

LQCD Computing at BNL

2014 USQCD All-Hands Meeting
JLAB
April 18, 2014

Robert Mawhinney
Columbia University

BGQ Computers at BNL

USQCD half-rack
(512 nodes)

2 racks of DD1
RBRC

1 rack of DD2
BNL



USQCD use of BNL DD2 BGQ

USQCD has 10% of the available time on the BNL DD2 BGQ (pre-production)

This time is included in the allocations by the SPC

Chris Kelly is currently running his SPC allocated time on 512 nodes of DD2. He will run for 1-2 months on 512 nodes to use USQCD 10%.

DD2 rack is running very well. Used extensively by BNL internal users.

USQCD 512 Node BGQ at BNL



USQCD 512 Node BGQ at BNL

Purchased with \$1.32 M from USQCD with FY13 Equipment Funds

Delivered in March, 2013

Install by IBM began on April 9, 2013

Turned over to users (Chulwoo) on Monday, April 15, 2013

Chulwoo ran DWF evolution of $32^3 \times 64 \times 24$ MDWF+ID strong coupling ensemble with $m_\pi = 140$ MeV for 1.5 days, with 100% reproducibility testing without problems

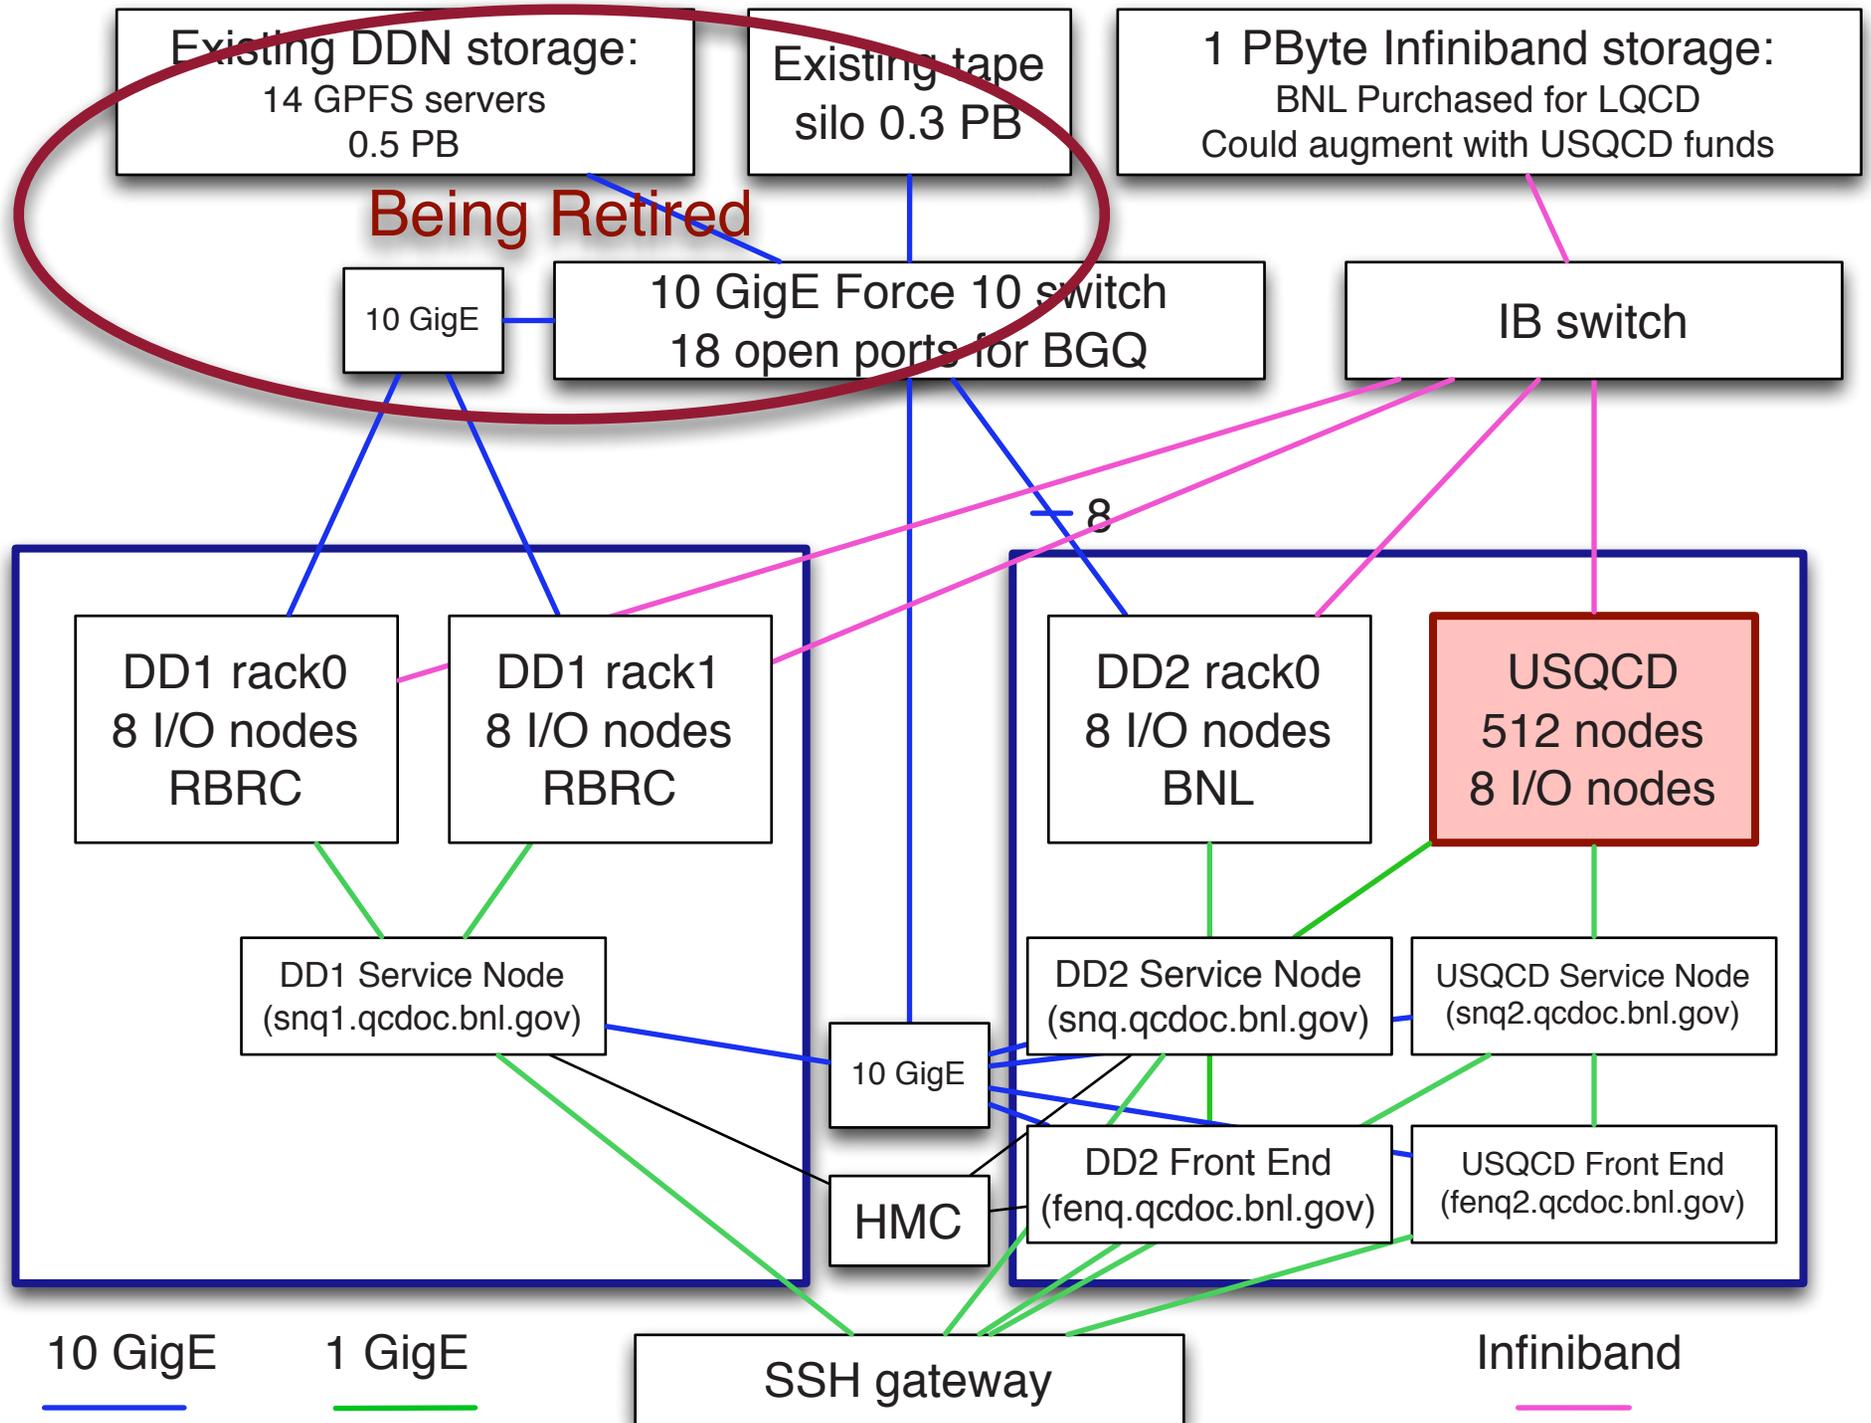
Began USQCD allocation running on July 1, 2013

3 allocated projects: Sugar 7.55, Mackenzie 18.65, Kelly 44.60 (M BGQ core-hours)

Usage as of 4/17/14: Sugar 5.71, Mackenzie 14.51, Kelly 28.78 (M BGQ core-hours)

Sugar running done early in allocation year, remaining time to be used by Mackenzie.

Expect to deliver full allocations to these three allocated projects.



USQCD BGQ Utilization at BNL

2013-2014 allocation month	Utilization	Comments
July	48%	Faulty compute node, IBM slow to diagnose. No hardware problems from March-June.
August	79%	2 day chilled water outage
September	90%	
October	91%	
November	83%	3 days lost to hardware failure
December	95%	
January	91%	Loadleveler hang
February	99%	

Almost all usage has been a single user running on 512 nodes full time.

Failure to achieve full utilization almost completely due to hardware problems.

Globus online access point at BNL being used by Detar to transfer MILC data. Are seeing intermittent freezing of transfers (a few hours of hang every 5 days or so). Trying to diagnose problem currently.

Near term BNL Possibilities

USQCD half-rack is supported by a total of 0.5 FTE at BNL. Cost effectiveness of computing increased by low personnel costs.

USQCD pays IBM for a service contract.

BGQ hardware now available at half price - possibility of filling up the rack

Important to consider buying more compute power and dropping service contract

- * Cost neutral option - trade remaining service contract costs for 512 more nodes.
- * BNL staff would maintain USQCD hardware, as they already maintain DD1 and DD2 hardware.
- * Might be able to add USQCD hardware in 2015, even with expected lower LQCD budget.

Summary

First year of USQCD BGQ running on track to deliver allocated computing time

Limited number of users - important that they be ready to run to keep machine full.

Cost neutral options for near term doubling of compute power

BNL has retired NY Blue, an IBM BG/L system.

- * Lab is engaged in seeking a replacement system - likely a phi cluster
- * Possibility for USQCD to augment such a system - more phi boards or next generation accelerators.