Cluster Consolidation at NERSC

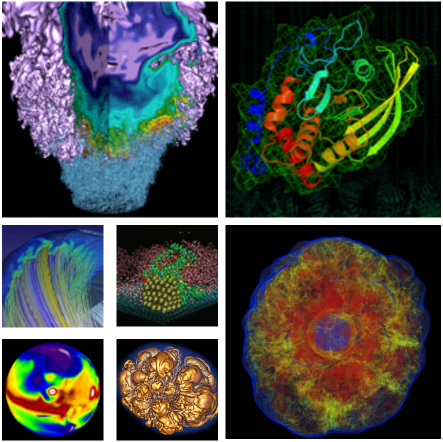

NERSC 40 YEARS at the FOREFRONT 1974-2014

Larry Pezzaglia
NERSC Computational Systems Group
lmpezzaglia@lbl.gov

HEPiX Spring 2014

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

Snapshot of NERSC

- Located at Lawrence Berkeley National Laboratory, NERSC is the production computing facility for the US DOE Office of Science
- NERSC serves ~5000 users, ~400 projects, and ~500 codes
  - Focus is on "unique" resources:
    - Expert computing and other services
    - 24x7 monitoring
    - High-end computing and storage systems
  - Known for:
    - Excellent services and user support
    - Diverse workload
  - NERSC provides Hopper (a Cray XE6), Edison (a Cray XC30), and three data-intensive systems: Carver, PDSF, and Genepool.

The NERSC Cluster Model

Cluster Expansion

- In 2012, NERSC purchased a new system, “Mendel” to systematically expand its cluster resources
  - 500+ Sandy Bridge nodes, 8000+ cores
  - FDR InfiniBand interconnect
- Mendel transparently expands production clusters and services
  - Carver, PDSF, and Genepool (the “parent systems”) schedule jobs on portions of Mendel
  - Mendel provides multiple software environments to match those on each parent system
- This model was presented at the 2013 Cray User Group meeting
  - http://cug.org/proceedings/cug2013_proceedings/includes/files/pap184-file1.pdf
  - http://cug.org/proceedings/cug2013_proceedings/includes/files/pap184-file2.pdf

Data-Intensive Systems

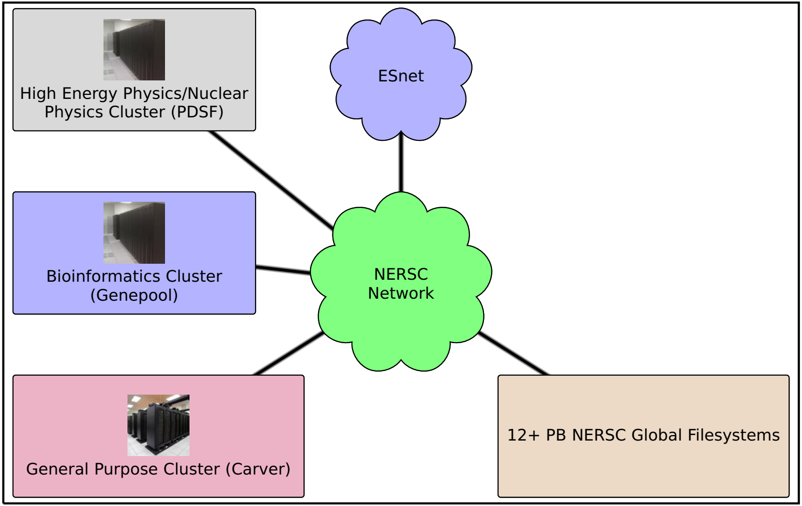

Data-Intensive Systems

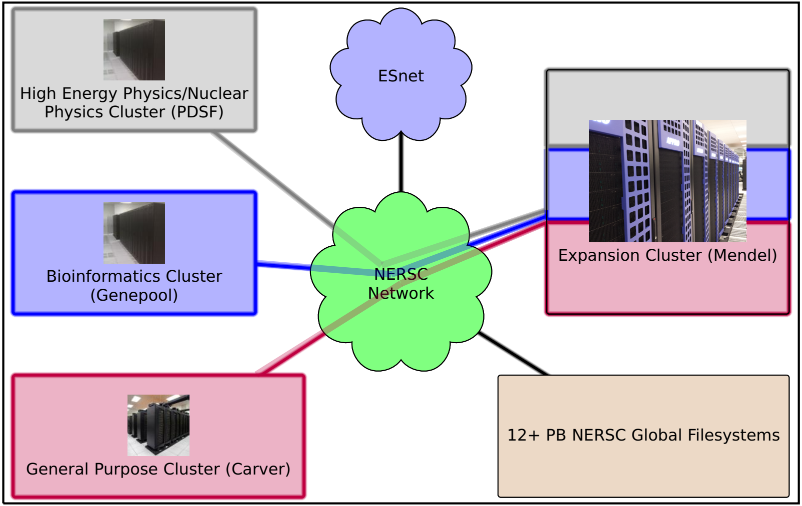

The Mendel Approach

- We use tools to construct convenient management abstractions and tuned user environments on top of this platform:
  - Familiar open-source software:
    - **xCAT** to provision and manage nodes
    - **Cfengine3** to provide configuration management (versioned with SVN)
  - NERSC-developed BSD-licensed software:
    - **avs_image_mgr** to handle xCAT image management and versioning

      http://github.com/lpezzaglia/avs_image_mgr
    - **CHOS** to provide multiple compute environments concurrently and seamlessly

      http://github.com/scanon/chos
    - **minimond** to collect trending data for troubleshooting and analysis

      http://github.com/lpezzaglia/minimond

NERSC Cluster Model

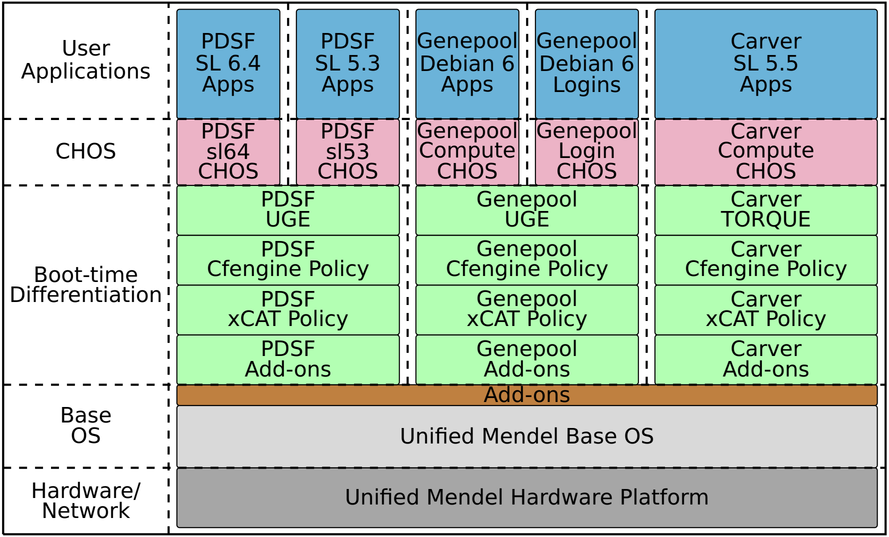

Extending the Model

Motivation for Consolidation

- Easy Mendel administration highlighted the operational burden of managing legacy clusters
  - Changing configurations with pdsh scales poorly
  - Mendel demonstrated the value of leveraging automation to manage complex systems
- We pursued further cluster consolidation efforts
  - Staff efficiency is highly valued
  - Legacy hardware also gains benefits through Mendel membership
  - Cost: Upfront effort to consolidate clusters and risk of user disruption
  - Reward: Reduced long-term sysadmin burden and increased system consistency

Extending the Model

- In spring 2014, we merged "Genepool", a legacy parent cluster, into Mendel's management system
- The combined cluster is now managed as a single integrated system with:
  - ~1000 nodes
  - Multi-generational, multi-vendor hardware
  - Multiple separate interconnects
  - A unified xCAT+Cfengine management interface
- Constrained by a 24x7, disruption-sensitive environment
- Change activated in a single all-day maintenance

Consolidated Cluster Model

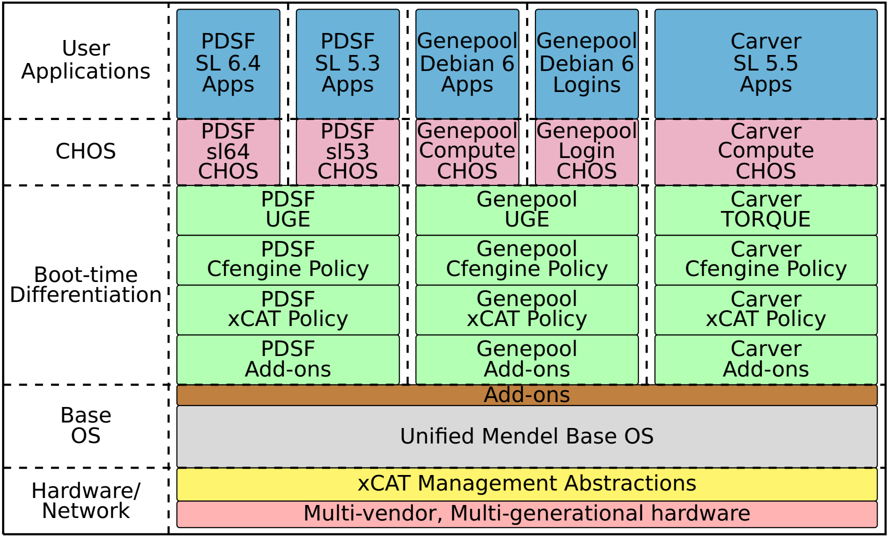

Specific challenges

Genepool and Mendel differ in several respects

	Mendel (New)	**Genepool (Legacy)**
Production interconnect	FDR InfiniBand	Gigabit Ethernet
Provisioning/IPMI network	Dedicated GbE	Dedicated IPMI network Provisions over production network
OS	SL 6.3 base OS with CHOS	Debian 6 without CHOS
Hardware	Homogeneous platform	Many hardware configurations

Approach

- xCAT’s hierarchical management features are suited to managing dissimilar hardware
- Expansion required changes to our software stack
  1. Expand hardware support in the base OS image
  2. Configure an xCAT service node
  3. Expand Cfengine rules and boot scripts
  4. Perform thorough testing
  5. Reboot the Genepool nodes through Mendel’s management system

Base OS modifications

- Support for all Genepool hardware was added to the base OS image.
  - Kernel modules for disk and network controllers
  - Initrd code to handle Genepool network characteristics
- An xCAT add-on was created for the xCAT service node
- Changes were made with **avs_image_mgr**
  - Provides a full revision history of every file
  - Provides the ability to roll back to any previous image

xCAT modifications

- An xCAT Service Node (SN) handles Genepool provisioning/management under the direction of the Management Node (MN)
- Only the SN requires connectivity to the Genepool networks
  - The MN only requires connectivity to the SN
  - The SN provides DHCP/TFTP/HTTP/xCATd services
  - xCAT commands, such as power/console operations, are transparently routed through the SN
- The SN is provisioned through the Mendel cluster model
  1. The Mendel base OS image is booted
  2. The xCAT add-on is activated
  3. Cfengine rules apply SN-specific configurations

Postscripts and Cfengine

- Postscripts were extended to support Genepool characteristics
  - GPFS cluster configuration
  - Multipath access to local disk arrays
  - Local filesystem configurations
  - Multiple hardware configurations
- Cfengine rules were augmented to support the additional node classes

The Combined Cluster

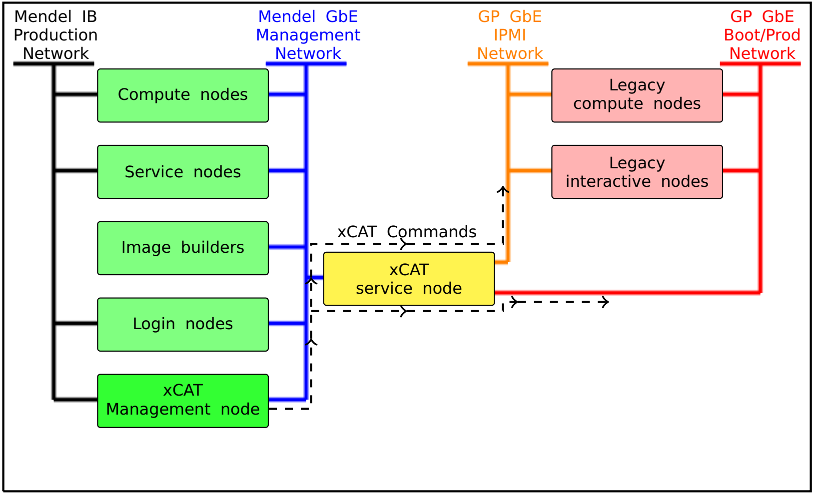

Cluster Automation

- The combined Mendel+Genepool system is complex
  - Many different node classes
  - Each node class represents a unique software/hardware combination
- Configuration complexity grows with the number of node classes.

Cluster Automation

- The quantity of node classes exceeds what a human administrator can hold in immediate memory
  - We must build abstractions to retain system manageability as complexity increases
  - Configuration management has become a necessity
- We manage a *single integrated system*, not a collection of nodes.
- Every change must be considered in a system-wide context
- Cfengine must broker change rollout

Development Updates

CHOS development

- CHOS enables concurrent support of multiple Linux environments on a single server
- A core component of the Mendel cluster model
- Under active development
- Recent changes include:
  - Ability to exit CHOS from within a CHOS environment
  - Build system improvements
  - **pam_chos** configurability enhancements
  - EL7 kernel support in a testing branch
- Planned features include:
  - Scripts to transform an installed EL system into a CHOS environment
  - A framework for user-supplied CHOS environments
  - Reduced kernel module scope

Data Collection

- Collecting trending data for historical analysis is growing increasingly important
- NERSC developed **minimond** to systematize this process
- Collects ~1000 statistics per node
- Modular framework for sending metrics to multiple data aggregation services
  - Supported output methods: plain text and Ganglia (via gmetric or EmbeddedGmetric)
  - AMQP support is planned
- Only absolute counter values are recorded. Calculation of derived statistics must be performed on a remote analysis server

Metrics Graphs

GPFS mmpmon Operations

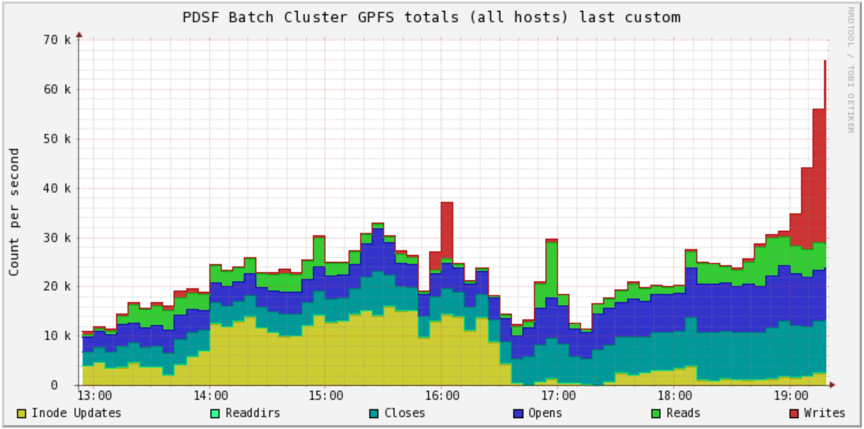

Metrics Graphs

GPFS mmpmon Throughput

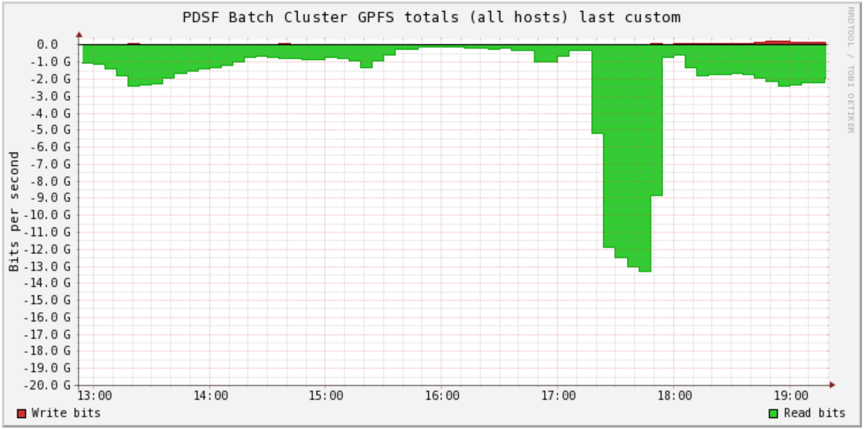

Metrics Graphs

IB gateway Throughput

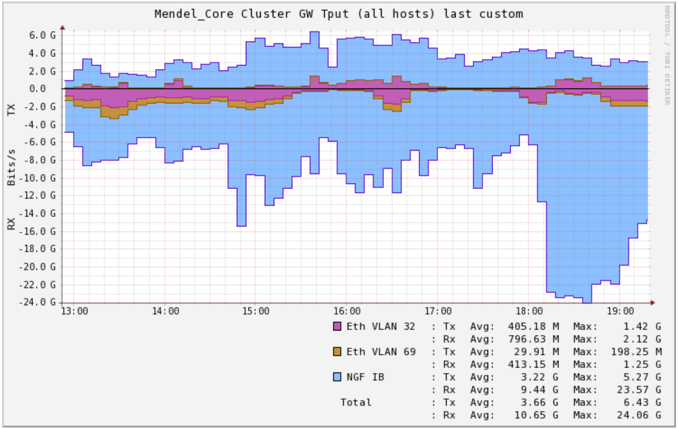

Conclusions

- The extended NERSC cluster model enables systematic management of several multi-vendor, multi-interconnect, and multi-generational clusters as a single integrated system
- A unified management interface abstracts away complex details
- Implementation involved minimal user disruption
- Extending the model was far easier than separately managing both clusters
- The new model dramatically simplifies operations

Acknowledgements

- This work was supported by the Director, Office of Science, Office of Advanced Scientific Computing Research of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

National Energy Research Scientific Computing Center