LQCD Computing at BNL

2013 USQCD All-Hands Meeting
BNL
April 19, 2013

Robert Mawhinney
Columbia University

BNL Computers used for QCD

12k node QCDSP, 600 GFlops, 1998-2005

2 ×12k node QCDOC, 20 TFlops, 2005-2011

2k node RBRC BGQ, 400 TFlops, 2012-
1k node BNL BGQ, 200 TFlops, 2012-

3k nodes RBRC/BNL BGQ, 600 TFlops, 2012-
0.5 k nodes USQCD BGQ, 100 TFlops, 2013-

USQCD use of BNL DD2 BGQ

- USQCD has 10% of the available time on the BNL DD2 BGQ (pre-production)
- Some non-RBC users have gotten accounts, but not used them
- RBC has been readily using the 10% of the DD2 for USQCD, primarily for pion/kaon measurements, both development and production.

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
- Machine shut down on report of detection of slow leak on Wed. morning, April 17. Reported to IBM and Joe Depace at BNL ran a calibration process on pressure sensors. Chulwoo restarted evolution job on 4/19/13.
- Standard BGQ production environment with load leveler for queuing and XL compilers.
- Currently mounting disks from front end node, awaiting new 1 PByte Infiniband system, expected in May.
- 1 PByte system was purchased by BNL, to be used primarily for LQCD. Should be augmented by USQCD funds, subject to general US budgetary issues.

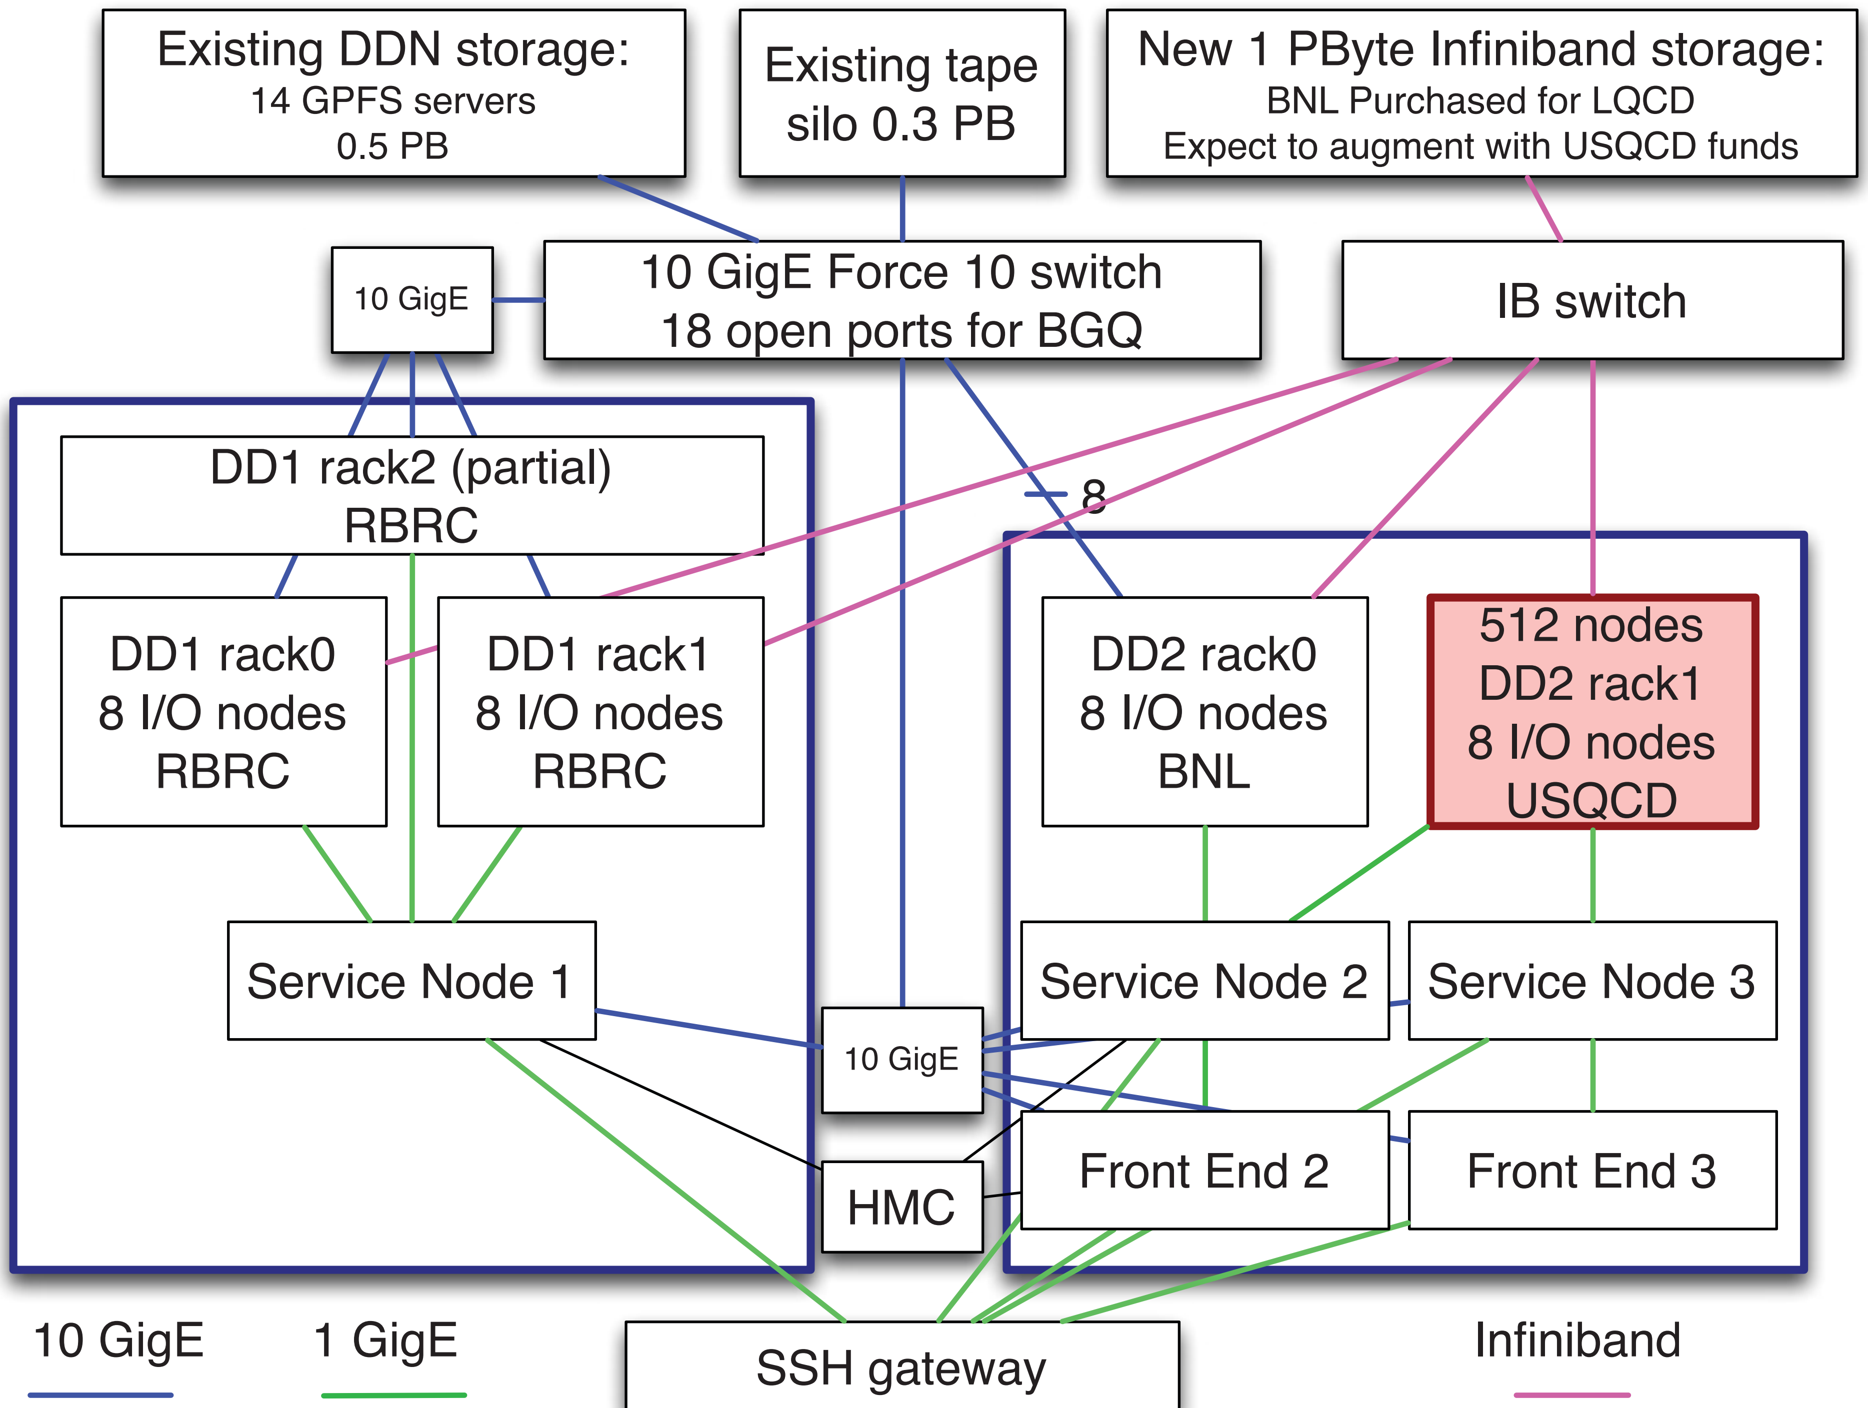
Existing DDN storage:
14 GPFS servers
0.5 PB
Existing tape
silo 0.3 PB
New 1 PByte Infiniband storage:
BNL Purchased for LQCD
Expect to augment with USQCD funds
10 GigE
10 GigE Force 10 switch
18 open ports for BGQ
IB switch
DD1 rack2 (partial)
RBRC
8
DD1 rack0
8 I/O nodes
RBRC
DD1 rack1
8 I/O nodes
RBRC
DD2 rack0
8 I/O nodes
BNL
512 nodes
DD2 rack1
8 I/O nodes
USQCD
Service Node 1
10 GigE
Service Node 2
Service Node 3
HMC
Front End 2
Front End 3
SSH gateway
10 GigE
1 GigE
Infiniband

More BGQ at BNL

- BNL can easily accomodate 1.5 more racks of BGQ for USQCD
- Current rack can be fully populated at any time. It has a heat exchanger between the cooling loop and the rack which can handle the load of a fully populated rack.
- Cooling and power is in place in the machine room for a second USQCD rack
  - * A second heat exchanger must be purchased
  - * A transfomer is required to convert existing power to voltage required for BGQ
  - * ≈ $100k infrastructure cost
- The current service node and front end can readily handle a second rack

LQCD Measurements

- Measurements on large volumes with deflation and all mode averaging can use large memory, long run times and tightly coupled architectures.
- Example: $48^3 \times 96 \times 24$ DWF simulations of RBC
  - DWF single precision even/odd preconditioned eigenvector is 12 GBytes
  - 600 single precision low modes takes 7.2 TBytes - must fit in memory to deflate
  - Deflated, sloppy solve (1e-4 stopping condition) takes 18 PFlop - fixes minimum machine size
  - If want solution in 1 hour, requires 5 TFlops sustained.
  - On 50 GFlops nodes this is 100 nodes, each with 72 GBytes of memory
  - Time for 96 solves (all times slices) is 96 hours or 4 days.
  - This doesn't include the time to generate the 600 low modes
  - For this example, more low modes would be better.
- RBC pion/kaon measurement package on $48^3 \times 96 \times 24$ takes 5.2 days on 1 rack BGQ. Rack-hours for a given statistical accuracy reduced 5-20× compared to earlier methods without deflation and/or low-mode averaging.

- 10x faster nodes requires 720 GBytes/node to hold mode for deflation.
- 0.4 days to solution, but memory size is prohibitive.
- Need sufficient network bandwidth between nodes to keep 10x faster node running.
  * Hyung-Jin Kim (BNL): Put $48^3 \times 96 \times 24$ DWF calculation on 72 GPUs
  * No deflation in this test, so memory is not an issue
  * Sustains 3547 GFlops, or 49.2 GFlops/GPU
  * Currently, GPU's not able to get good performance for this size lattice
- 10x as many nodes is viable, since then memory is 7.2 GBytes/node, but require a network which can support local CPU speed without stalling.
  * A 1000 node cluster or a BGQ rack is a reasonable size
  * Need multiday reliability, including no dropped bits, to avoid excessive I/O

Other Algorithms

- Domain decomposition, inexact deflation, and/or multigrid do not require as much memory
- Working examples for Wilson/clover fermions.
- DWF: attempts (so far) not viable. Most CPU time ends up in little Dirac operator
- This can be a very dense matrix
  * Parallelization of this can require handling many small messages
  * BGQ network is has low latency and can handle the many small messages neede to get good performance on little Dirac operator
  * Peter Boyle is pursuing this direction for DWF on BGQ
- Future is hard to predict, but network, reliability and memory of BGQ makes it very competitive, particularly for measurement jobs which would have to span many 10's of GPUs.

Summary

- BNL has successfully managed QCDSP, QCDOC, BG/L, BG/P and now BG/Q
- USQCD half-rack operational - initial burn in phase underway
- Should be available to interested USQCD members in a month or so. Allocations start July 1, 2013.
- BNL can readily add 1.5 more BGQ racks, with minimal costs beyond the racks themselves.
- Opportunity for substantial increase in USQCD resources for both generating lattices and large evolution jobs
- Future:
  - Precision measurements can be done $\approx 10\times$ faster with deflation and all mode averaging, provided machines have sufficient memory and reliability
  - Large volume work requires a powerful network
  - Argues for continued USQCD access to BGQ-style machine and its successors.
  - BNL is obvious location to continue to locate these machines.