JLab Scientific Computing: Theory HPC & Experimental Physics

Thomas Jefferson National Accelerator Facility
Newport News, VA
www.jlab.org

Sandy Philpott
HEPiX Annecy – May 2014

Outline

Accelerator Upgraded!

Resources at JLab

- compute: CPU & Accelerated (GPU, MIC)
- disk, tape, and networking

Banking Core Hours / Load Balancing

- share resources between theory and experimental physics

Lustre Upgrade & Partitioning

Plans for next 2+ years

6 -> 12 GeV Accelerator Upgrade

- Upgraded accelerator has delivered beam!
  - With one new experimental Hall D - GlueX experiment
  - And will double the 6 GeV data rates in existing Halls A,B,C

www.jlab.org/news/releases/beam-target-cebaf-accelerator-achieves-12-gev-commissioning-milestone

- Initial external IT/Computing review in June 2012, reviewed again in November 2013
  - Ensure readiness of data acquisition and analysis on day 1
  - Both reviews ended with positive results – we're on target!
- Virtual field trip for high schoolers

https://www.youtube.com/watch?v=JyiBe5g17eI

Jefferson Lab Resources

Conventional Clusters

- **12s:** 276 nodes 4416 cores
  - 16 cores, 32 GB per node, Sandy Bridge CPU
  - full fat tree QDR fabric with no oversubscription
- **9q, 10q** 544 nodes 4352 cores
  - 8 cores, 24 GB per node , Nehalem and Westmere CPU
  - partitioned into 13 * 32 node and 1 * 128 node partitions, no oversubscription
- **farm**: 120 nodes, 1400 cores
  - Mix of 8 and 16 core nodes with recycled SDR and DDR
  - Single node jobs only

Jefferson Lab Resources

Accelerated Clusters

- **12k:** 42 nodes, 168 GPUs, 4.7 GB / gpu
  - host: 16 cores, 128 GB per node, full FDR IB
  -
- **9g, 10g,11g:** 114 nodes, 456 GPUs, 1.5-2.7 GB / gpu
  - host: 8 cores, 48 GB per node
  - 70% in 3 generations of gaming cards:
    - GTX-285, GTX-480 / GTX-580, GTX-690
  - 30% Tesla 2050, “Fermi” GPU, 2.7 GB / gpu
  - Infiniband: mix of ½ SDR (i.e. in half bandwidth slot), ½ DDR, ½ QDR, full QDR for GTX-690 (dual GPU cards, leaving full bandwidth slots open)
  - GTX-285’s will be retired June 30, GTX-480’s at most a year later

- **12m**: 16 nodes, Xeon Phi for R&D
  - Host: 16 cores, 128 GB per node, full FDR IB

Scientific Computing Infiniband Fabric

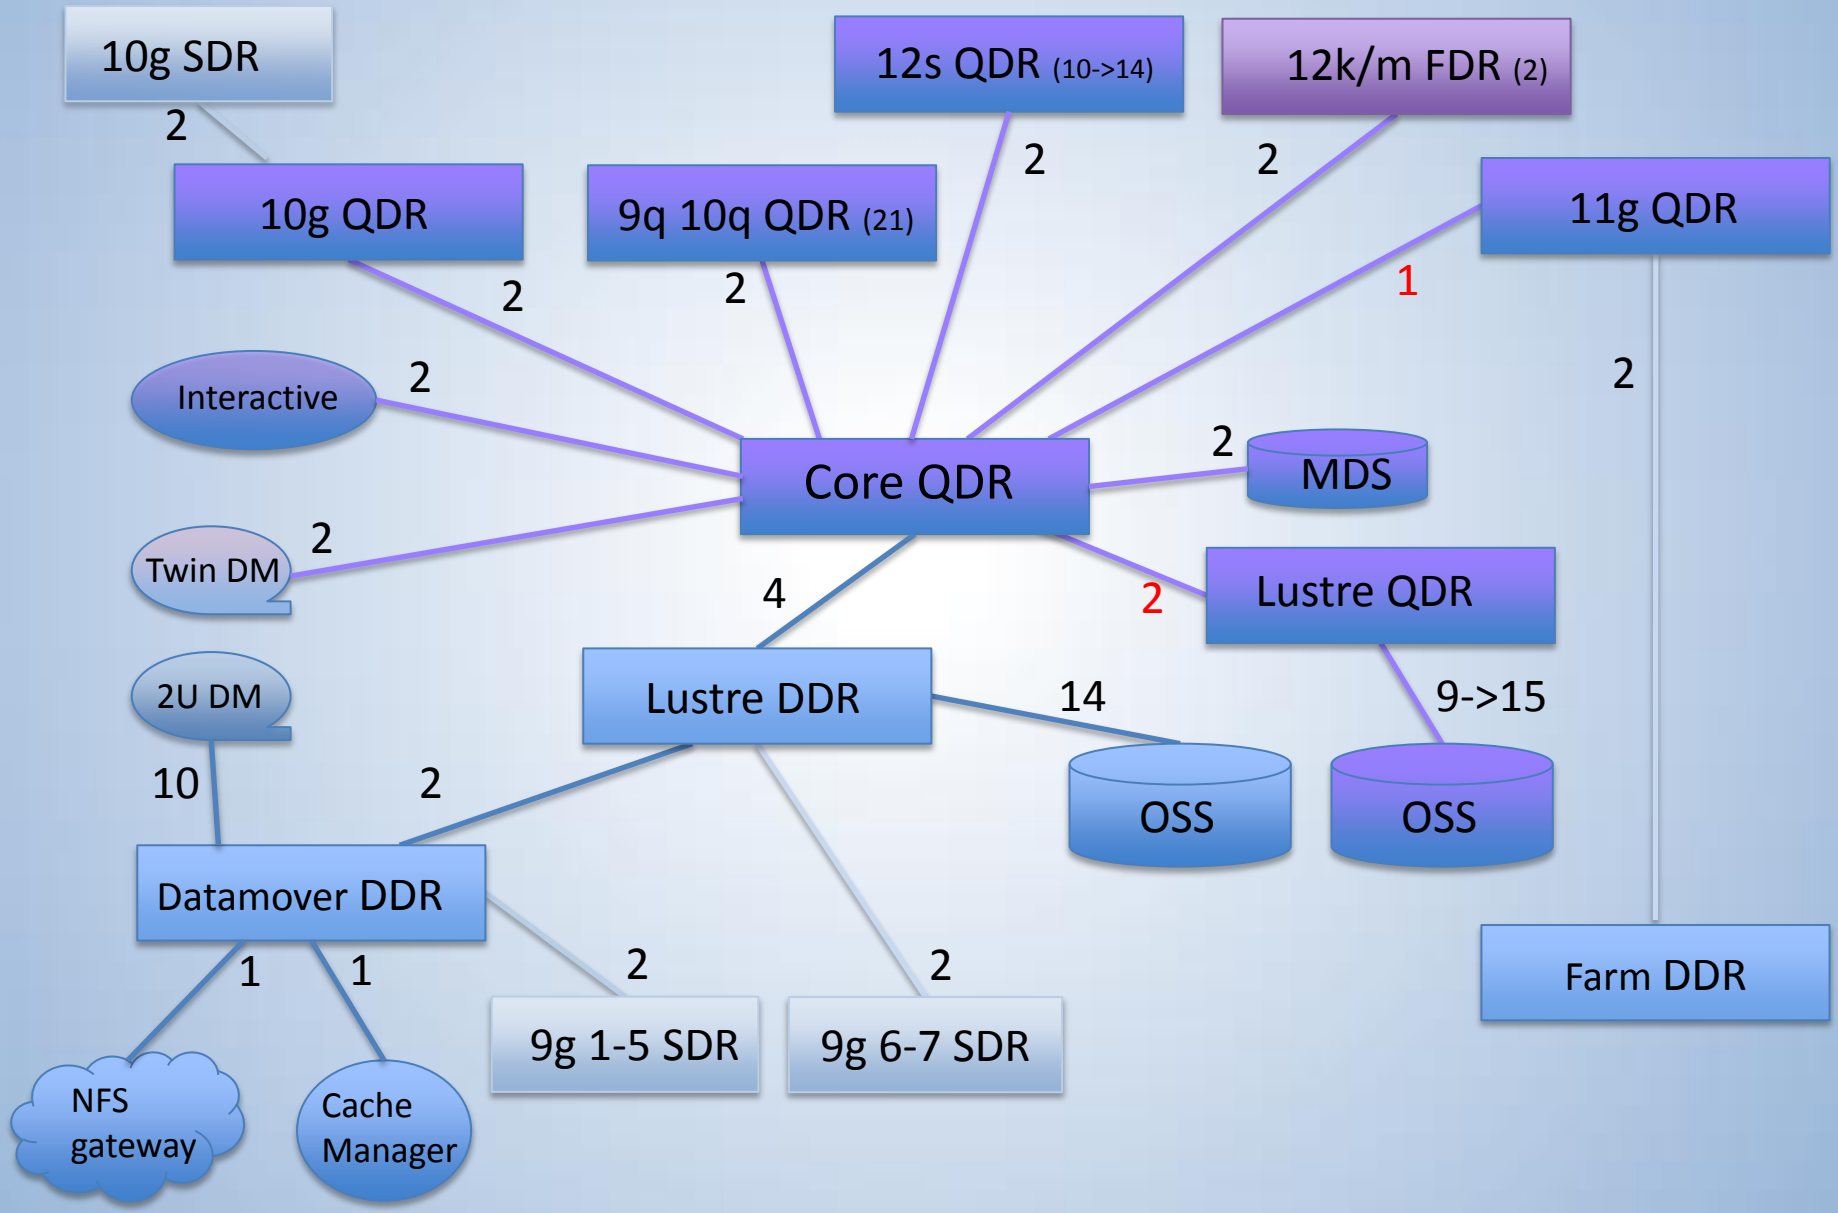

Resource Sharing (banking hours)

Farm and HPC clusters have similar nodes – easy to share

→ Move capacity between LQCD and Experimental Physics

Goals:

- Never waste core hours, even when users are otherwise busy
- Support Hall D GlueX & Hall D CLAS12 data challenges (using LQCD cores)

Experimental Physics loaned LQCD 512 cores for several months, banking ~1M core hour credit …

then borrowed 2048 cores for their 3 week long data challenge to exhaust the credit

Theory Infiniband Clusters: 8800 cores

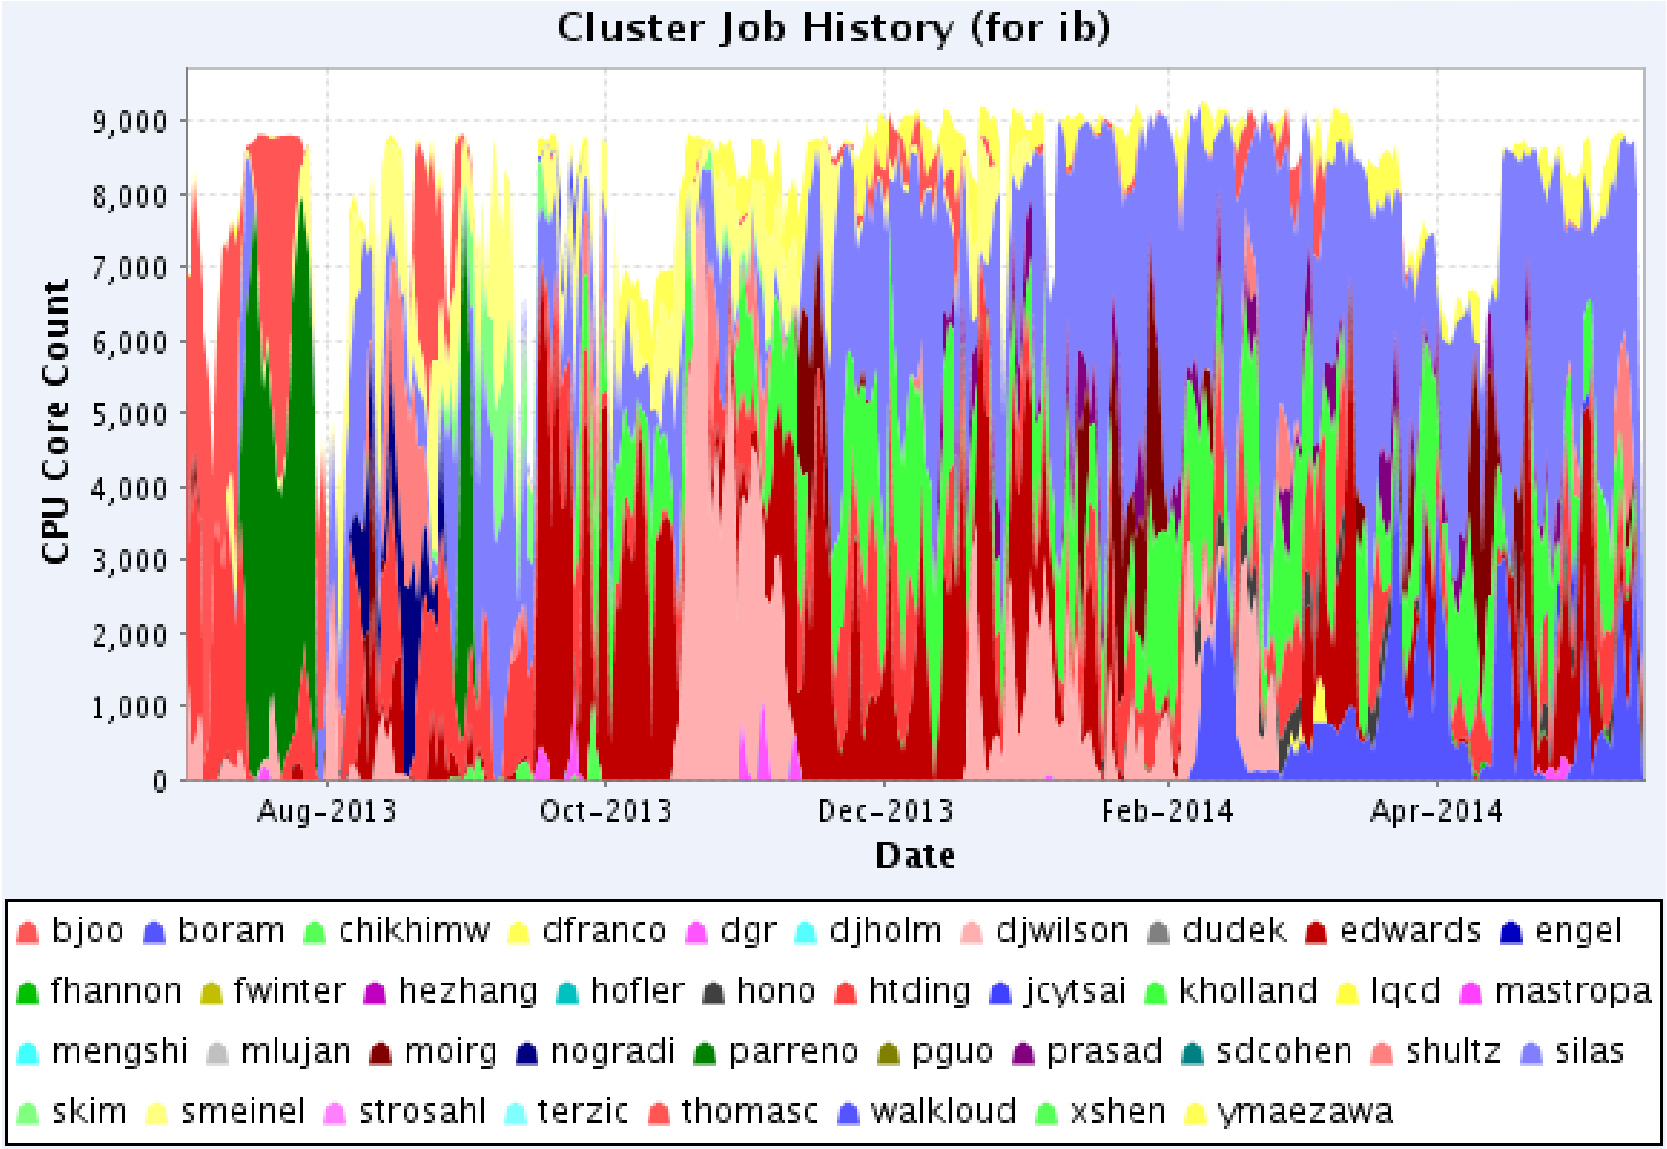

Farm Cluster: 1400 cores, hyperthreaded

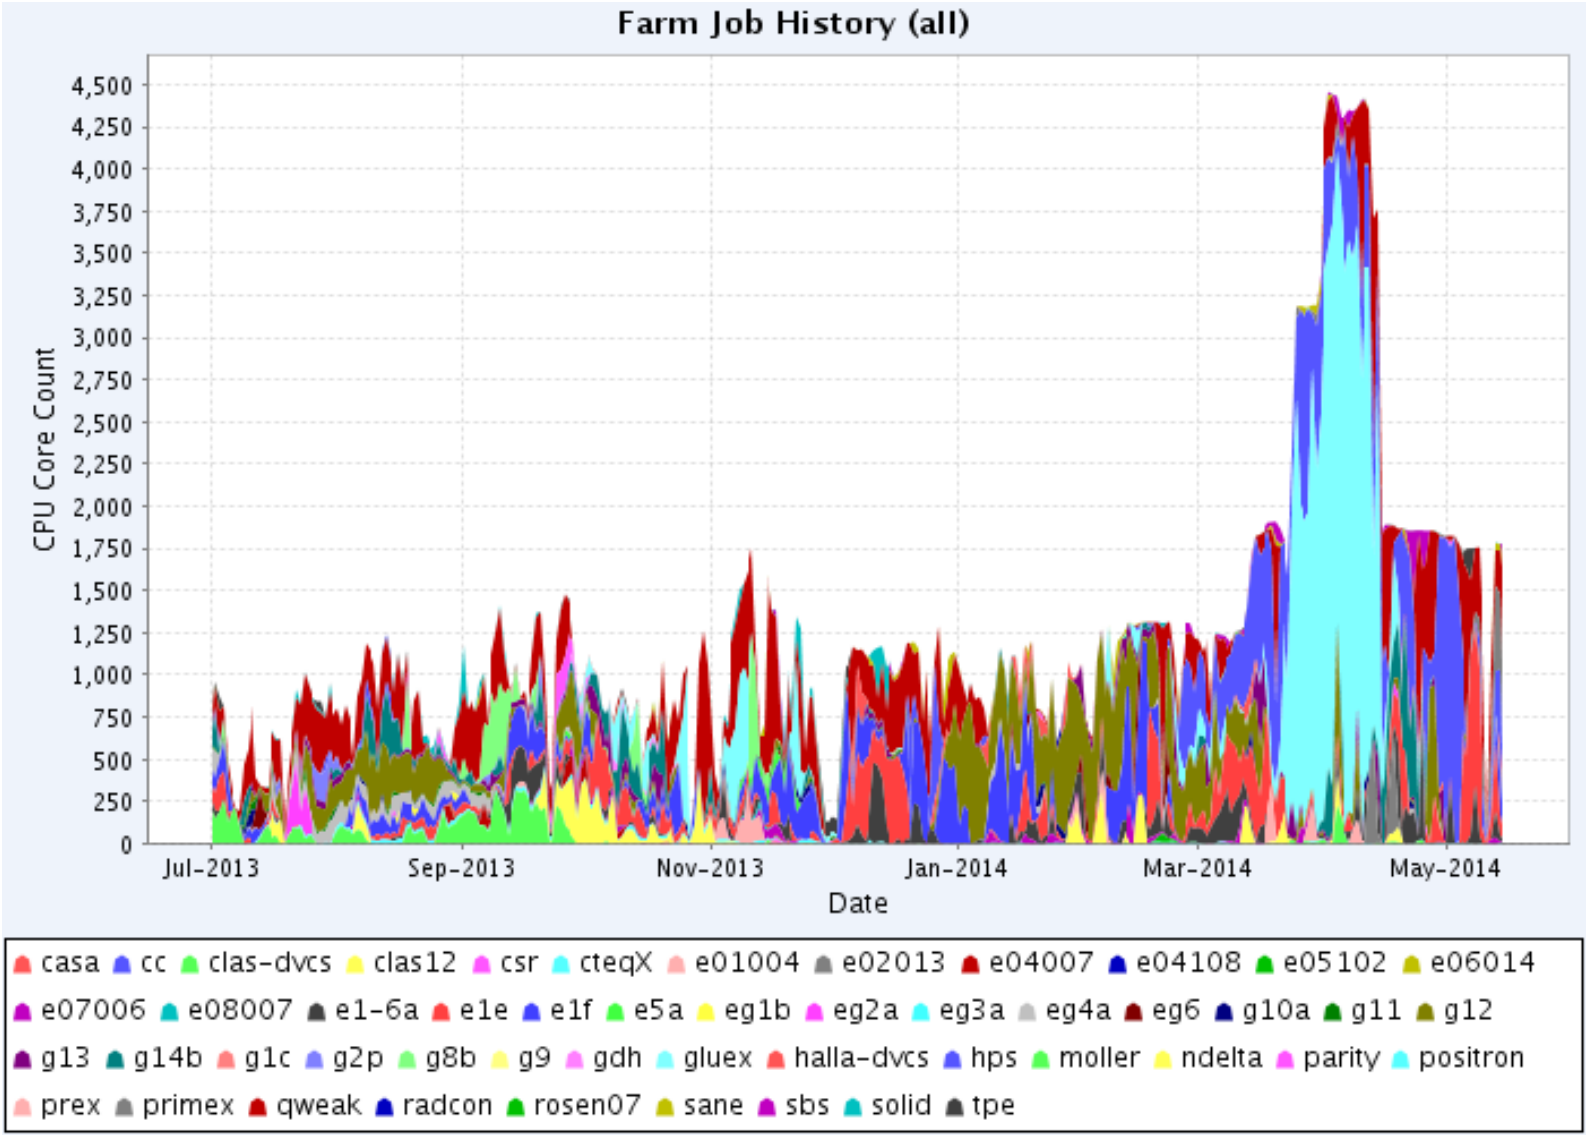

Win-Win

Future Impacts

- As long as the farm is busy, LQCD users don't have to worry as much about losing any of their allocations due to wasted core hours
- Physics can do data challenges at much larger scale than their resources for this year and next (delaying purchases until needed for production running)

We anticipate similar scale data challenges late summer (when LQCD typically lags) and again next winter (just after a farm upgrade).

The swings might grow larger as we try to use LQCD as a flywheel for the peaks and valleys of commissioning and early running of the GlueX detector. (We will work with users to ensure this won't harm LQCD productivity).

Storage and Networking

Disk Systems

- Lustre 1 PB on 30 OSSs each with 30 * 1/2/3 TB disks, 3 8+2 RAID6
  - /cache/project… write-through cache to tape & delete to stay under quota
  - /volatile/project… large scratch; delete to stay under quota
  - 8.1 GB / sec aggregate bandwidth
  - this will soon be increased ~20%
  - 100 MB/s – 1 GB/s single stream
- ZFS servers 250 TB
  - /work/project not backed up
  - /home/… nightly backup
  - **Move to ZFS on Linux**
    - retire 5 year old SunFire Thors; continue using our 2 year old Oracle appliance

Tape Library – IBM TS3500

- 14 frames; can add 1 or 2 more when needed (1300 slots each, 2.5 TB / slot) and can upgrade low density frames to high density, or buy 2nd library
- 2 LTO-6 drives, 10 LTO-5 drives, 2 LTO-4 drives => 1.6 GB/sec to / from tape

Offsite Gateway

- Globus Online works well for large bulk data transfers, up to ~ 5 Gbps (on 10Gbps link)
- Restricted shell implemented on the host (same as JLab login hosts)

Storage Evolution

Lustre Upgrade and Partitioning

- Each generation of server is higher capacity, higher performance
- Rolling upgrades: retire 5 year old systems, maintain capacity
- Upgrade from Lustre 1.8 to 2.5, partition by performance
  - Prototype Lustre with 2 pools: fastest/newest, and older/slower
  - Begin using striping, and all stripes will be fast (or all slow)
  - Use retired servers as a testbed, along with a **new higher performance MDS (ready now, what to buy??)**
  - By the end of 2014, this will be in production, with “inactive” projects moved from the main partition into the older, slower partition, freeing up highest performance disk space for active projects

- Question: how to name this new slow partition?
  - Option 1: no new name, inactive projects on slow, active on fast, and slow ones can be accessed via links: /cache/inactiveA -> /slowcache/inactiveA
  - Option 2: give users access to both fast and slow, and let them optimize, and “charge” half as much for slower servers?

Future Plans

2014

Increase the farm by ~ double as Halls come online. Upgrade Lustre; deploy 4 new OSSs (30 * 4 TB RAID6); move to ZFS on Linux. Begin using Puppet. Deploy workflow tool for farm. Continue to automate core sharing / load balancing.

2015 - 2016

Operate current HPC resources (minus oldest gaming cards): run the late Fall 2009 clusters through June 2015, and mid 2010 clusters through June 2016 -- longer than usual due to absence of hardware for 2015.

Experimental Physics grows to match the size of LQCD, enabling efficient load balancing (with careful balance tracking).

2016 – 2017

JLab will be the deployment site for the first cluster of LQCD-ext II.

This resource will be installed in the current location of the 9q / 10q clusters (same power and cooling, thus lower installation costs).

Continue to grow physics farm to meet 12GeV computing requirements. Final configuration will use ~20,000 cores.