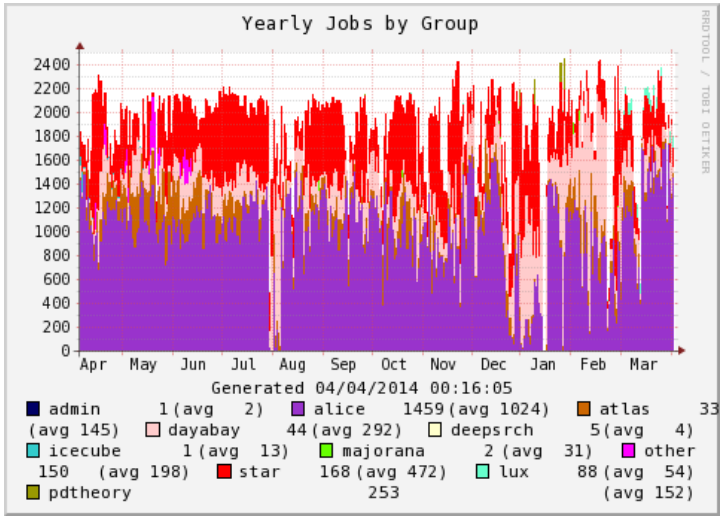

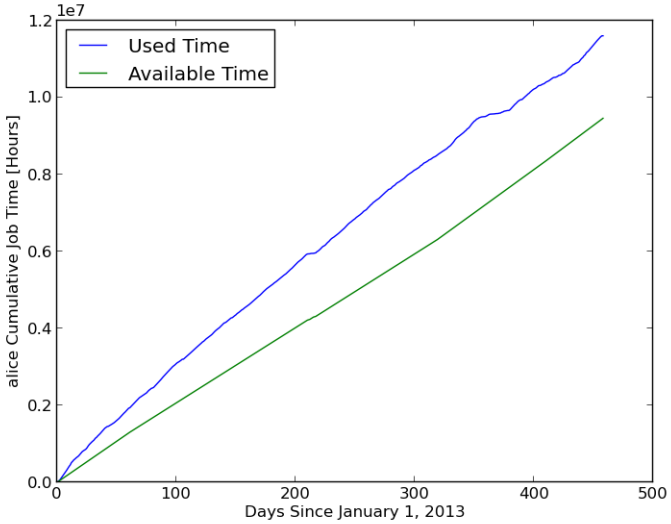

- **Alice user 9.3M hours in 2013**

PDSF Monitoring and Operations support

- **Nagios monitoring of hardware and services**
- **Ganglia monitoring – used in diagnostics and problem resolution.**
- **24x7 Alarm monitoring by operations**
  - Operators trained to resolve simple issues
  - 24x7 escalation to CSG Systems Engineer on call

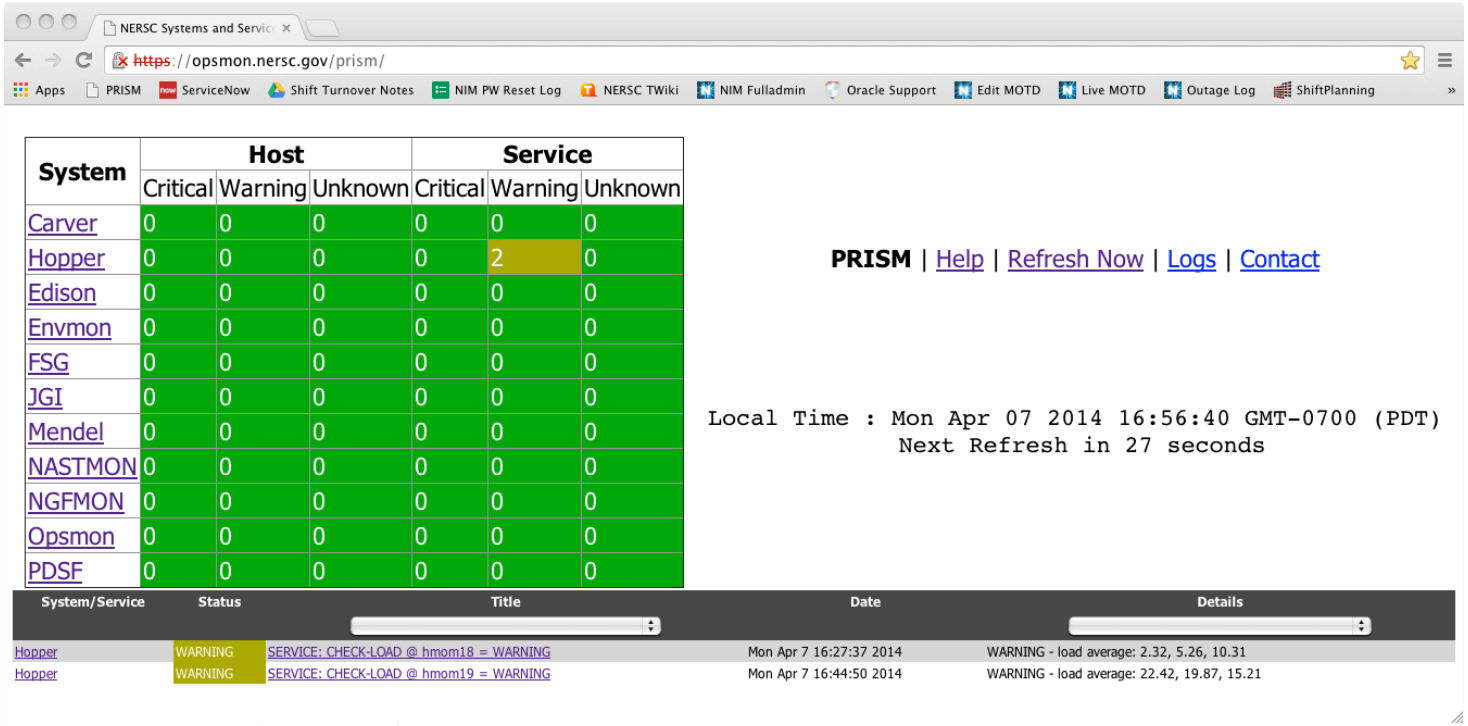

System	Host			Service		
	Critical	Warning	Unknown	Critical	Warning	Unknown
Carver	0	0	0	0	0	0
Hopper	0	0	0	0	2	0
Edison	0	0	0	0	0	0
Envmon	0	0	0	0	0	0
FSG	0	0	0	0	0	0
JGI	0	0	0	0	0	0
Mendel	0	0	0	0	0	0
NASTMON	0	0	0	0	0	0
NGFMON	0	0	0	0	0	0
Opsmon	0	0	0	0	0	0
PDSF	0	0	0	0	0	0

System/Service	Status	Title	Date	Details
Hopper	WARNING	SERVICE: CHECK-LOAD @ hmom18 = WARNING	Mon Apr 7 16:27:37 2014	WARNING - load average: 2.32, 5.26, 10.31
Hopper	WARNING	SERVICE: CHECK-LOAD @ hmom19 = WARNING	Mon Apr 7 16:44:50 2014	WARNING - load average: 22.42, 19.87, 15.21

Alice OSG Monitoring

- **Gratia reporting forwarded to WLCG**
  - Lisa Gerhardt took over development and maintenance of the SGE probe to assure accurate reporting
- **RSV monitoring (connected to the NERSC monitoring system triggers alarms 24x7)**
- **BDII Monitoring (raw data gets to OSG)**
- **Manual monitoring of Alice web pages with job information (Lisa/Jeff) –to be automated soon.**

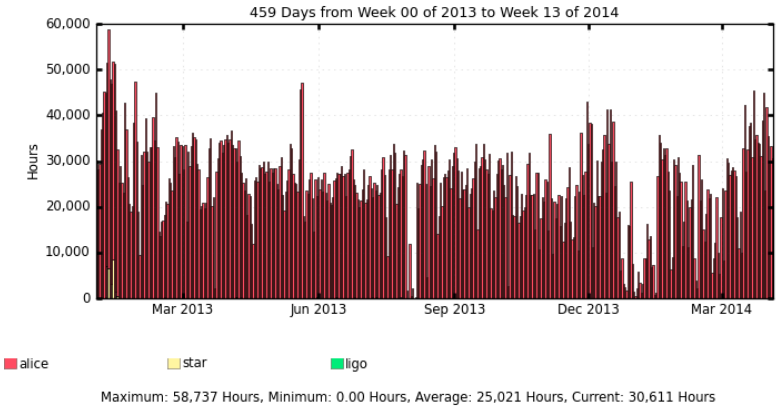

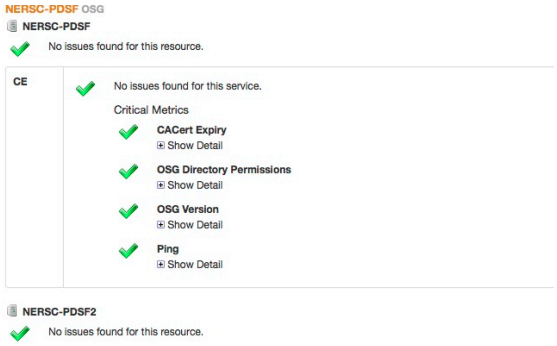

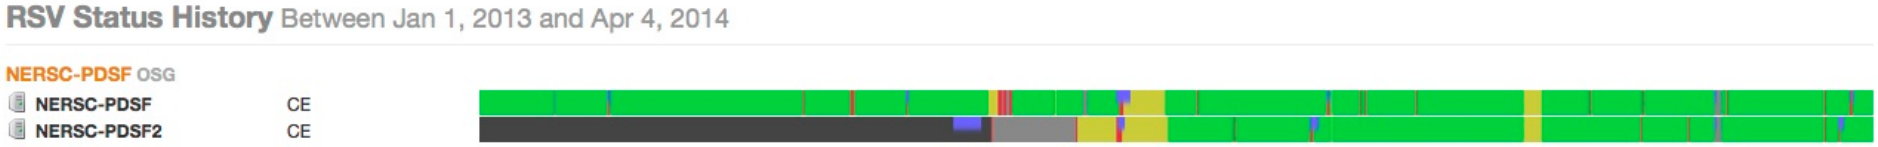

Alice Job Flow at PDSF

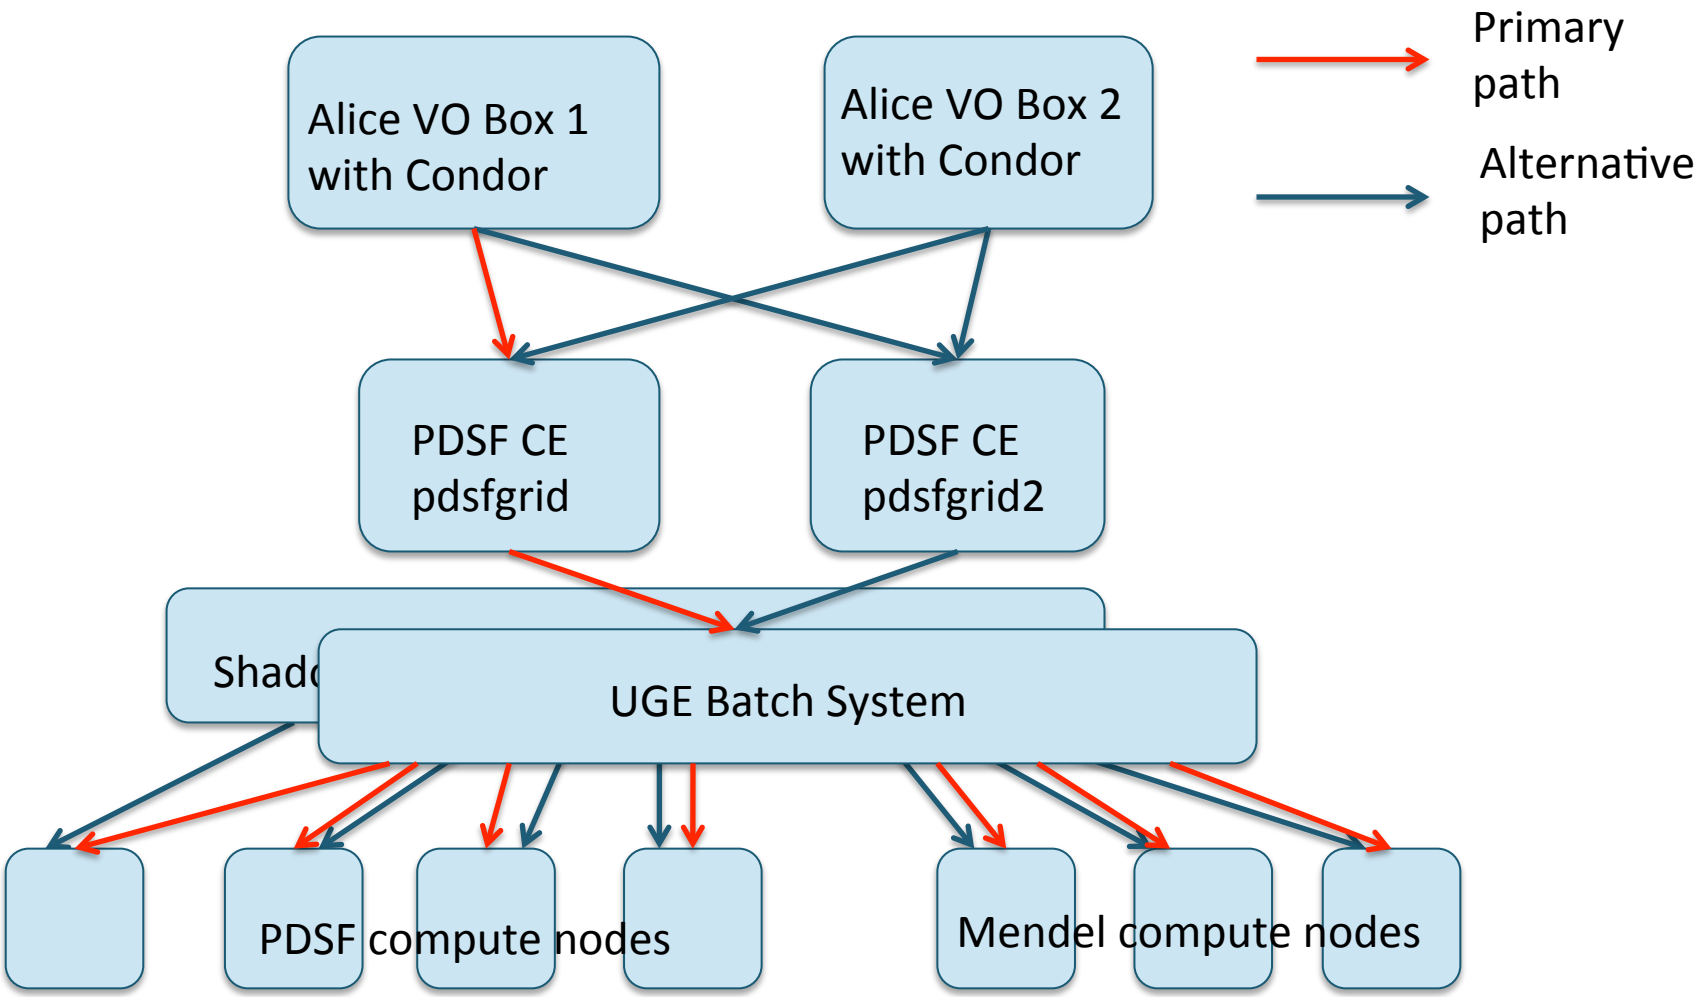

Move to the Hill (CRT)

- **Goal – execute the move with minimal disruption to production services**
- **Preparation activities early Summer 2014**
  - Move backup services from the "old" PDSF to Mendel
    - Shadow batch scheduler
    - Secondary VO box
    - OSG Infrastructure
  - Setup minimal second SE XRootD cluster within Mendel infrastructure
- **Aim for new storage procurement with delivery in Jan/Feb 2015**
- **Connectivity between OSF and CRT – 100Gb/s (might be 400Gb/s)**
- **Run the "old" hardware at OSF until PDSF@Mendel is fully functional and then send to salvage.**
  - Savings in moving expenses
  - All of the "old" hardware out of warranty.
    - 50% of storage warranty expired 10/2013
    - 50%of storage warranty expires this summer
    - Using compatible drives from decommissioned computes to support failing storage.
    - 75% of computes – warranty already expired >6months
    - 25% of computes – warranty will expire November 2014
    - Only 35% of HS06 in the old hardware.

PDSF Move Schedule

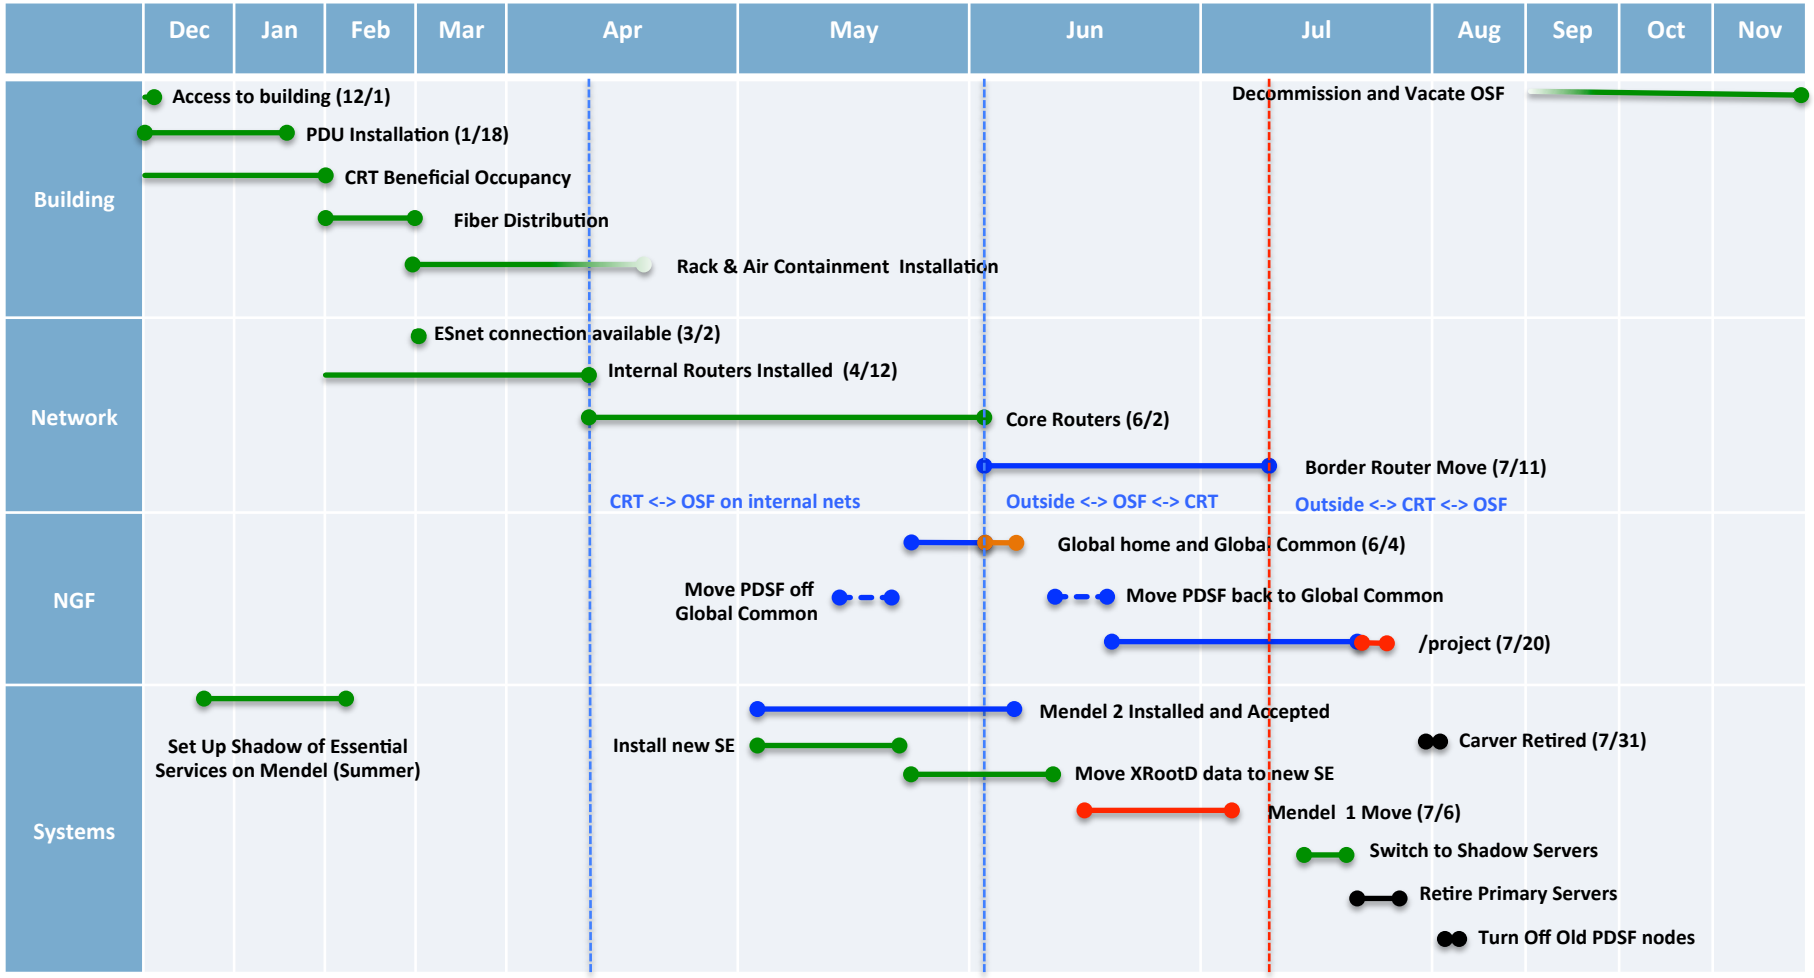

Dual-homed XRootD servers

- **Dual – homed XRootD servers to separate inside access (over 36 net) from the outside one.**
- **Second 10 Gb port on the servers configured.**
- **Switch set up and all the fibers pulled.**
- **To do:**
  - Set rp_filters
  - Set routes on computes
  - Register servers on the external interface with the redirector

Dual Homed Xrootd Servers

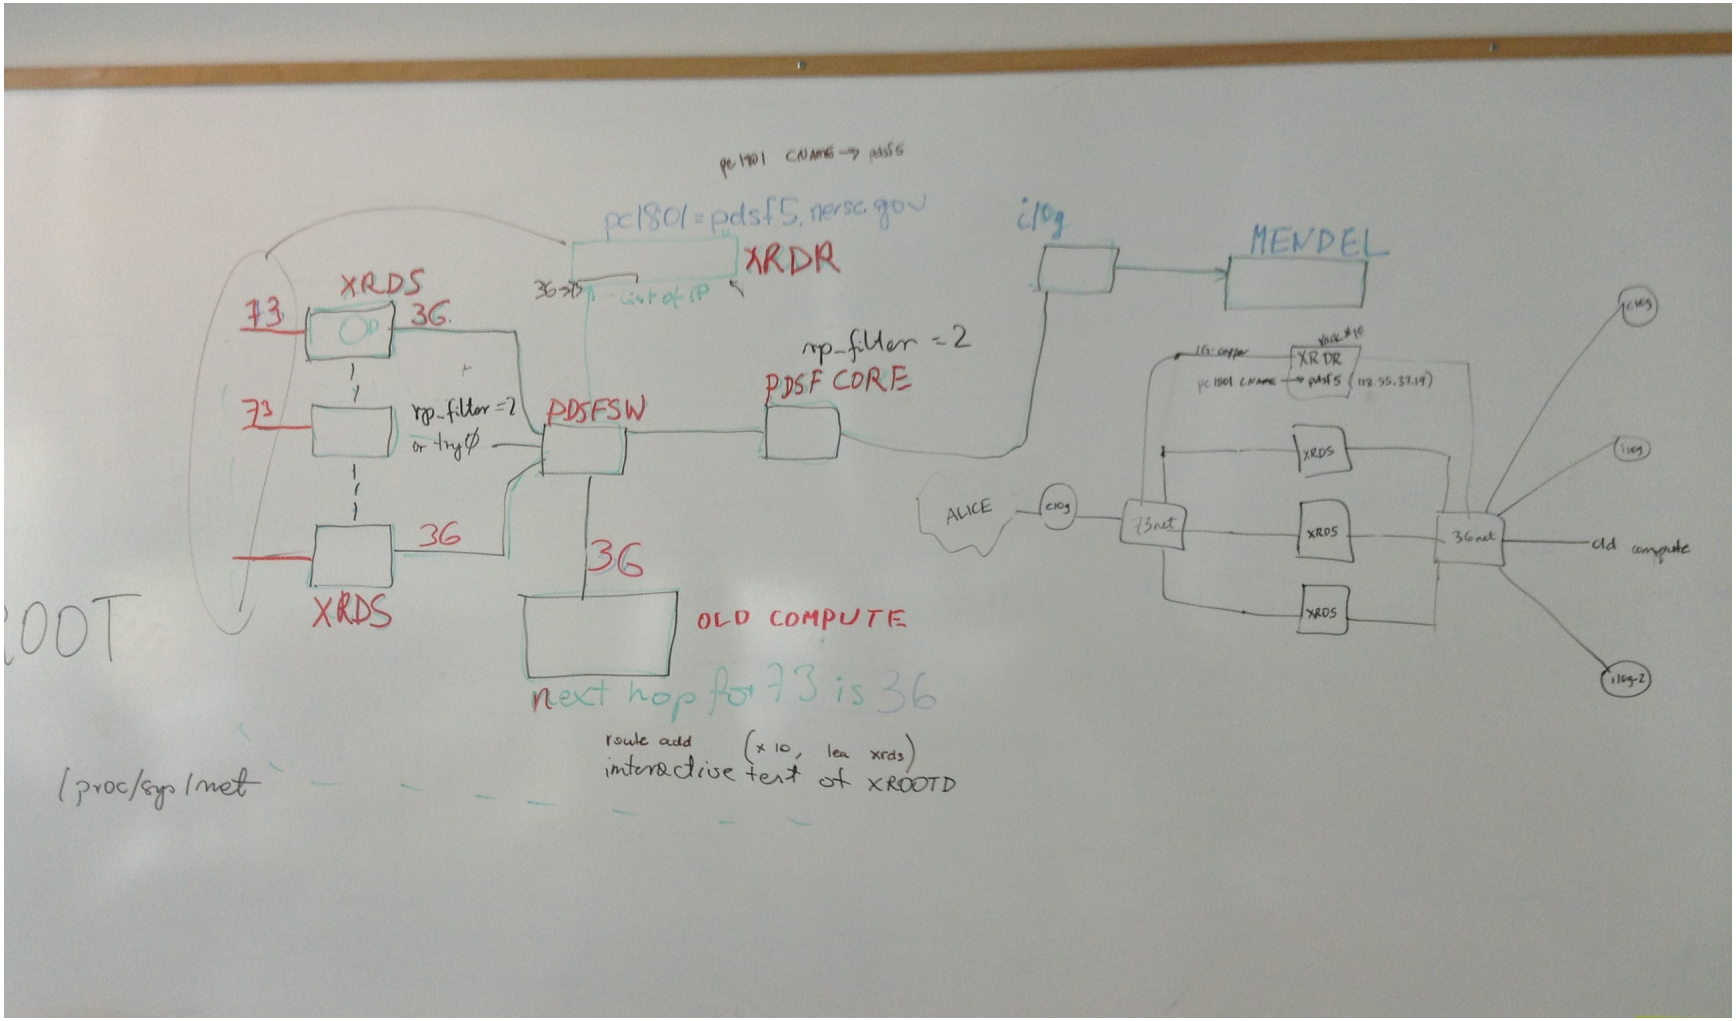

Dual Networking for PDSF XRootD Servers

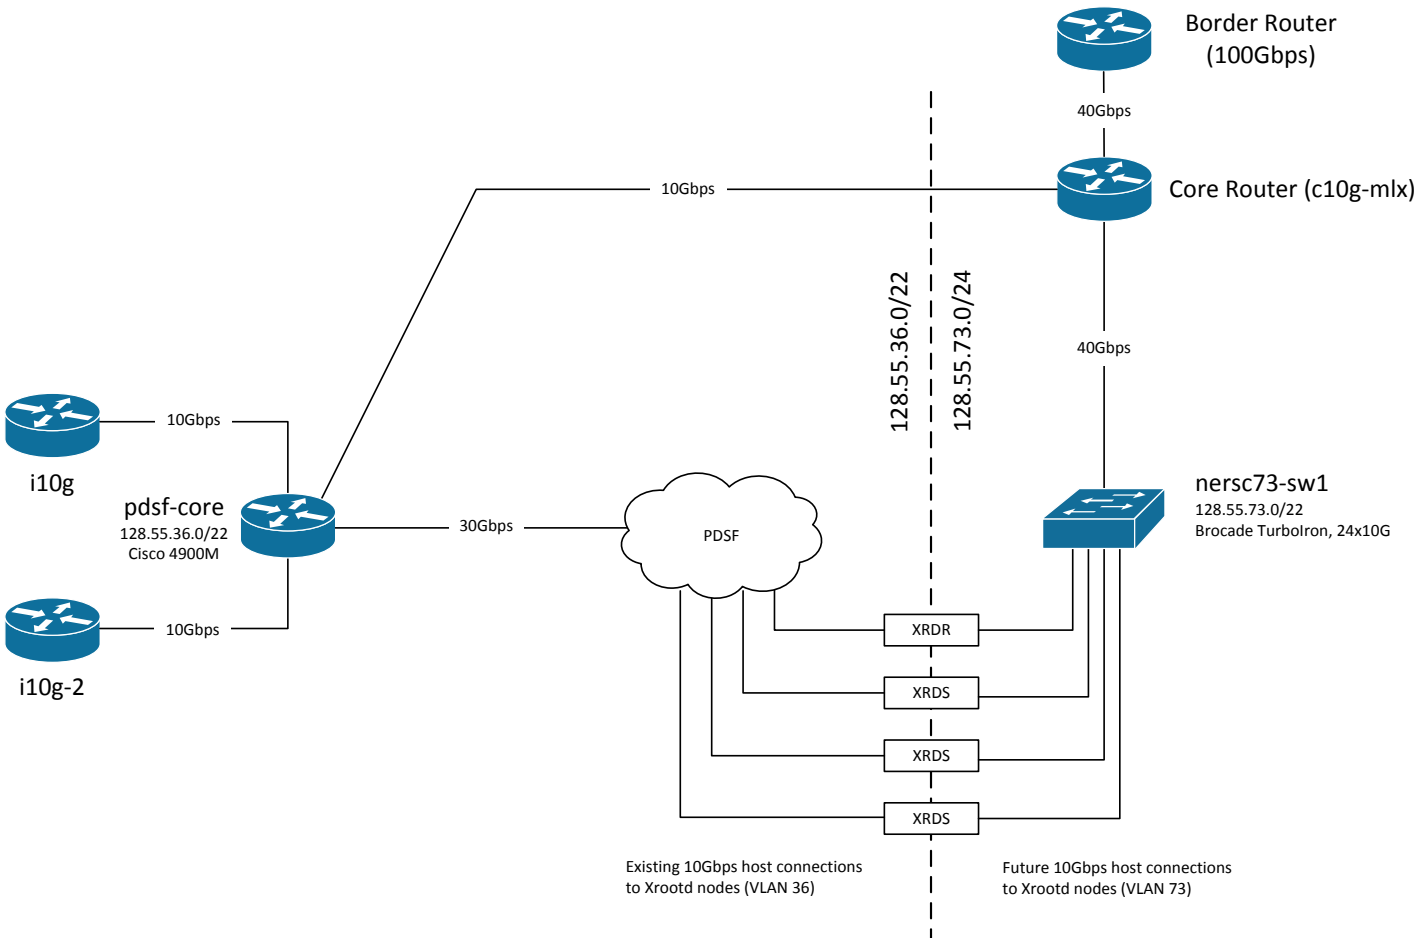

Expanding Production to the Crays

- **Working to get parrot online on Cray computes**
  - Turns cvmfs calls into http calls
  - Enables cvmfs interface without cvmfs
  - Needs local proxy inside NERSC domain
- **SAGA job submitter**
  - Following work at Titan on Cray version of ALICE job submitter