

LQCD Computing at BNL

2014 DOE Annual Review of LQCD Computing Project
FNAL
May 15-16, 2014

Robert Mawhinney
Columbia University

BGQ Computers at BNL

Half-rack of GA
USQCD

2 racks of DD1
RIKEN-BNL Research Center

1 rack of DD2
BNL



DD1 and DD2 are internal IBM revision labels for prototype BGQ hardware.

The DD2 hardware is very close to the final GA (General Availability) hardware.

The DD1 has reduced I/O bandwidth, slower network links and less memory.

USQCD 512 Node BGQ at BNL



Purchased with \$1.2 M from USQCD with FY13 Equipment Funds.

Delivered in March, 2013, first user began running April 15, 2013.

Early user running until July 1, 2013, including reproducibility testing

Standard BGQ environment with load leveler for queuing and XL compilers.

May 2013, 1 PByte Infiniband disk system purchased by BNL for the RBRC and BNL BGQ, primarily for LQCD. Shared with USQCD BGQ.

USQCD not funding for long-term archival storage at BNL.

USQCD use of BNL DD2 BGQ



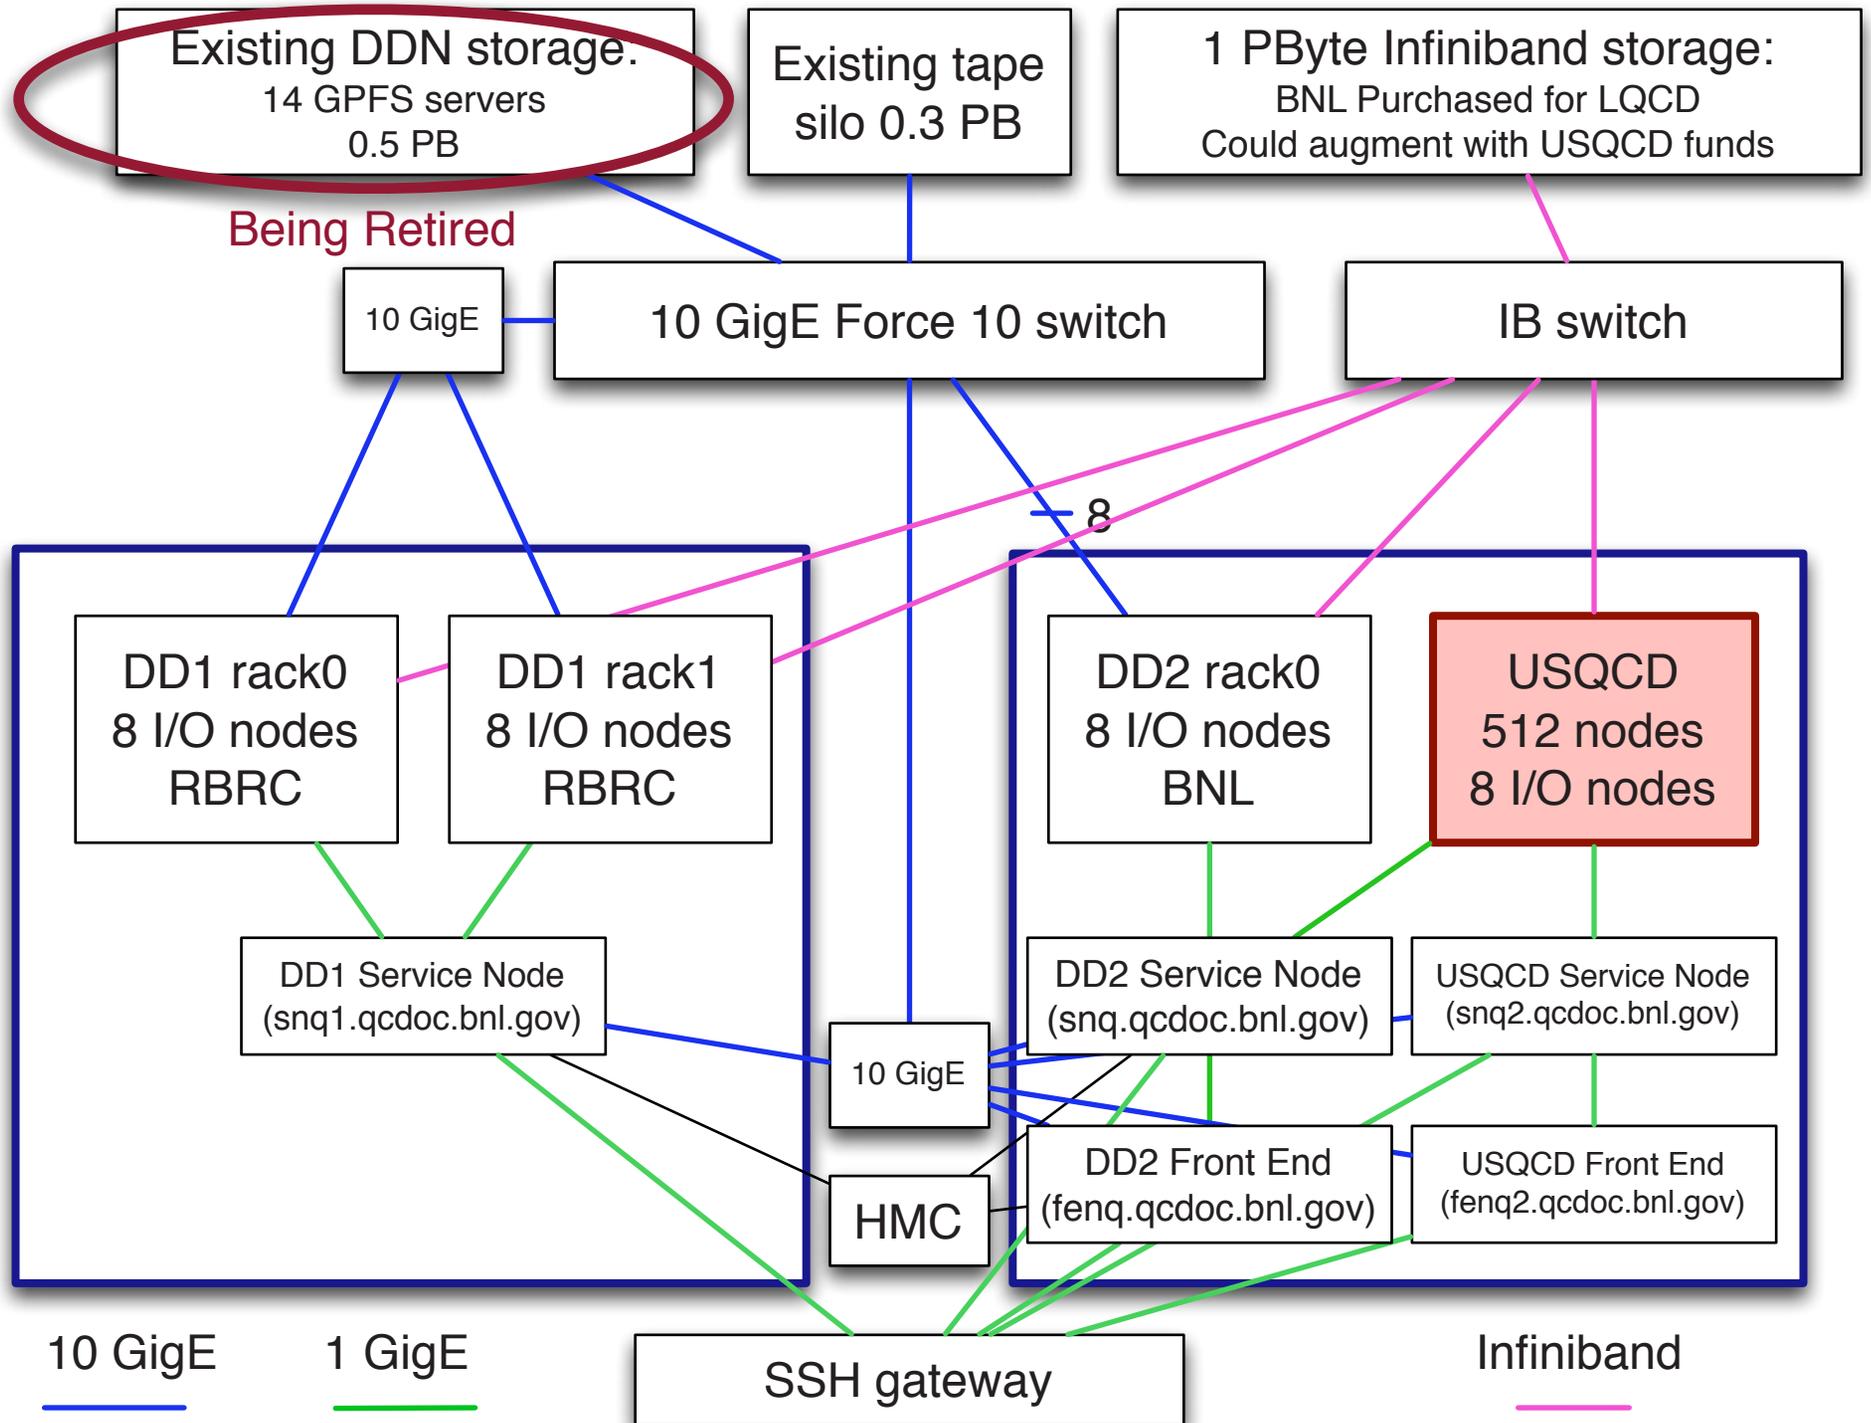
USQCD has 10% of the available time on the BNL purchased DD2 BGQ (pre-production) and provides funds for operation of the DD2.

This time is included in the allocations by the SPC.

DD2 rack runs very well. Used extensively by BNL internal users.

During early 2013, BNL LQCD users regularly ran single jobs taking 5.5 days on full rack.

During 2013-2014 allocation year, USQCD 10% will come from using 50% of the rack for 2+ months.



USQCD BGQ Utilization at BNL

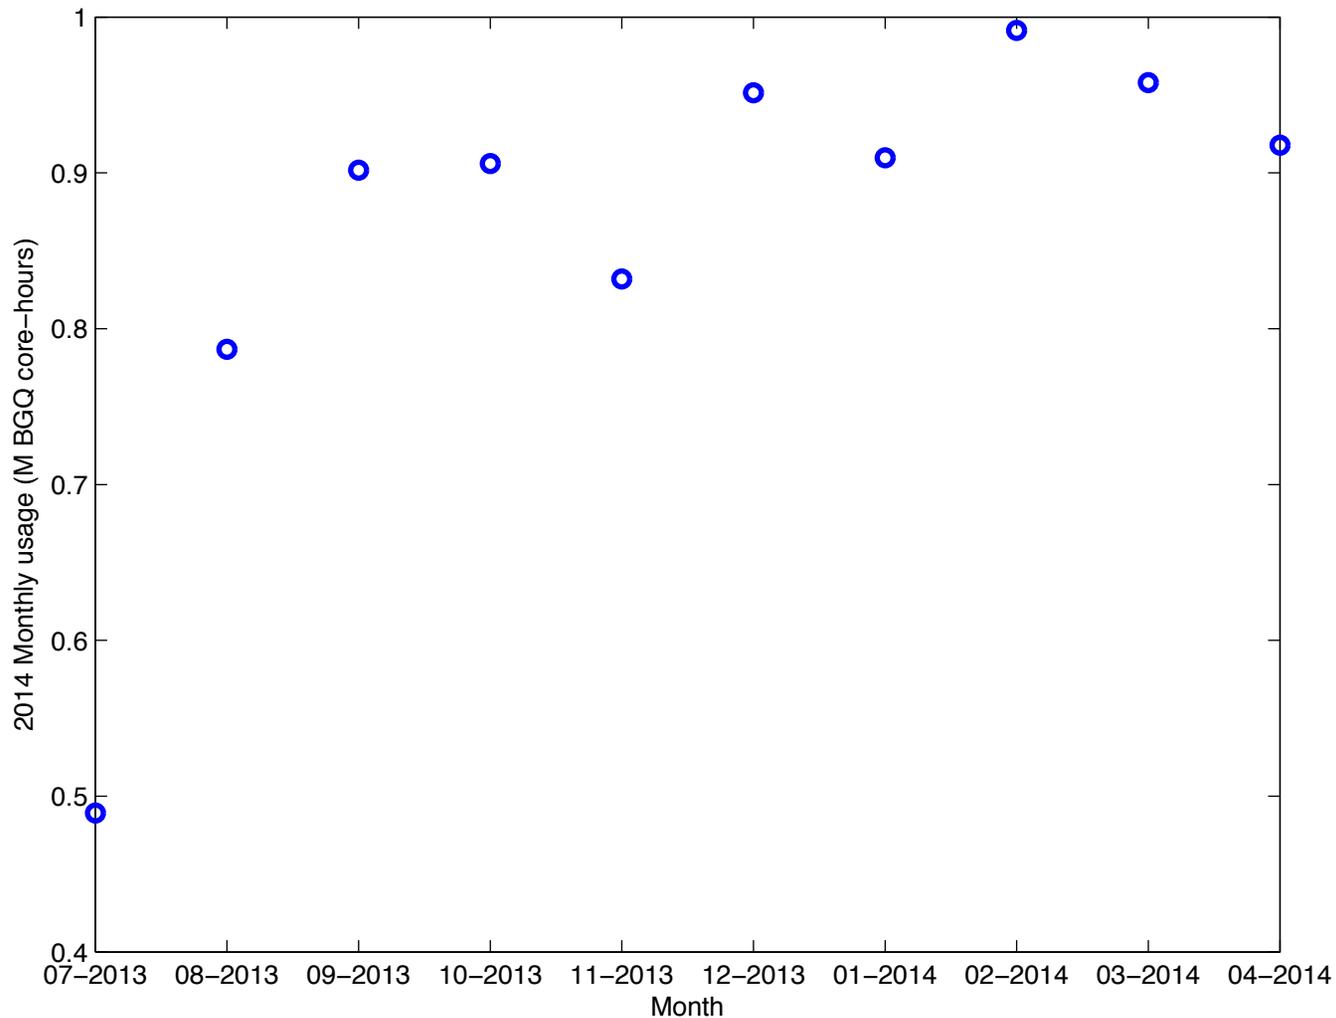
Month	Utilization	Comments
May, 2013	87%	Early user time
June, 2013	80%	Early user time
July, 2013	48%	Faulty compute node, IBM slow to diagnose.
August, 2013	79%	2 day chilled water outage
September, 2013	90%	
October, 2013	91%	
November, 2013	83%	3 days lost to hardware failure
December, 2013	95%	
January, 2014	91%	Loadleveler hang
February, 2014	99%	
March, 2014	96%	
April, 2014	92%	Filters cleaned. I/O drawer software error

Utilizations are fractions of total node-hours possible in a month, a very strict metric.

Almost all usage has been a single-user-at-a-time running on 512 nodes full time. No lost time for scheduler/queue.

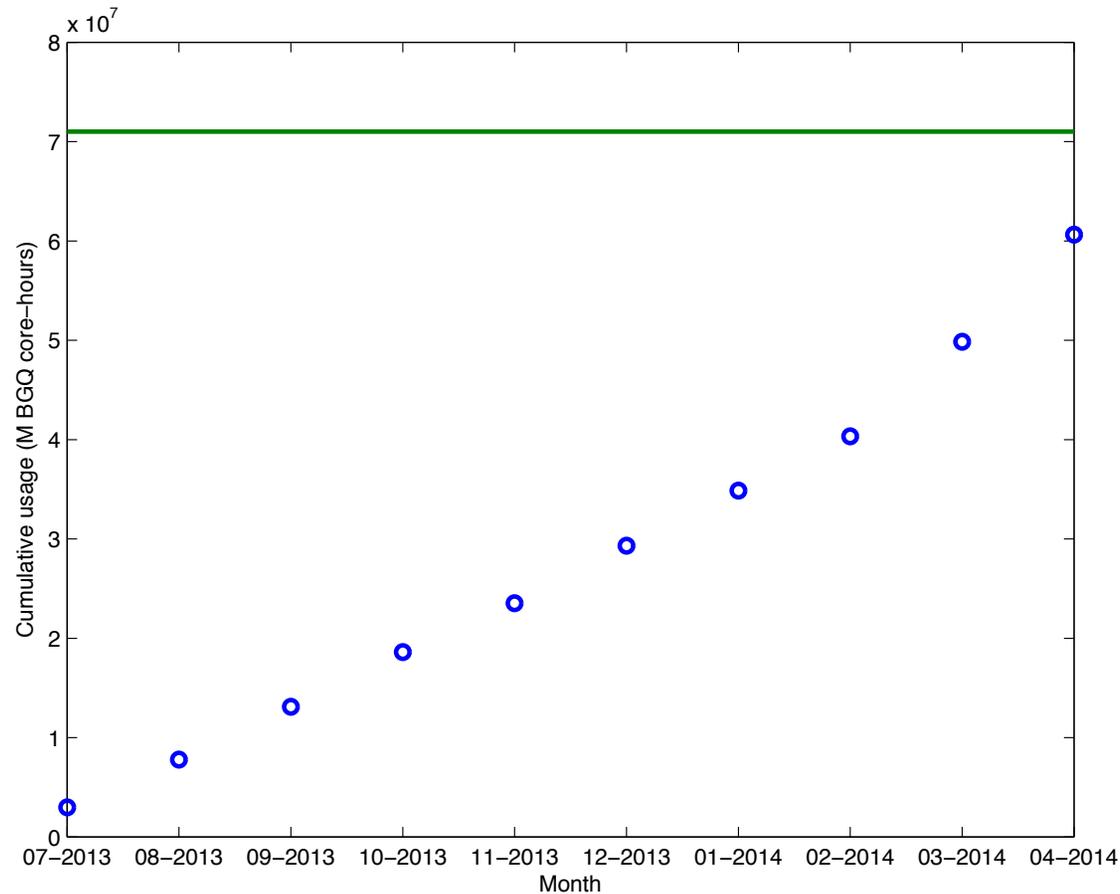
High volume data transfer to BNL from other sites required by one user to keep jobs running. Using Globus Online and working well, after some initial, intermittent connection authentication problems.

Monthly Utilization of USQCD half-rack



Utilizations are fractions of total node-hours possible in a month, a very strict metric.

Status of 2013-2014 Allocations



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 5/9/14: Sugar 5.71, Mackenzie 18.57, Kelly 39.38 (M BGQ core-hours)

Summary

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

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.

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

The USQCD, RBRC and BNL BGQ's have been extremely important for testing and developing methods and programs that are run on these BGQs and the larger machines at the ALCF and Livermore.