

Fermilab Site Report Spring 2013 HEPiX

Keith Chadwick

Grid & Cloud Computing Department
Fermilab

Work supported by the U.S. Department of Energy under contract No. DE-AC02-07CH11359

Acknowledgements

I am reporting on the accomplishments of the many individuals and organizations of the Fermilab Computing Sector,

Please credit them with all the accomplishments and assign to me any misunderstandings.

ITIL Service Management

Fermilab has commissioned all of the ITSM processes =>

Fermilab has a job opening for a Service Level Manager,

- Jack Schmidt has moved to ANL

New people have been assigned to ITIL roles:

- Capacity Manager is now Robert Harris,
- Service Continuity and Availability Manager is now Mike Diesburg.



ISO 20000

On 13-Feb-2013, Fermilab became the first DOE laboratory to earn ISO 20000 certification for excellence in information technology (IT) service management processes.

This certification was the culmination of a five-year program led by Fermilab's Computing Sector.

More information is available at:

- http://www.fnal.gov/pub/today/archive/archive_2013/today13-02-14.html
- <http://cd-docdb.fnal.gov/cgi-bin/ListBy?topicid=190>



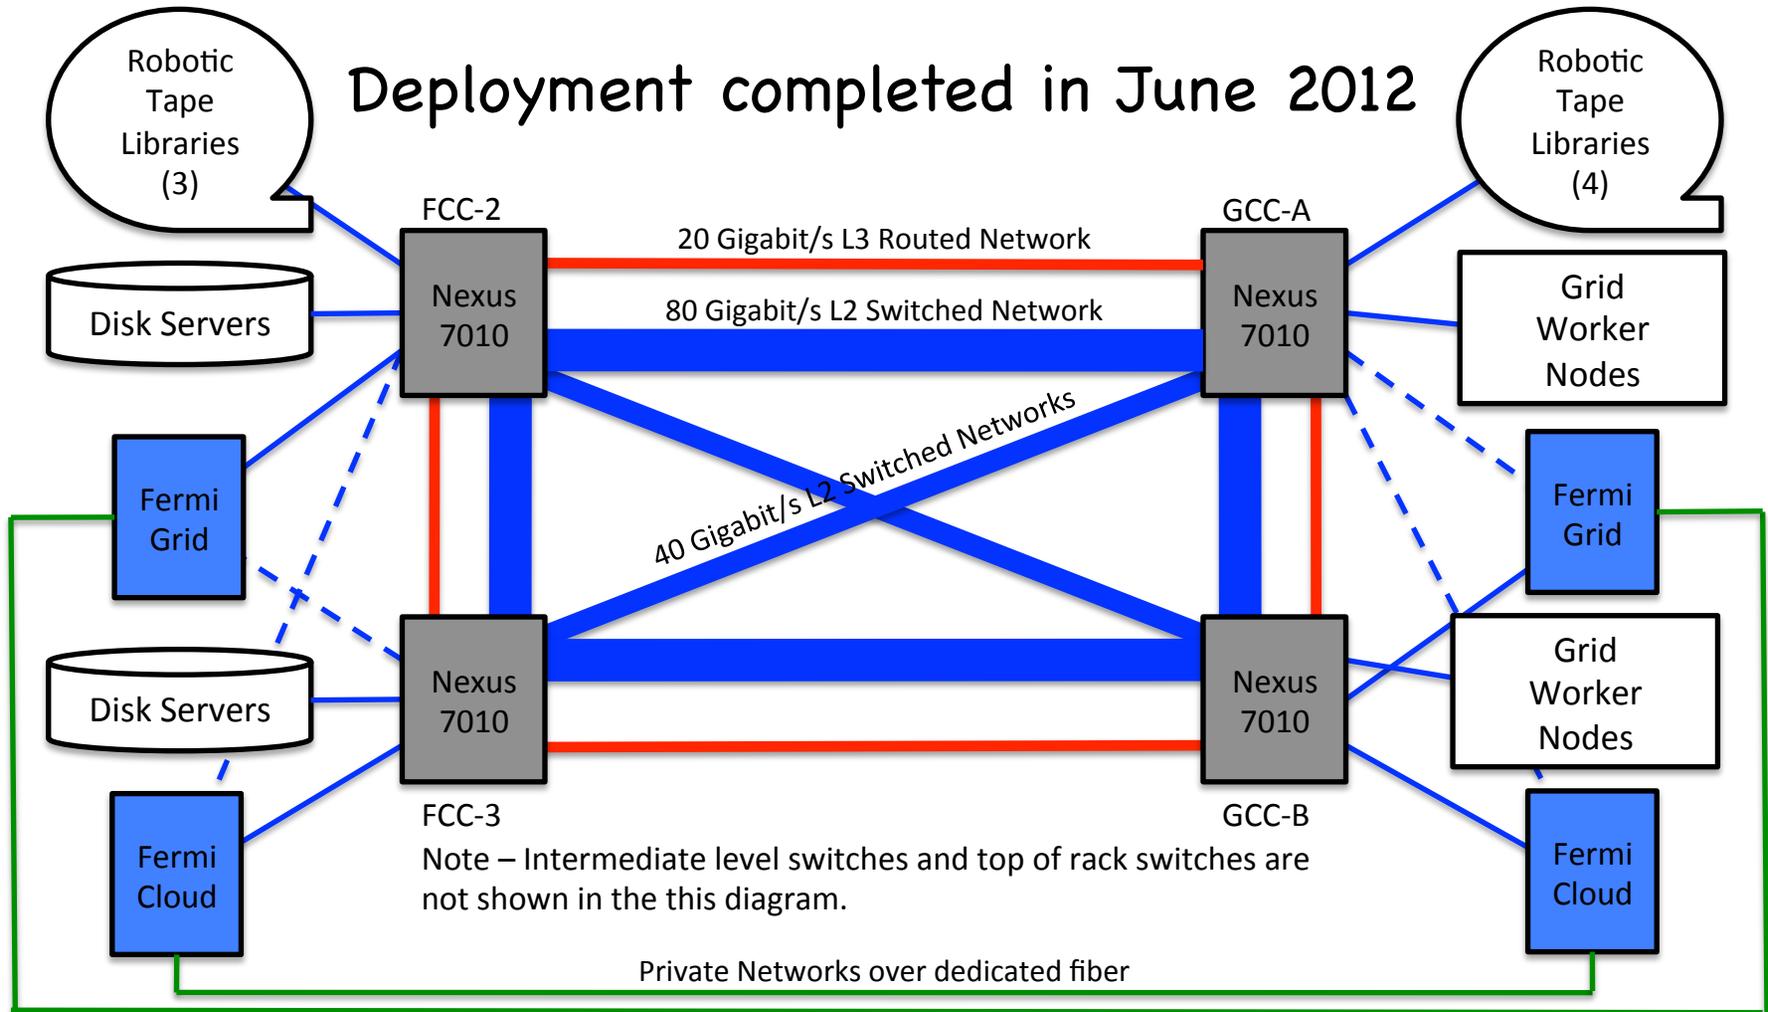
The ISO 20000 Certification Award Ceremony



ISO 20000 Team



Distributed Network Core Provides Redundant Connectivity



Current Fermilab WAN Capabilities

Metropolitan Area Network provides 10GE channels:

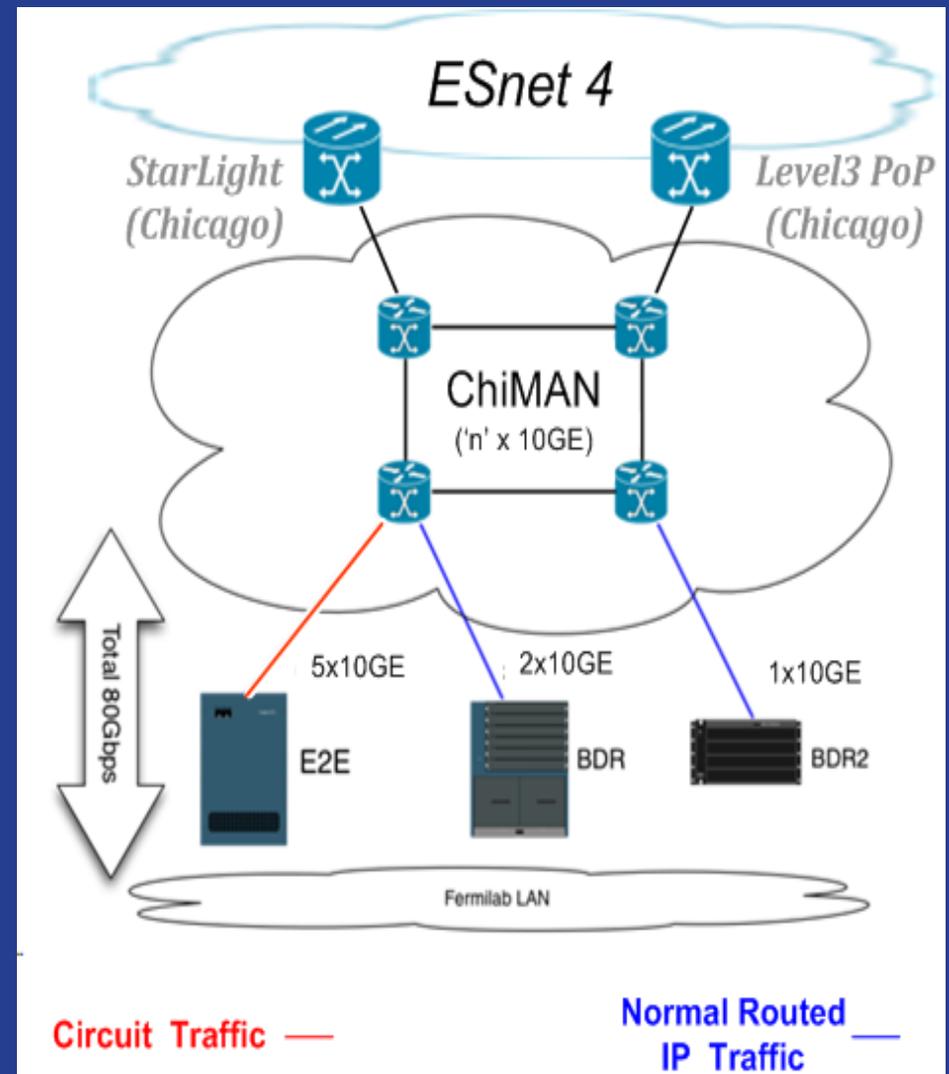
- Currently 8 deployed

Five channels used for circuit traffic

- Supports CMS WAN traffic

Two used for normal routed IP traffic

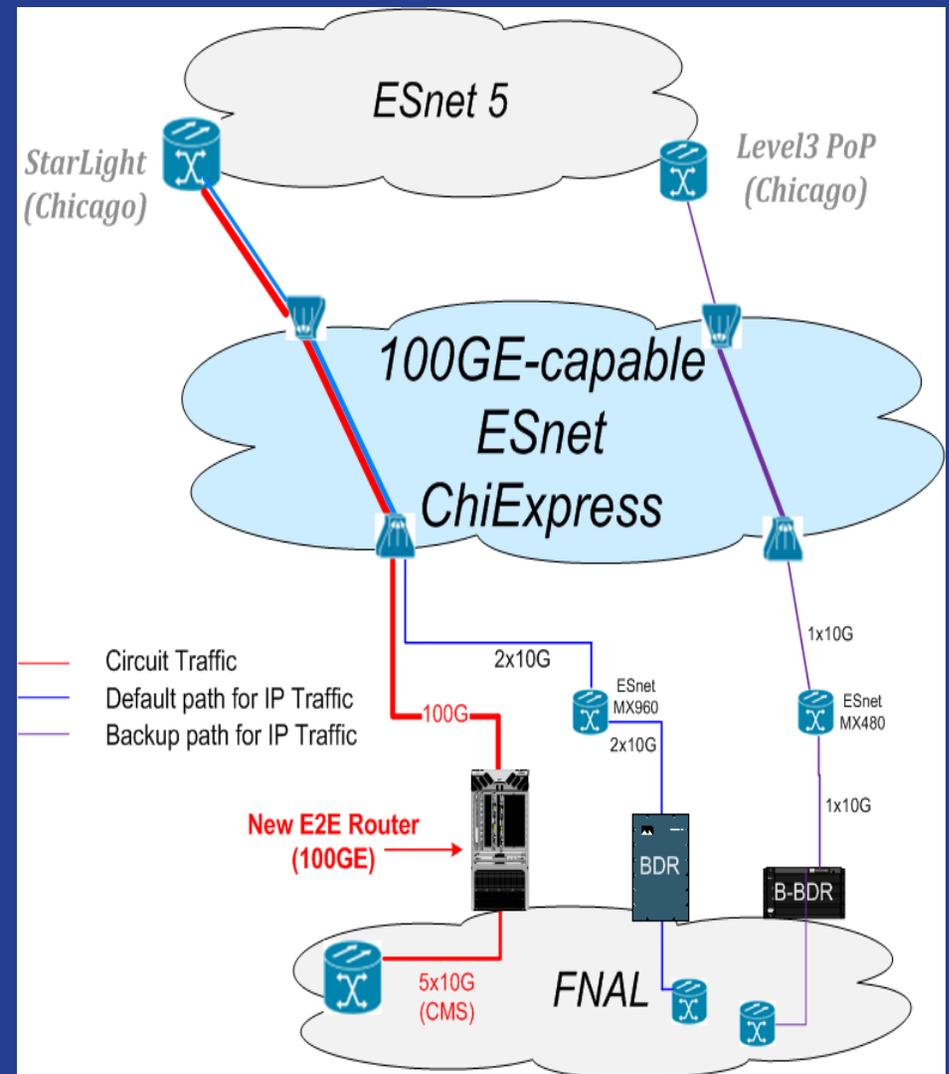
- Backup 10GE for redundancy
- Circuits fail over to routed IP paths



Near-Future Fermilab WAN Capabilities

ESnet ChiExpress MAN:

- **One 100G channel**
 - Circuit-based high impact science data traffic
 - Network R&D activities
- **Three 10G channels**
 - For default routed IP traffic
- **Full geographic diversity within MAN**
- **Production deployment in spring of 2013**



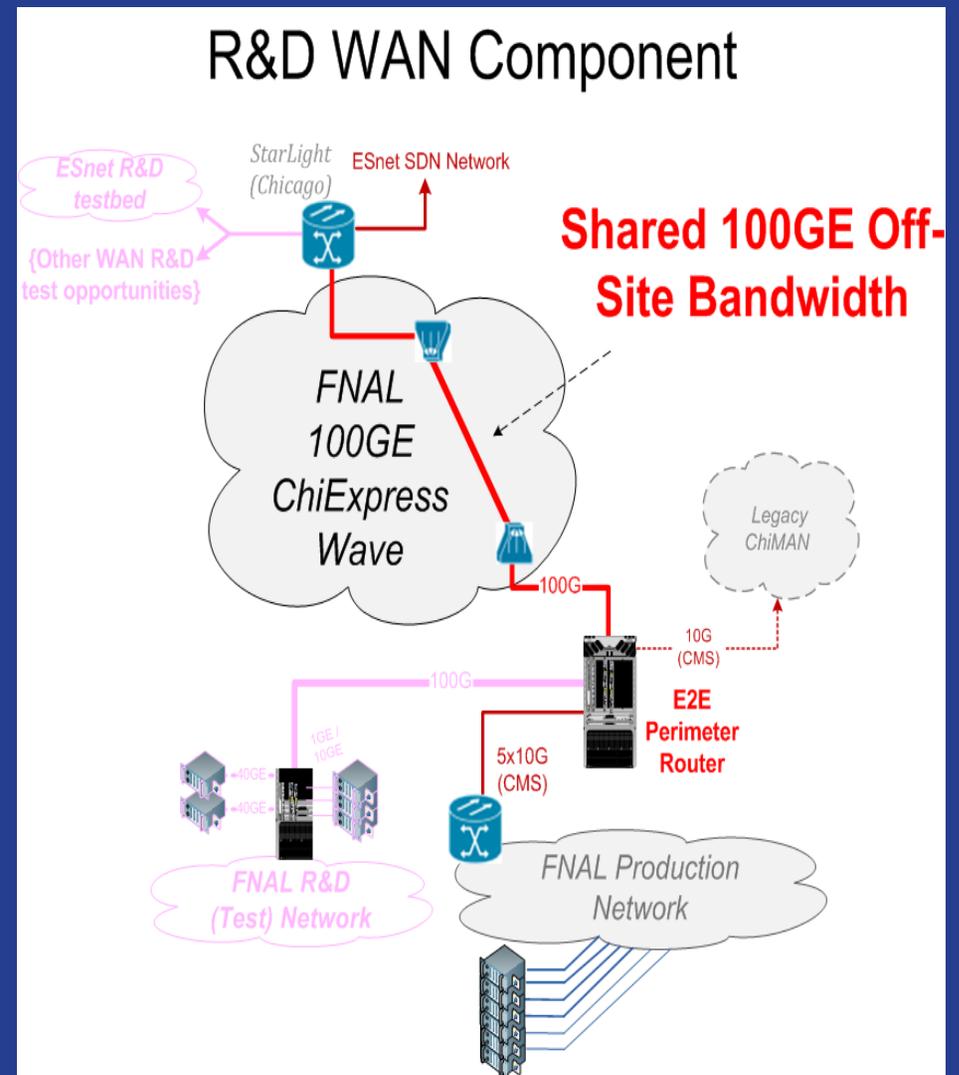
Use of 100G Wave for FNAL R&D Test Bed

100G wave will support
50G of CMS traffic,

Leaving ~50G for FNAL
R&D network (Potentially
higher when CMS traffic
levels are low),

Planning WAN circuit into
ESnet 100G testbed:

- Potential for circuits to
other R&D collaborations



A dedicated R&D Network facility

- 40GE/100GE R&D
- Production-like env for tech eval
- Testing of firmware upgrades

100 Gb/s capable switch:

- 2-port 100GE module
- 6-port 40GE module
- 10GE copper module

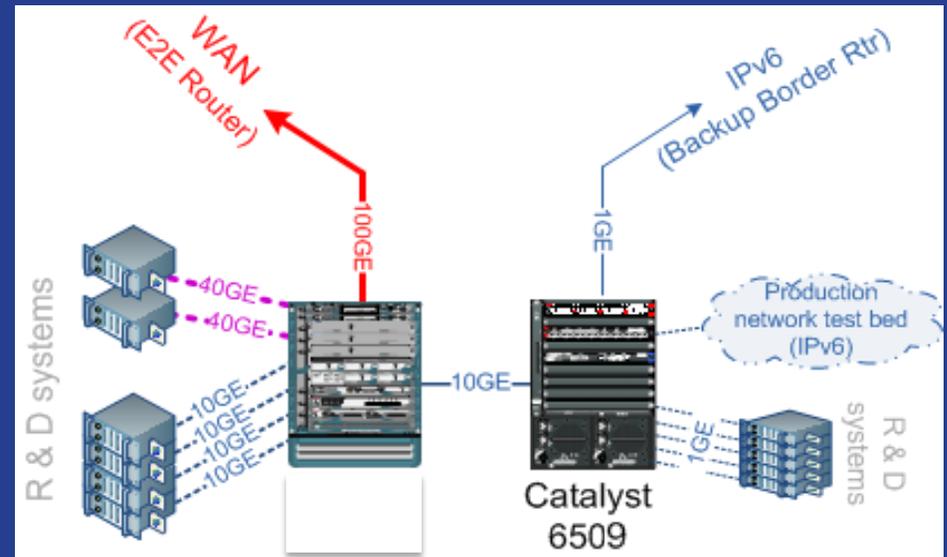
12 nodes w/ 10Gb Ethernet:

- Intel X540-AT2 (PCIe) / 8 cores / 16 GB RAM

2 nodes w/ 40Gb Ethernet:

- Mellanox ConnectX®-2 (PCIe-3) / 8 cores w/ Nvidia M2070 GPU

Full implementation delayed until end of year due to limitations of current 40GE/100GE modules



Catalyst 6509E for 1GE systems:

- IPv6 tests / F5 load balancer / Infoblox DNS, Palo Alto firewall

Computing Facilities

FCC:

- The FCC3 computer rooms have been operating without any disruption since commissioning in December 2010,
- FCC2 computer room had 6 hours of scheduled downtime, corresponding to >99.9% uptime.

GCC:

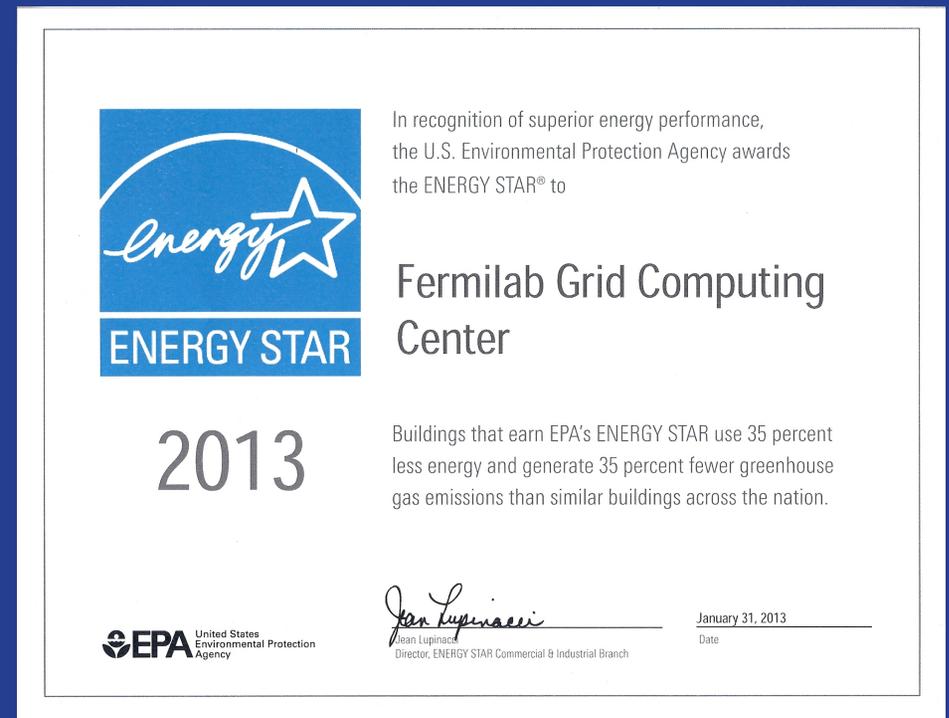
- GCC-A had 8 hours of scheduled downtime,
- GCC-B and GCC-C were impacted by "Load Shed" events the summer of 2012,
- GCC-B and GCC-C are expected to have additional "Load Shed" events for the summer of 2013.

LCC:

- LCC had 2.6 hours of downtime, corresponding to >99.9% uptime.

The GCC computer center received an energy star certificate for 2013.

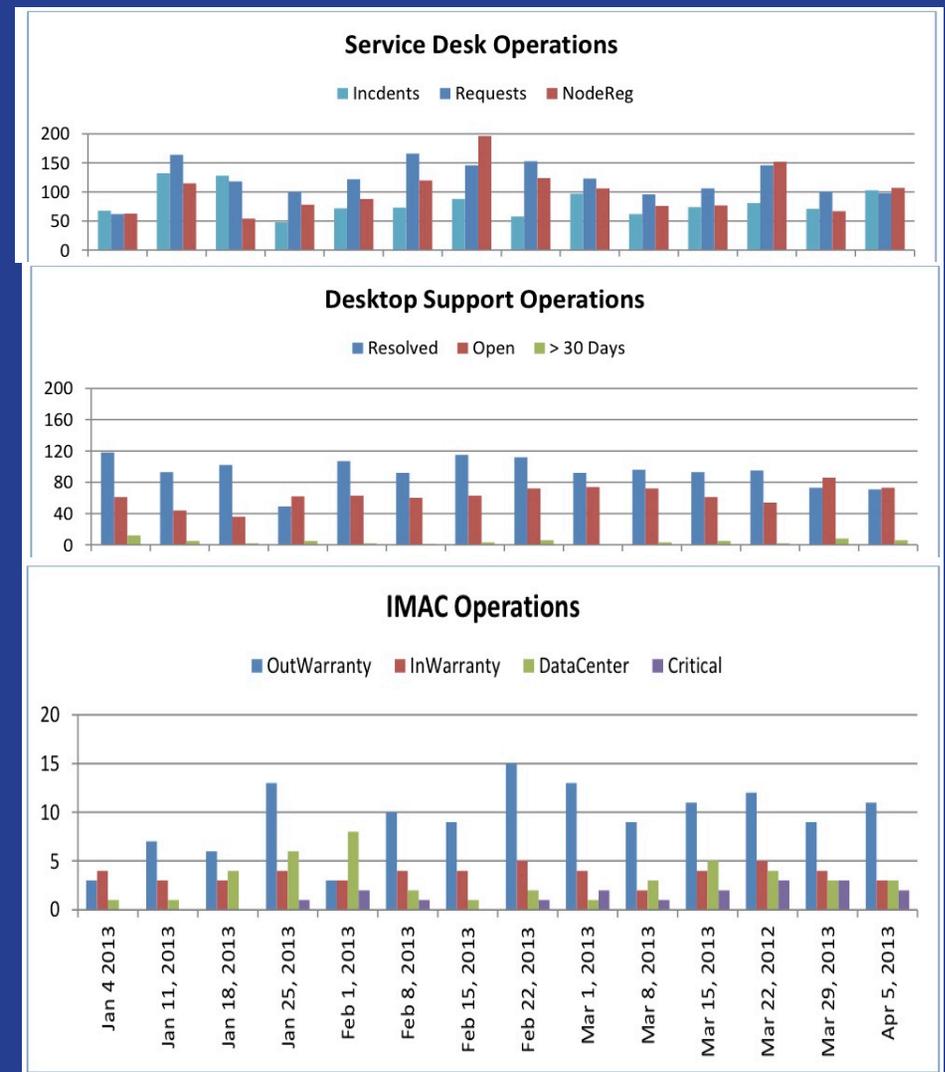
- This is the third year in a row that this building has received the award.



Dell Managed Services

Dell provides the following services:

- Service desk operations
- Desktop support (windows+apple)
- Printer support
- IMAC (Installs, Moves, Adds, Changes)
- Network infrastructure
- PREP+Logistics



Data Movement and Storage

Closing in on 70 PB of data on tape in the robots at Fermilab,

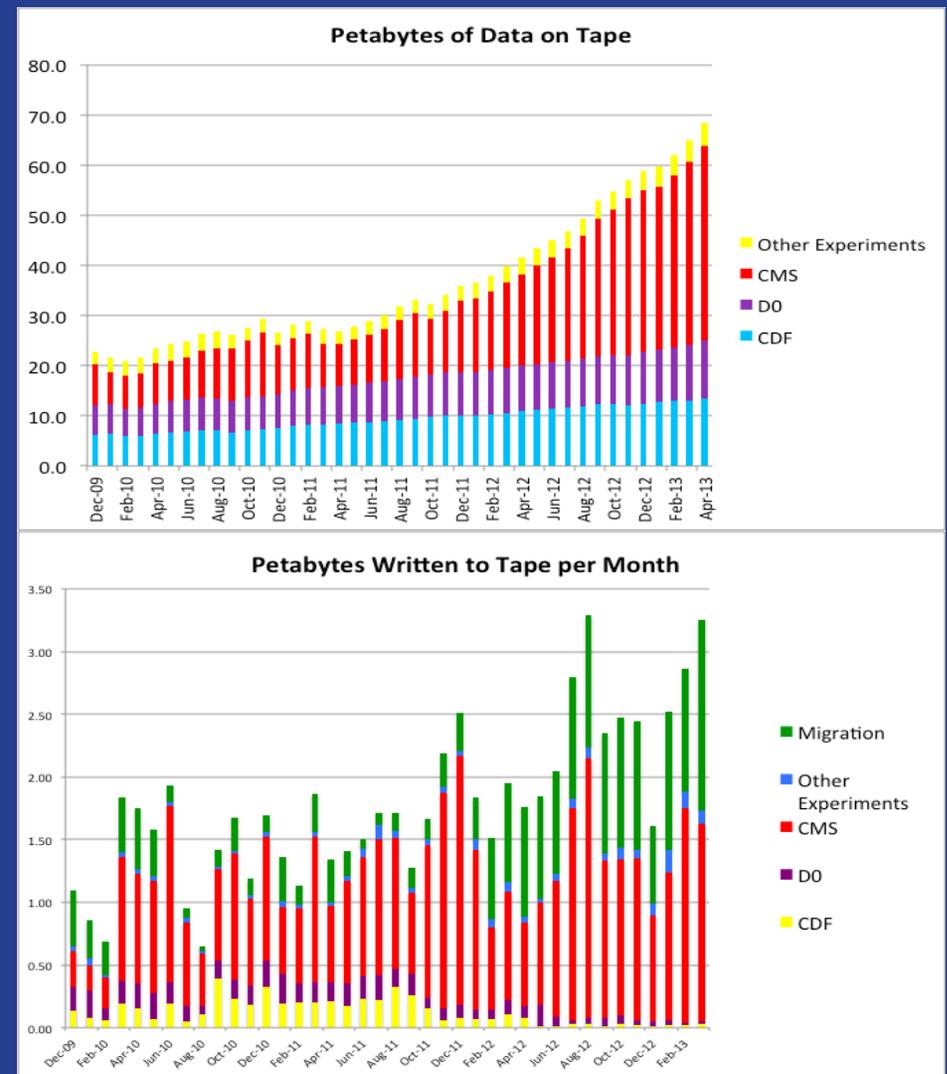
- CMS accounts for ~50% of total tape consumption.

We recently encountered excessive write errors on new tapes. The cause was tracked down to a batch of tapes that were contaminated with debris in the manufacturing process.

We worked with the vendors for several months to resolve the issues.

The resolution consisted of tape drive firmware improvements, tape manufacturing changes, and replacement of affected unused tapes.

We have now been running error free for a few weeks.



HPC and LQCD

Six HPC clusters in production,

Four large clusters dedicated to LQCD:

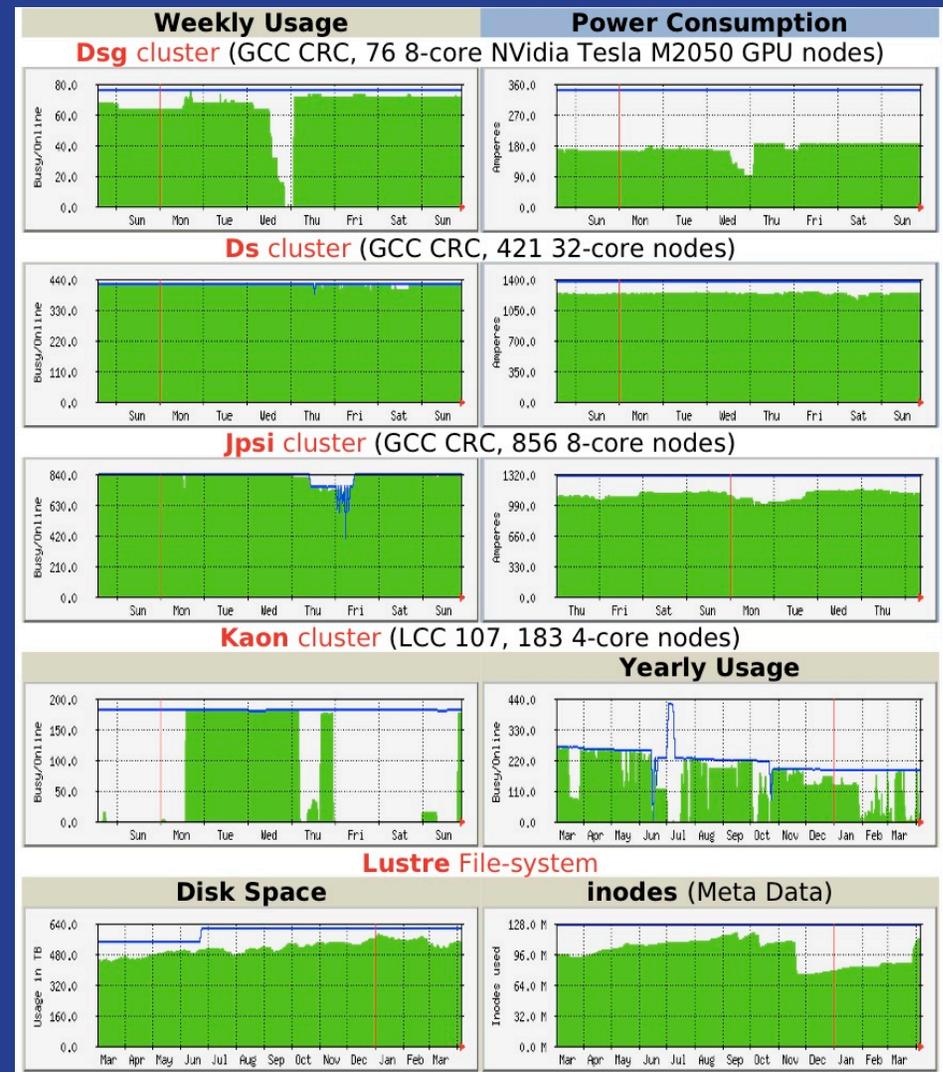
- Dsg (nvidia graphics)
- Ds
- Jpsi
- Kaon

Two smaller clusters for other HPC clients:

- Computational Cosmology
- Wilson

New HPC cluster based on Intel phi is in the process of acquisition – contract signed with vendor, waiting for delivery.

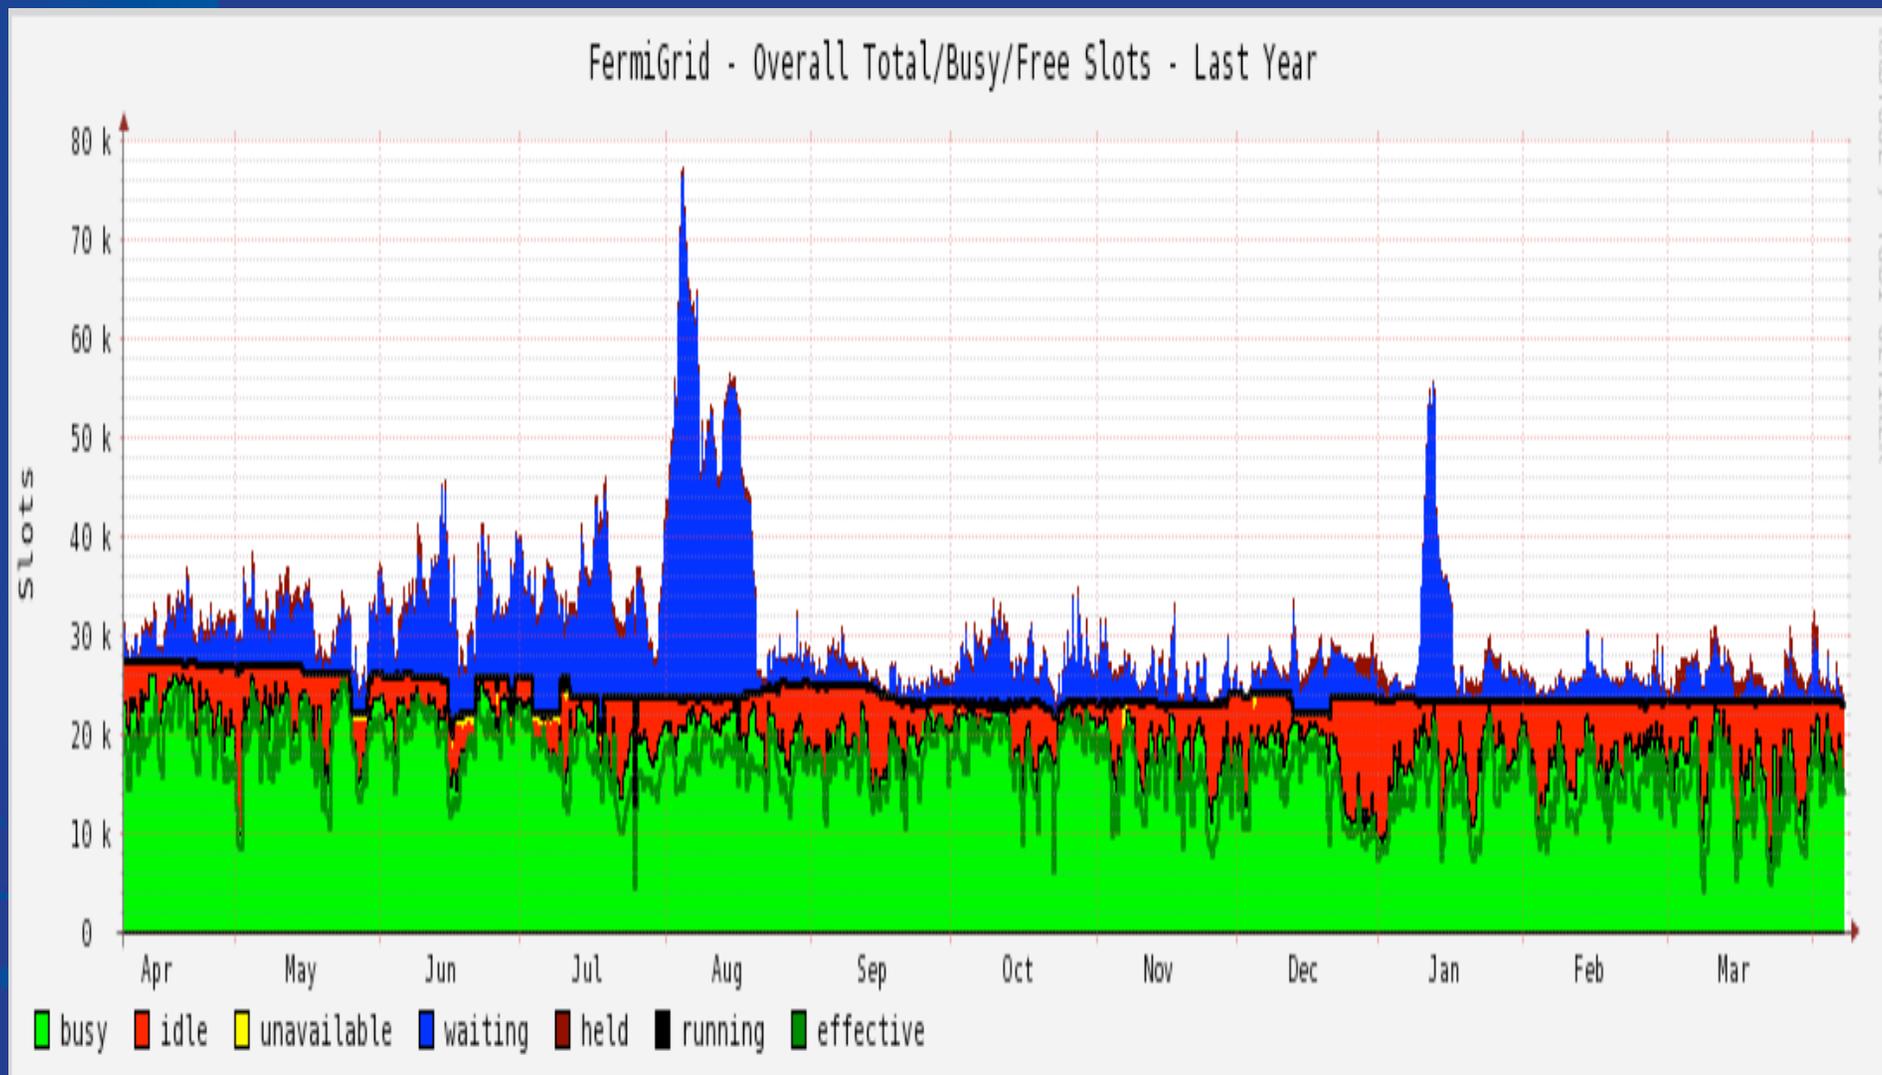
- More on this at the Fall 2013 HEPiX.



Current Fermilab Campus Grid Statistics (as of April 2013)

Cluster(s)	Batch System	Job Slots	Raw Occupancy	Effective Utilization
CDF (Merged)	Condor	5268	81.8	70.0
CMS T1	Condor	6,272	90.3	85.5
D0 (Merged)	PBS	5,920	74.2	53.7
GP Grid	Condor	5,474	77.3	61.8
_____		_____	_____	_____
Overall-Today		22,934	80.4	68.3
Last Year		27,184	89.4	75.8

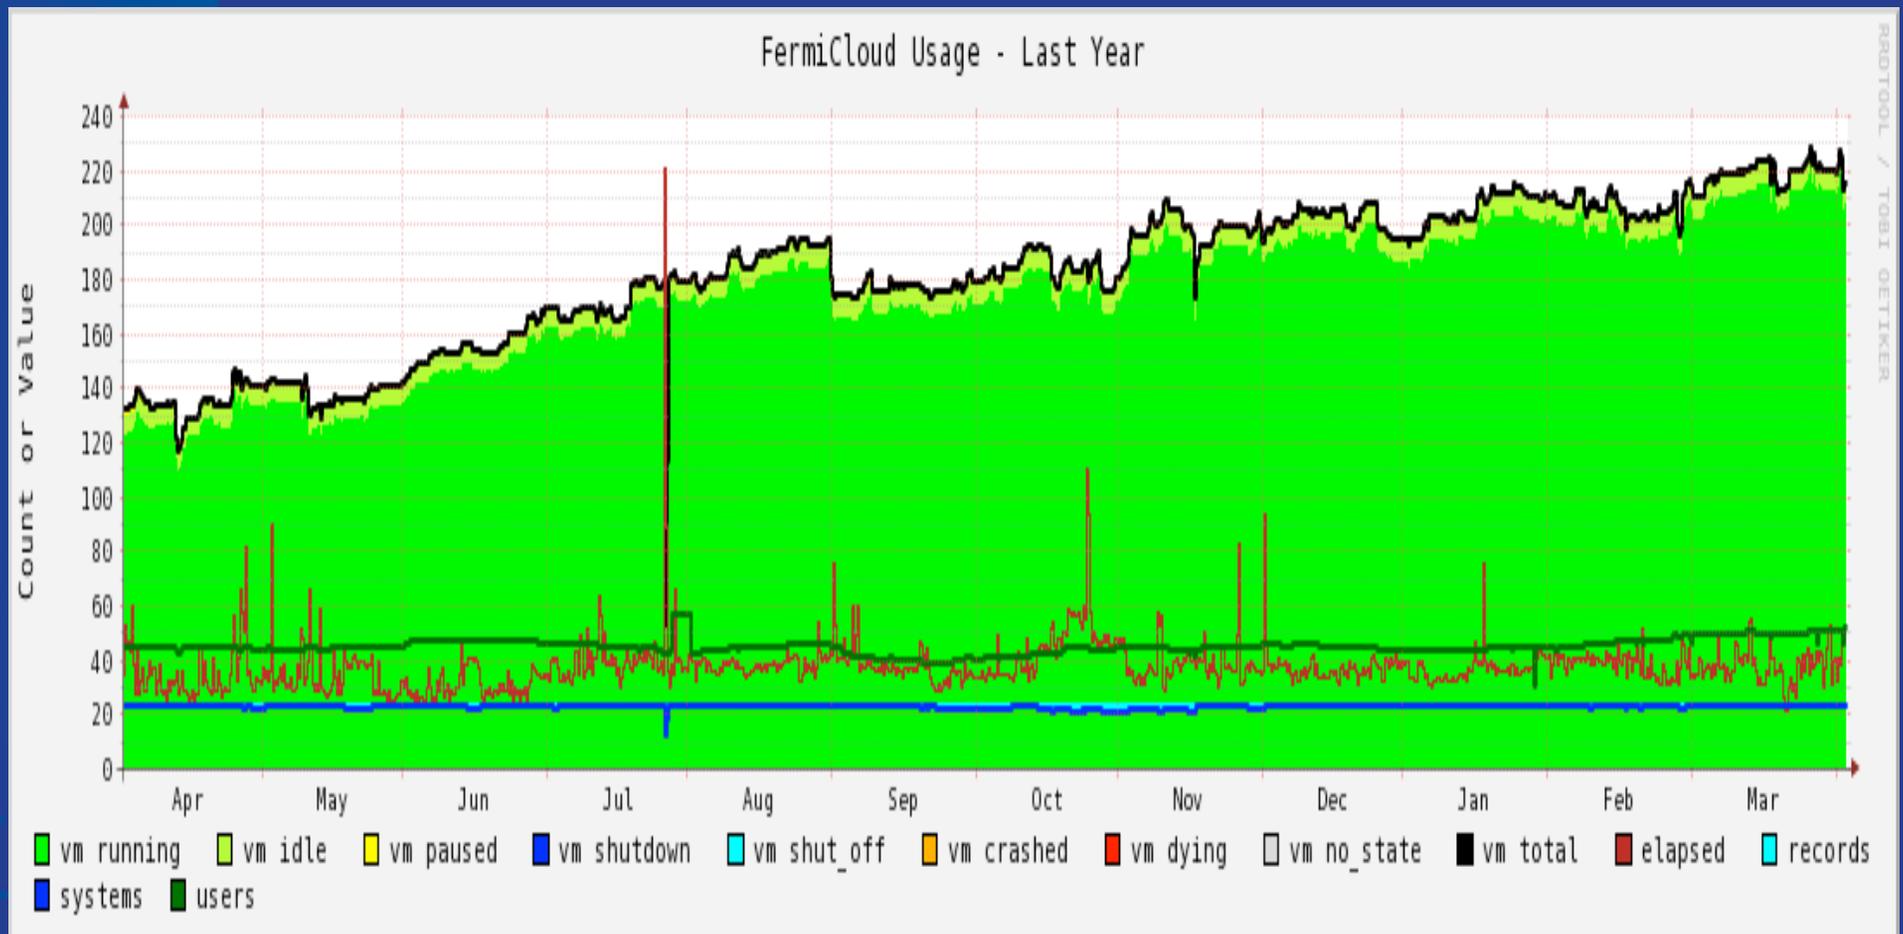
Fermilab Campus Grid



FermiGrid Service Availability (measured over the past year)

Service	Raw Availability	HA Configuration	Measured HA Availability	Minutes of Downtime
VOMS – VO Management Service	99.864	Active-Active	99.988%	60
GUMS – Grid User Mapping Service	99.863	Active-Active	100.000%	0
SAZ – Site AuthoriZation Service	99.863	Active-Active	100.000%	0
Squid – Web Cache	99.817	Active-Active	99.988%	60
MyProxy – Grid Proxy Service	99.781	Active-Standby	99.874%	660
ReSS – Resource Selection Service	99.915	Active-Active	99.988%	60
Gratia – Fermilab and OSG Accounting	99.229	Active-Standby	99.945%	300
MySQL Database	99.687	Active-Active	100.000%	0

FermiCloud Usage Continues to Grow

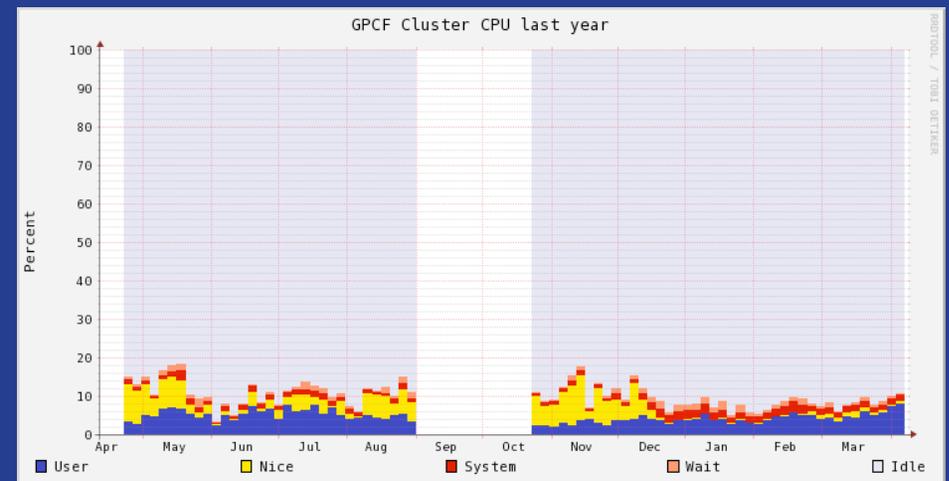
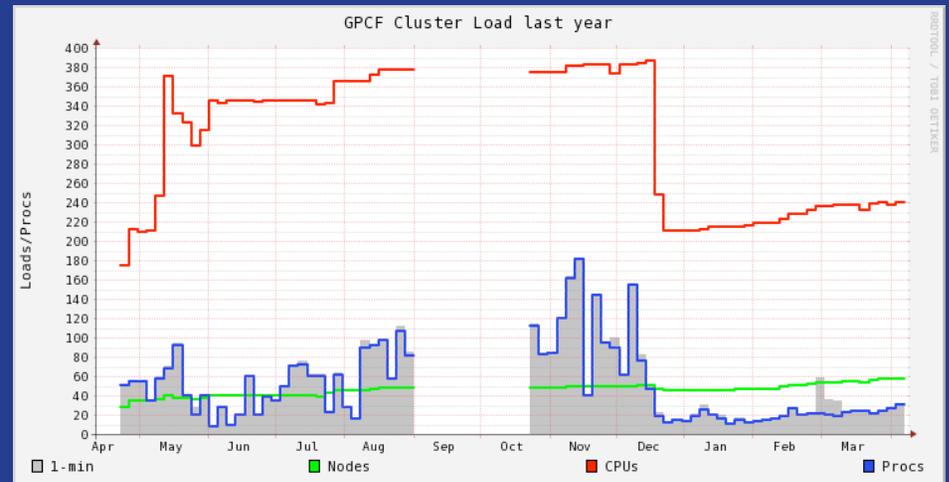


See my talk later in the week...

General Physics Compute Facility (GPCF)

GPCF provides statically deployed virtual machines:

- GPCF is a critical component of the Intensity and Cosmic Frontier experiments,
- As well as others who need similar services,
- It continues to operate well.



VMware Virtual Services

Primary focus is the support of Fermilab core computing division services,

Authentication

Fermilab has commissioned an internal task force to recommend changes to the currently deployed Fermilab authentication infrastructure:

- MIT Kerberos + Kerberos Certificate Authority
- Windows "Fermi" Active Directory Domain
- Windows "Services" Active Directory Domain (used for email, service-now, etc.).

Thank you to those individuals that responded to my request for information about authentication infrastructures that are deployed at the various HEPiX institutions!

I expect to be able to give a formal presentation on Fermilab's plan for the future at the Fall 2013 HEPiX in Ann Arbor.

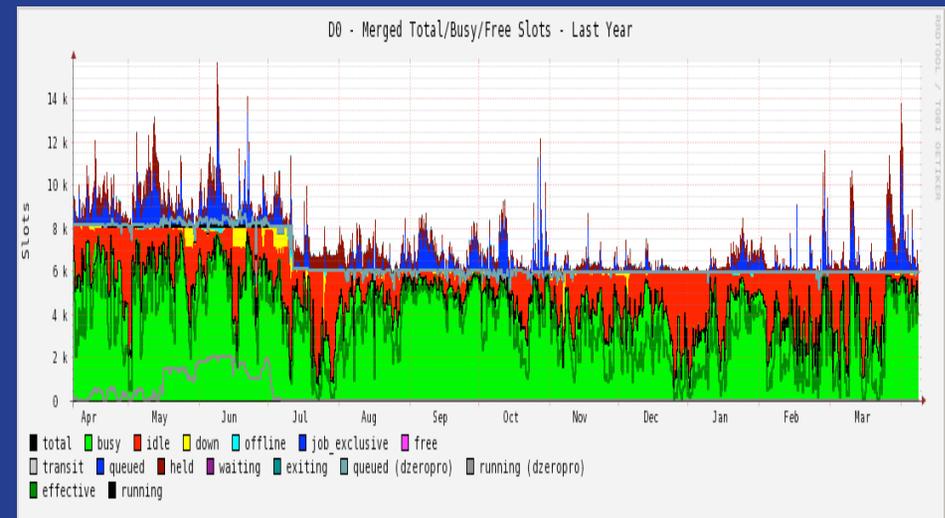
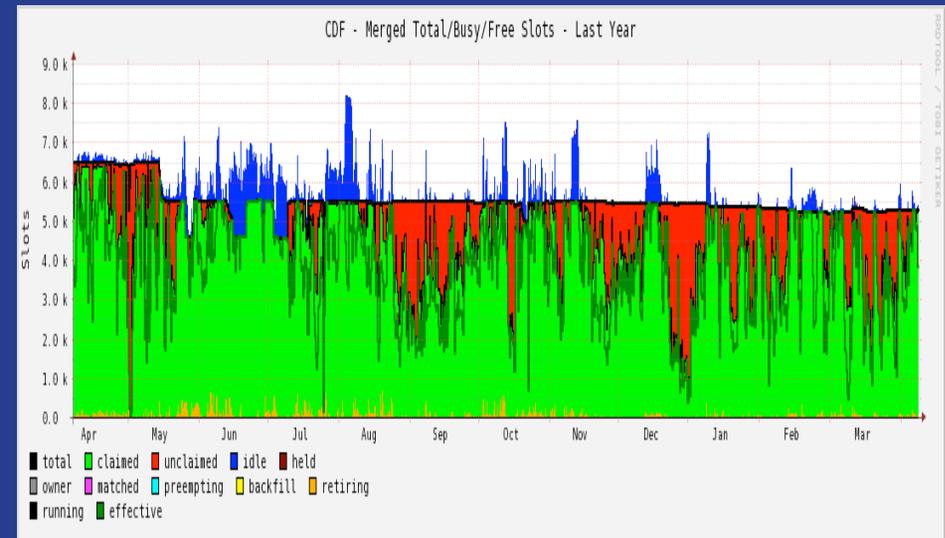
Energy Frontier – CDF & D0

CDF and D0 continue to perform analysis of their datasets,

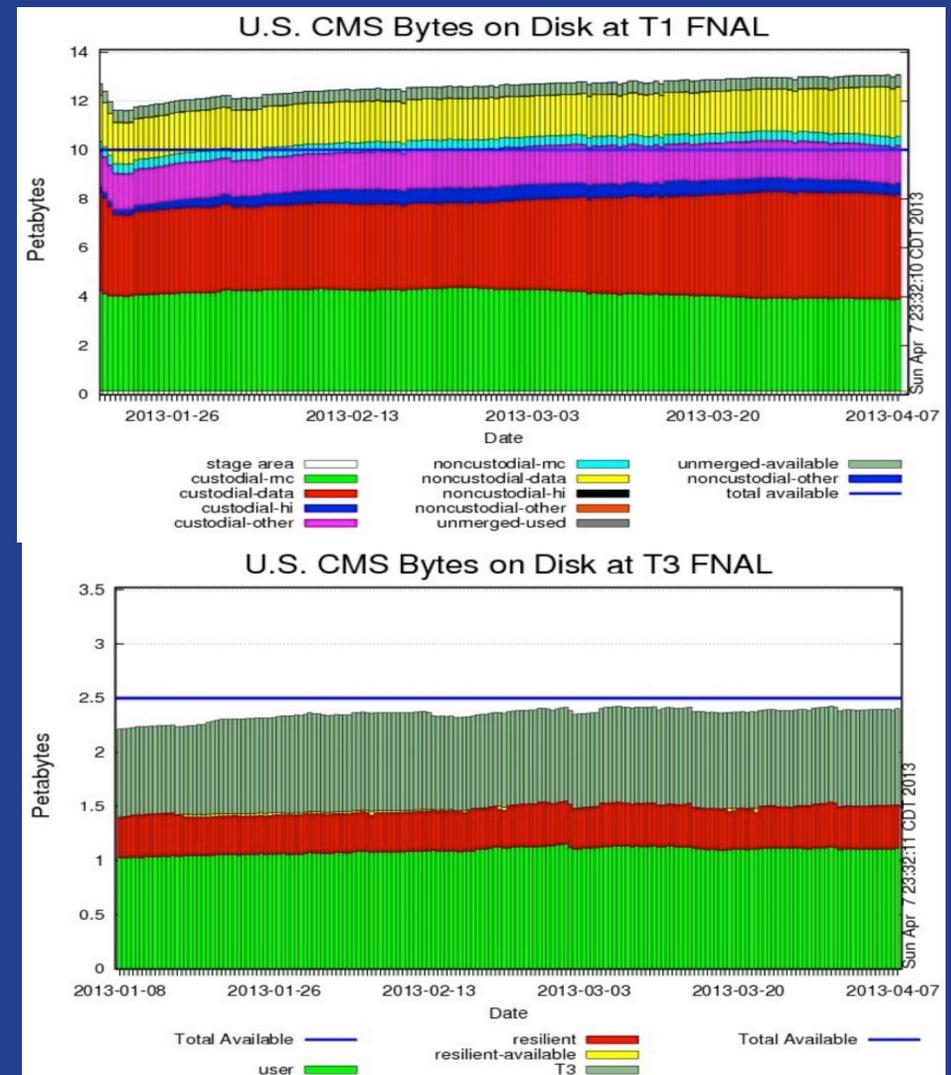
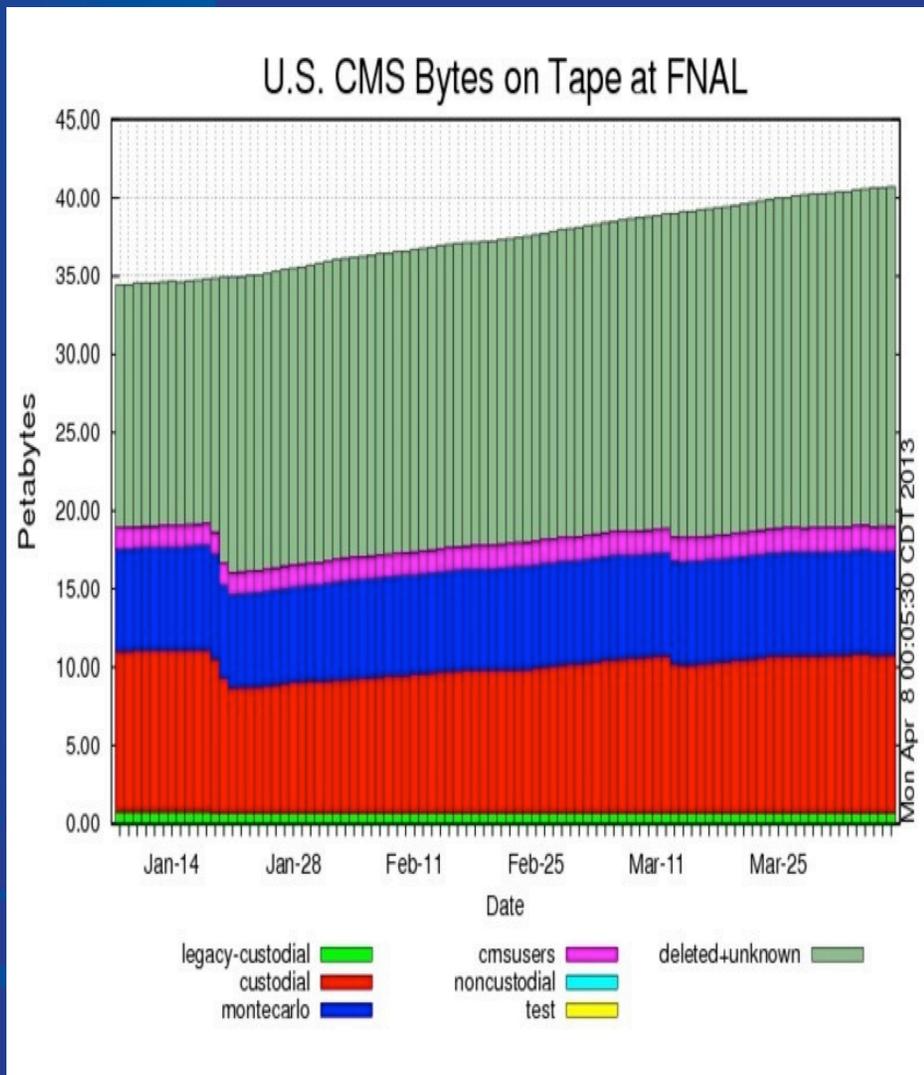
- CDF results on Exotic, Top, Higgs, Bottom, and QCD physics,
- D0 results on Bottom, Electroweak, New Phenomena, Higgs, QCD, and Top physics

The Run II Data Preservation efforts are gathering steam,

- An implementation using a combination of virtualization and cloud computing looks likely...

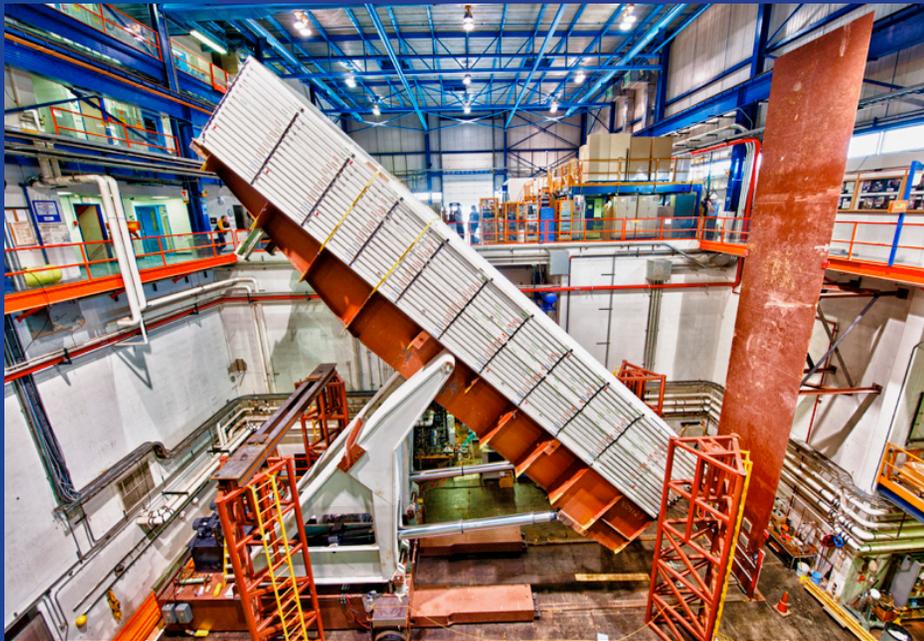


Energy Frontier – CMS



Intensity Frontier NOvA Block Pivoter and First Block

Block Pivoter:



Block pivoter animation:

- <http://www.youtube.com/watch?v=YtuMqjCiyMQ>

First block of the NOvA detector on the block pivoter painted black (to prevent light from entering the detector):

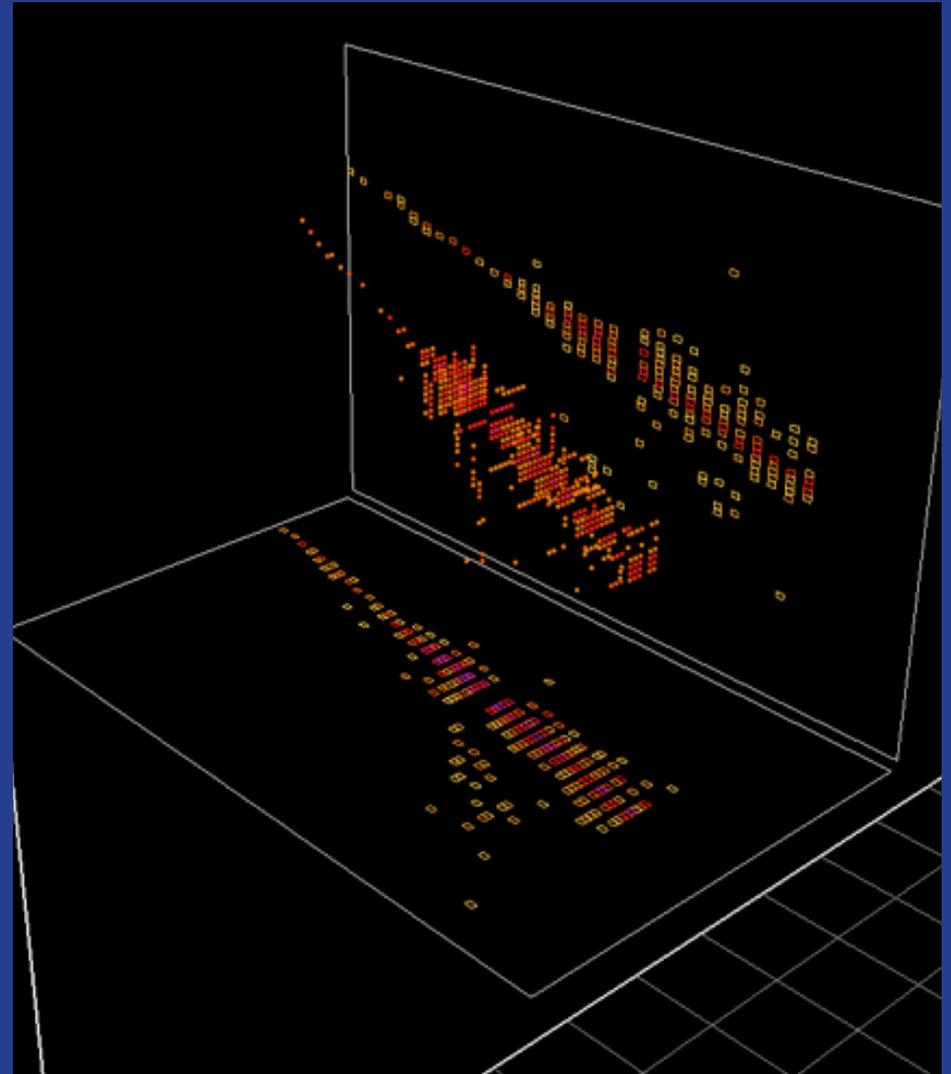


Intensity Frontier – NOvA

The NOvA neutrino detector under construction in Ash River, Minnesota has recorded the first 3-D particle tracks from cosmic rays – particles that are produced by a constant rain of atomic nuclei falling on the Earth's atmosphere from space
(see image =>)

The currently active section of the detector is about 12 feet long, 15 feet wide and 20 feet tall,

The full detector will measure more than 200 feet long, 50 feet wide and 50 feet tall.



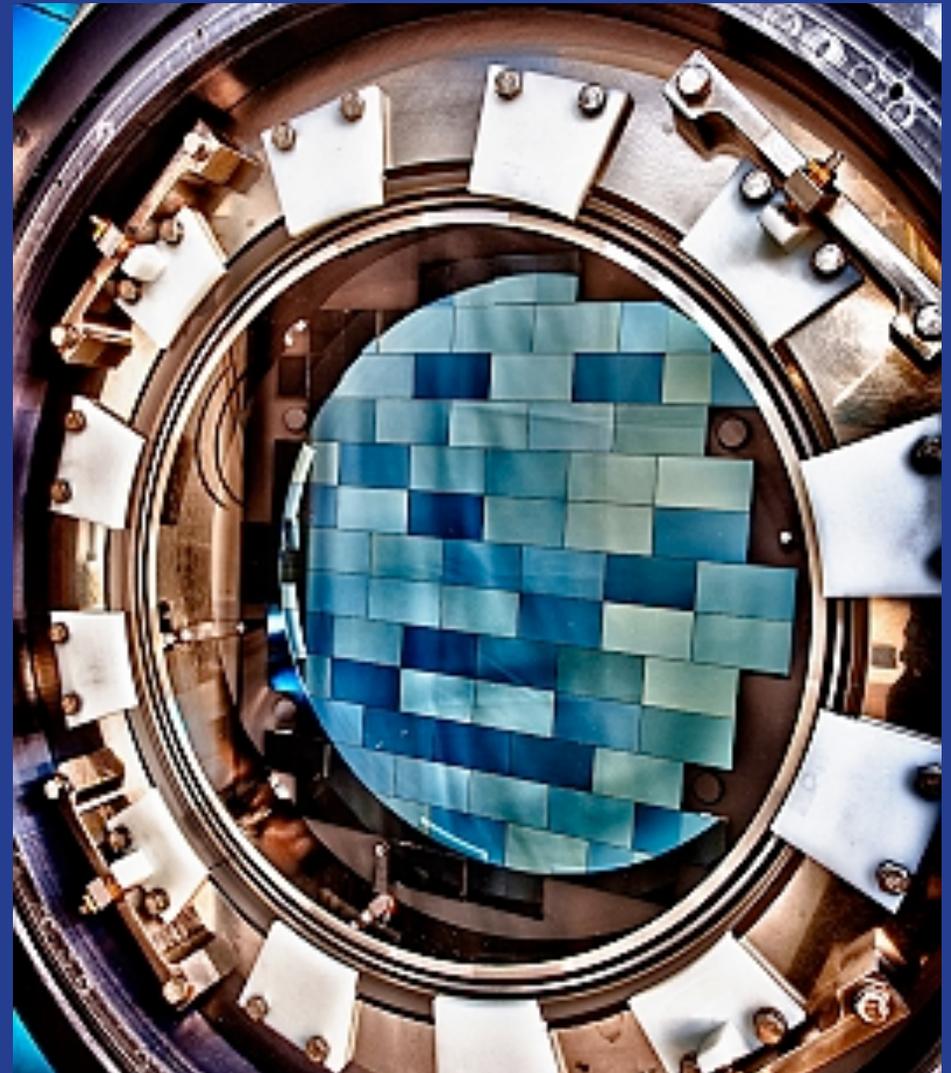
Cosmic Frontier – DES

Dark Energy Survey:

- The 570-megapixel Dark Energy Camera (DECam) is the world's most powerful digital imaging device, built at Fermilab and installed on the Blanco 4-meter telescope at the Cerro Tololo Inter-American Observatory in Chile.
- DECam was constructed for the Dark Energy Survey, a five-year effort to map a portion of the southern sky in unprecedented detail.
- First light of DECam was November 2012, and since that time, the DES collaboration has spent 50 nights completing the science verification phase of the experiment.
- Analysis pipelines are being “shaken down” using FermiGrid, OSG and XCEDE resources,
- The DES camera is getting ready for science,

More information:

- <https://www.darkenergysurvey.org/>



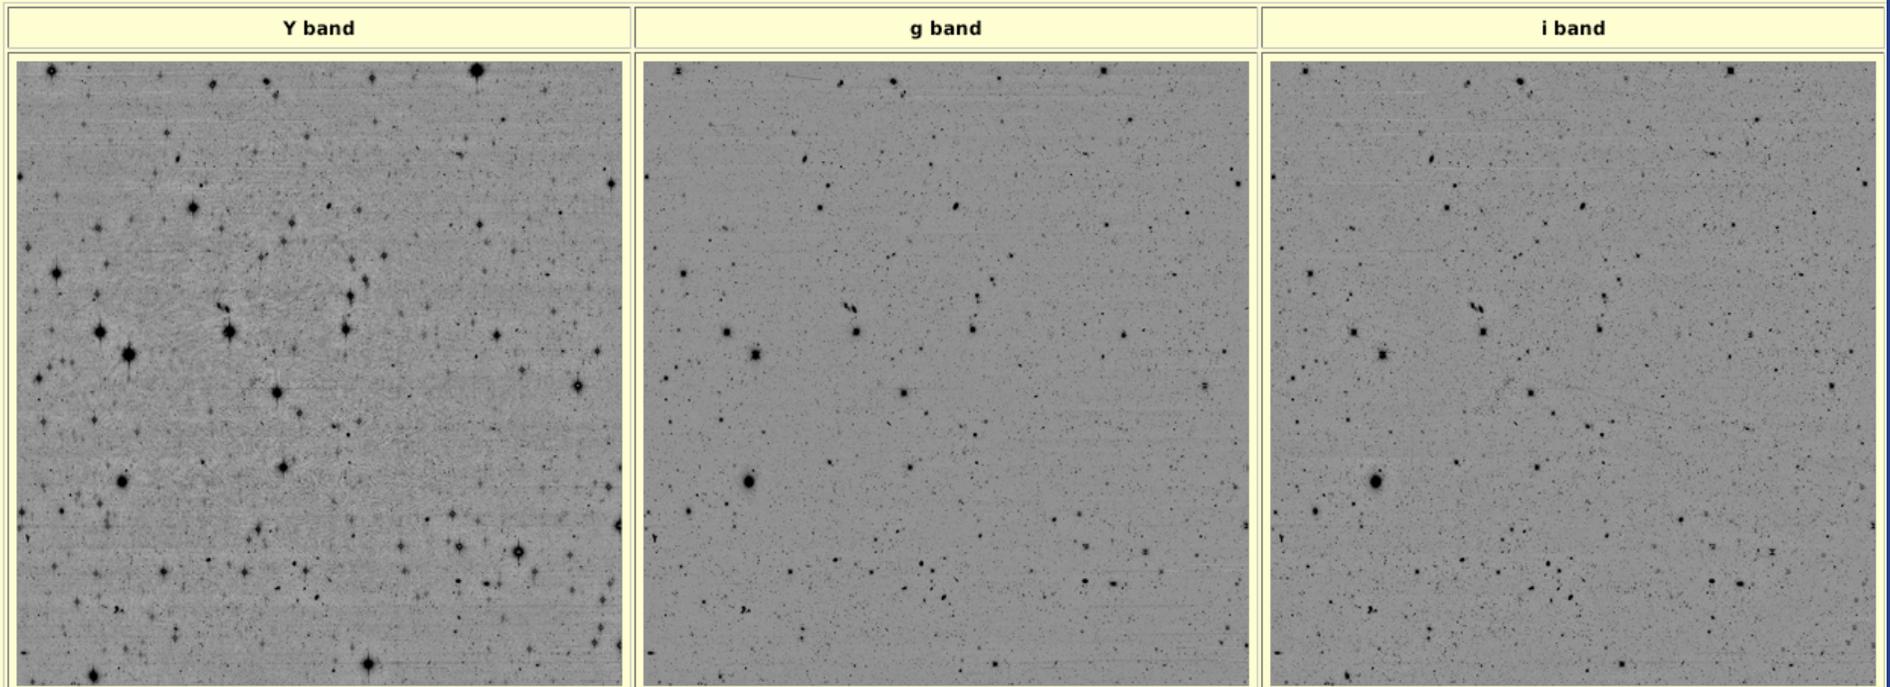
Blanco Telescope



DECam Images

Coadd Images for tile DES0102-4914

Run: 20130227000003 DES0102-4914
Source nites: 20121124, 20121207
Source runs: 20130225200732_20121124, 20130225212202_20121207
Bands: Y, g, i, r, z



Illinois Accelerator Research Center



Summary

It is an exciting time to be supporting science at Fermilab in all of the three frontiers of science:

- Energy Frontier
- Intensity Frontier
- Cosmic Frontier

And there is no danger of running out of work supporting and extending the computer systems and services that provide the foundations of supporting science.

Thank You!

Any Questions?