Report on QCDOC at BNL

All Hands' Meeting

US Lattice QCD Collaboration Meeting

FNAL, April 6-7, 2006

Stratos Efstathiadis

BNL

Brookhaven Science Associates

U.S. Department of Energy

BROOKHAVEN NATIONAL LABORATORY

OUTLINE

- Introduction
- QCDOC Overview

  (Architecture, Available Partitions, Current Status)
- User Environment and Monitoring

  (Network, Batch System, Resource Monitoring…)
- Conclusions

Introduction

- QCDOC: Quantum ChromoDynamics On Chip
- Optimized for Lattice QCD Calculations

  Capable of strong scaling: use many nodes on a small problem:
  - Large inter-node bandwidth
  - Small communications latency
- Designed for optimal price/performance, low power, modular design, high scalability and reliability.
- Collaboration between IBM, Columbia, UKQCD, RBRC and BNL
- Three large QCDOC machines in production.

QCDOC Collaboration

- Columbia (DOE)
  - Norman Christ
  - Saul Cohen
  - Calin Cristian
  - Zhihua Dong
  - Changhoan Kim
  - Ludmila Levkova
  - Xiaodong Liao
  - Meifeng Lin
  - Guofeng Liu
  - Robert Mawhinney
  - Shu Li
  - Azusa Yamaguchi
- BNL (SciDAC)
  - Chulwoo Jung
  - Konstantin Petrov
  - Stratos Efstathiadis
- UKQCD (PPARC)
  - Peter Boyle
  - Mike Clark
  - Balint Joo
- RBRC
  - Shigemi Ohta
  - Tilo Wettig
- IBM
  - Dong Chen
  - Alan Gara
  - Design groups:
    - Yorktown Heights, NY
    - Rochester, MN
    - Raleigh, NC

QCDOC Chip

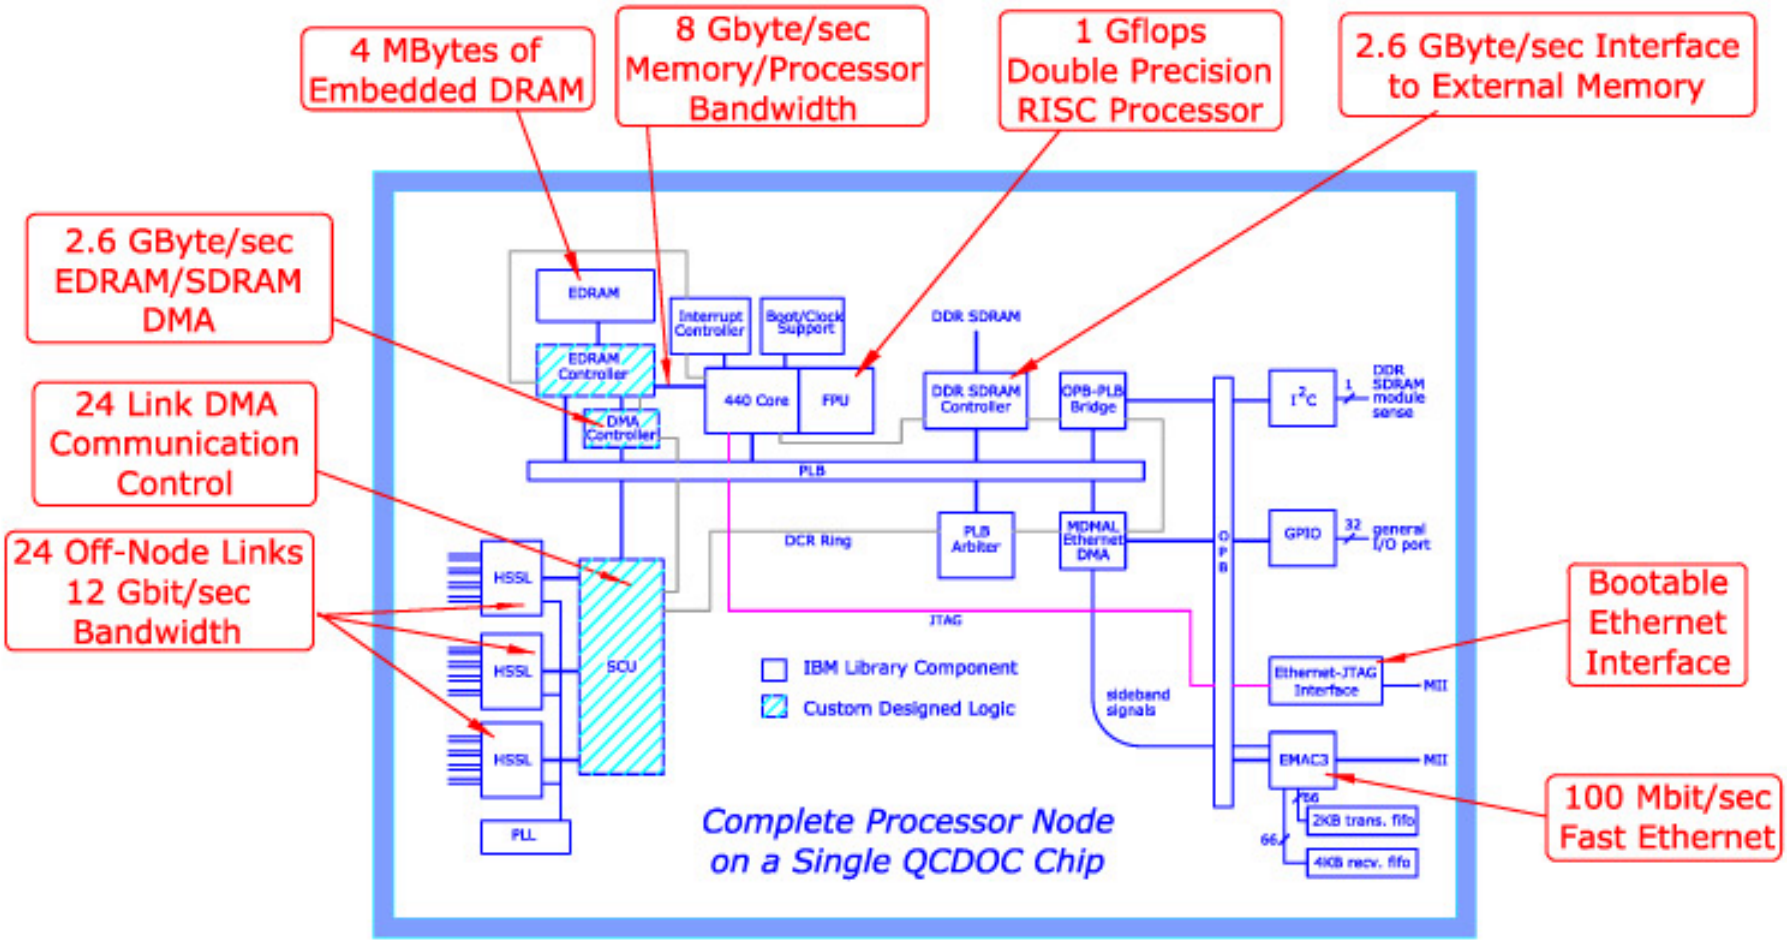
4 MBytes of Embedded DRAM
8 Gbyte/sec Memory/Processor Bandwidth
1 Gflops Double Precision RISC Processor
2.6 GByte/sec Interface to External Memory
2.6 GByte/sec EDRAM/SDRAM DMA
24 Link DMA Communication Control
24 Off-Node Links 12 Gbit/sec Bandwidth
Bootable Ethernet Interface
100 Mbit/sec Fast Ethernet
EDRAM
EDRAM Controller
DMA Controller
Interrupt Controller
Boot/Clock Support
440 Core
FPU
DDR SDRAM
DDR SDRAM Controller
OPB-PLB Bridge
PLB
PLB Arbiter
DCR Ring
MDMAL Ethernet DMA
HSSL
HSSL
HSSL
SCU
PLL
JTAG
IBM Library Component
Custom Designed Logic
sideband signals
I2C
DDR SDRAM module sense
GPIO
general I/O port
Ethernet-JTAG Interface
MII
EMAC3
MII
2KB trans. fifo
4KB recv. fifo
Complete Processor Node on a Single QCDOC Chip

QCDOC Construction

An ASIC (node). ~5 Watt at 400MHz

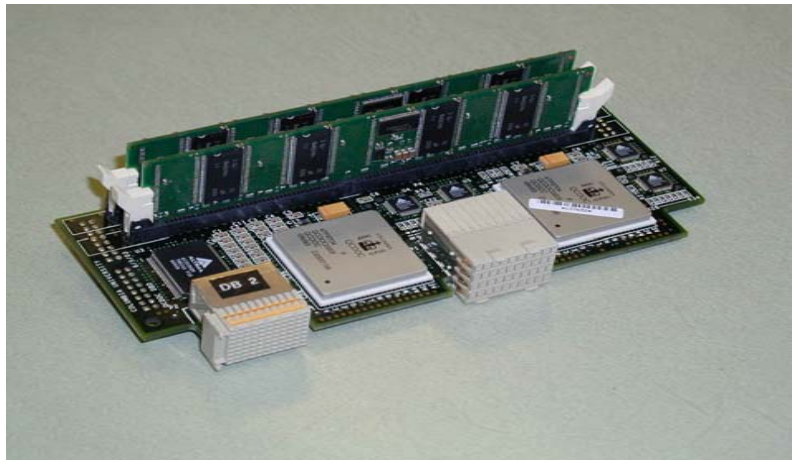

A daughterboard with two nodes and the vertically mounted DDR SDRAMs (128MB at BNL)

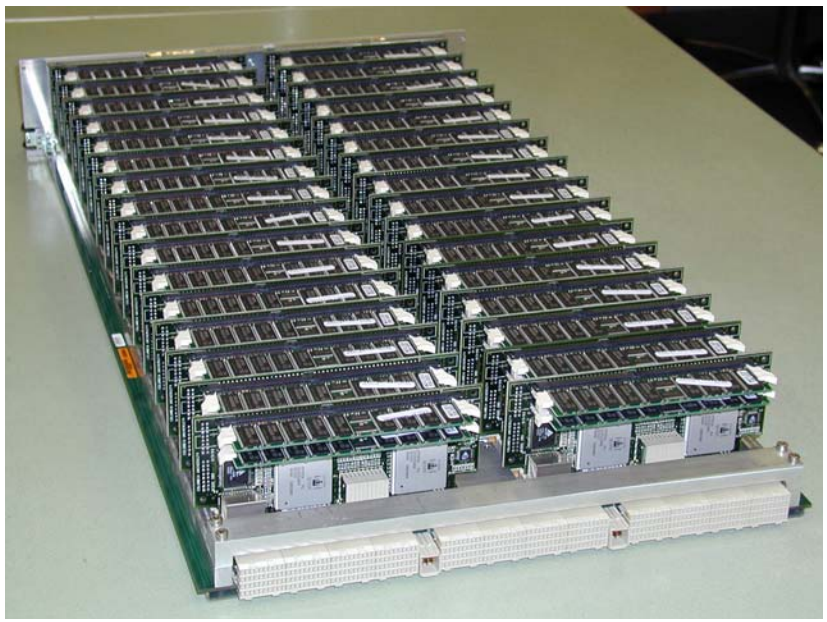

A single motherboard. Two rows of 16 daughterboard with 2 nodes each provide a total of 64 nodes. 14.5in x 27 in

A water-cooled rack containing 16 MBds with 1024 nodes. The upper compartment holds Ethernet switches

Networks:

- Black squares are nodes
- Blue Boxes are Mother boards
- Red Lines are communication links
- Green lines are Ethernet connections
- Green boxes are Ethernet switches
- Pink boxes are host CPU processors

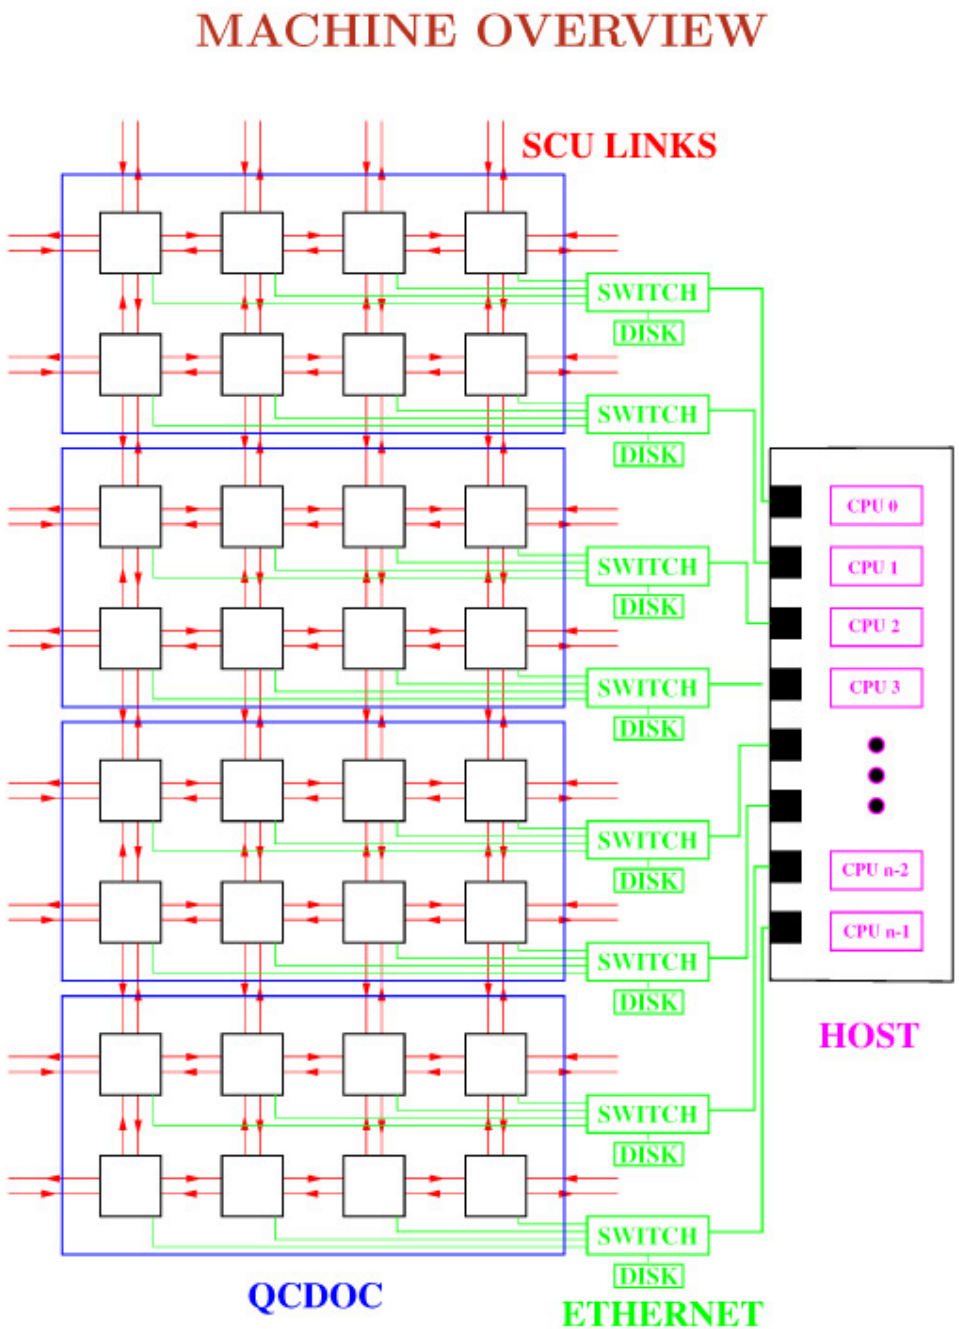

RBRC (right) and DOE (left) 12K-node QCDOC machines

USDOE QCDOC
Dedication Ceremony
November 30th 2005.

12 water cooled racks (12288 nodes),
$5.1M

QCDSP

Air Cooled Crates

Single Slot Back Plane

Current QCDOC Usage

	Machine partition	Nodes	Wired	Use
MILC (Dru Renner, Carleton Detar)	Rack 24-27	4096	8x8x8x2x2x2	48^3 x 144 lattice generation
	Rack 22-23	2048	16x4x4x2x2x2	40^3x96 lattice generation
RBC (Norman Christ, MF Lin, S.Cohen)	Rack 20-21	2048	16x4x4x2x2x2	DWF
LHPC (R. Edwards)	Rack 19	1024	8x4x4x2x2x2	Hadron Structure on MILC Lattices
Thermodynamics (Karsch Frithjof, Chulwoo Jung)	Rack 18	1024	8x4x4x2x2x2	P4 Thermodynamics
	Rack 17 – crate1 Rack 17 – crate 0	512 512	4x4x4x2x2x2 4x4x4x2x2x2	
LHPC (Dru Renner)	Rack 16 (14 MBs)	896	14x4x2x2x2x2	DWF Propagator Generator
	Rack 16 – crate 0 – slot 0 Rack 16 – crate 1 – slot 0	64 64		

The QCDOC Network Enclave

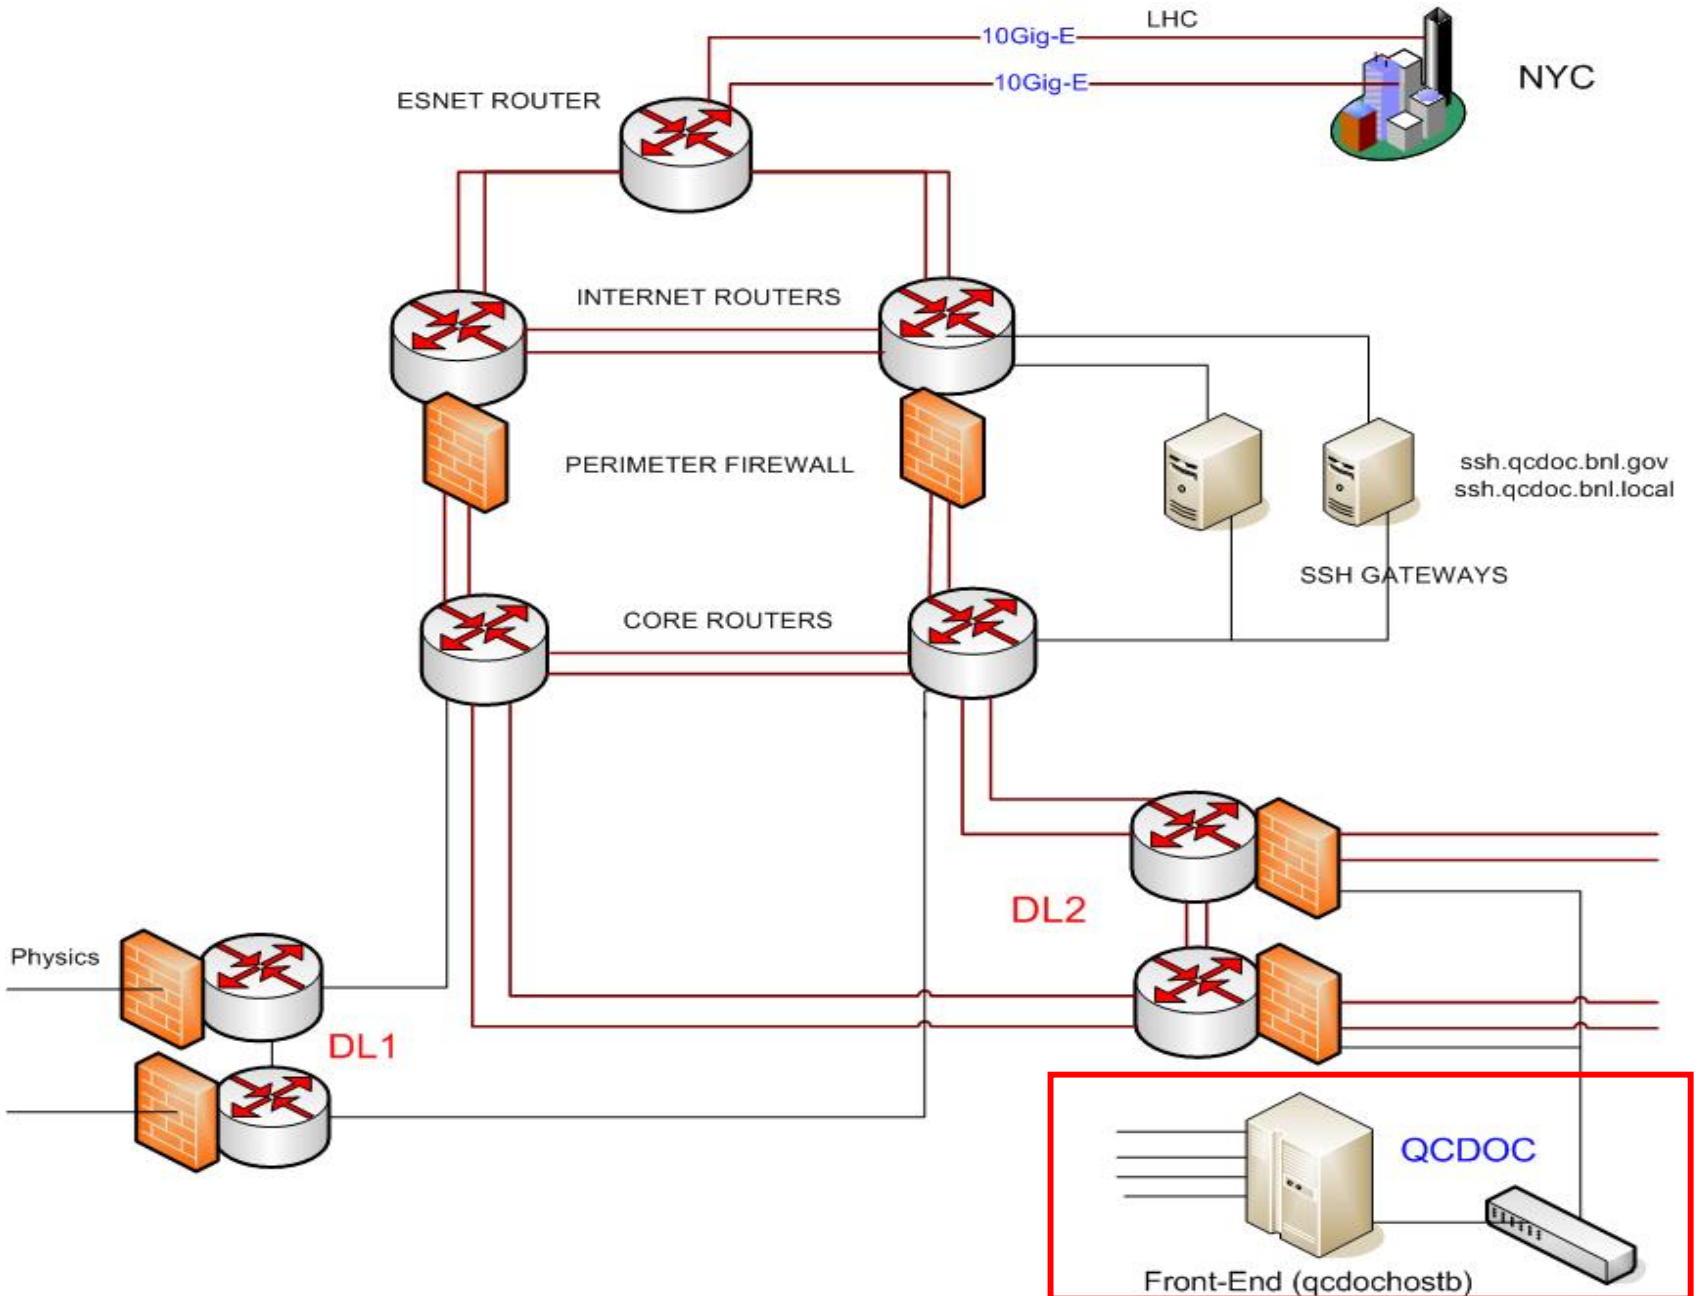

The QCDOC Network Enclave

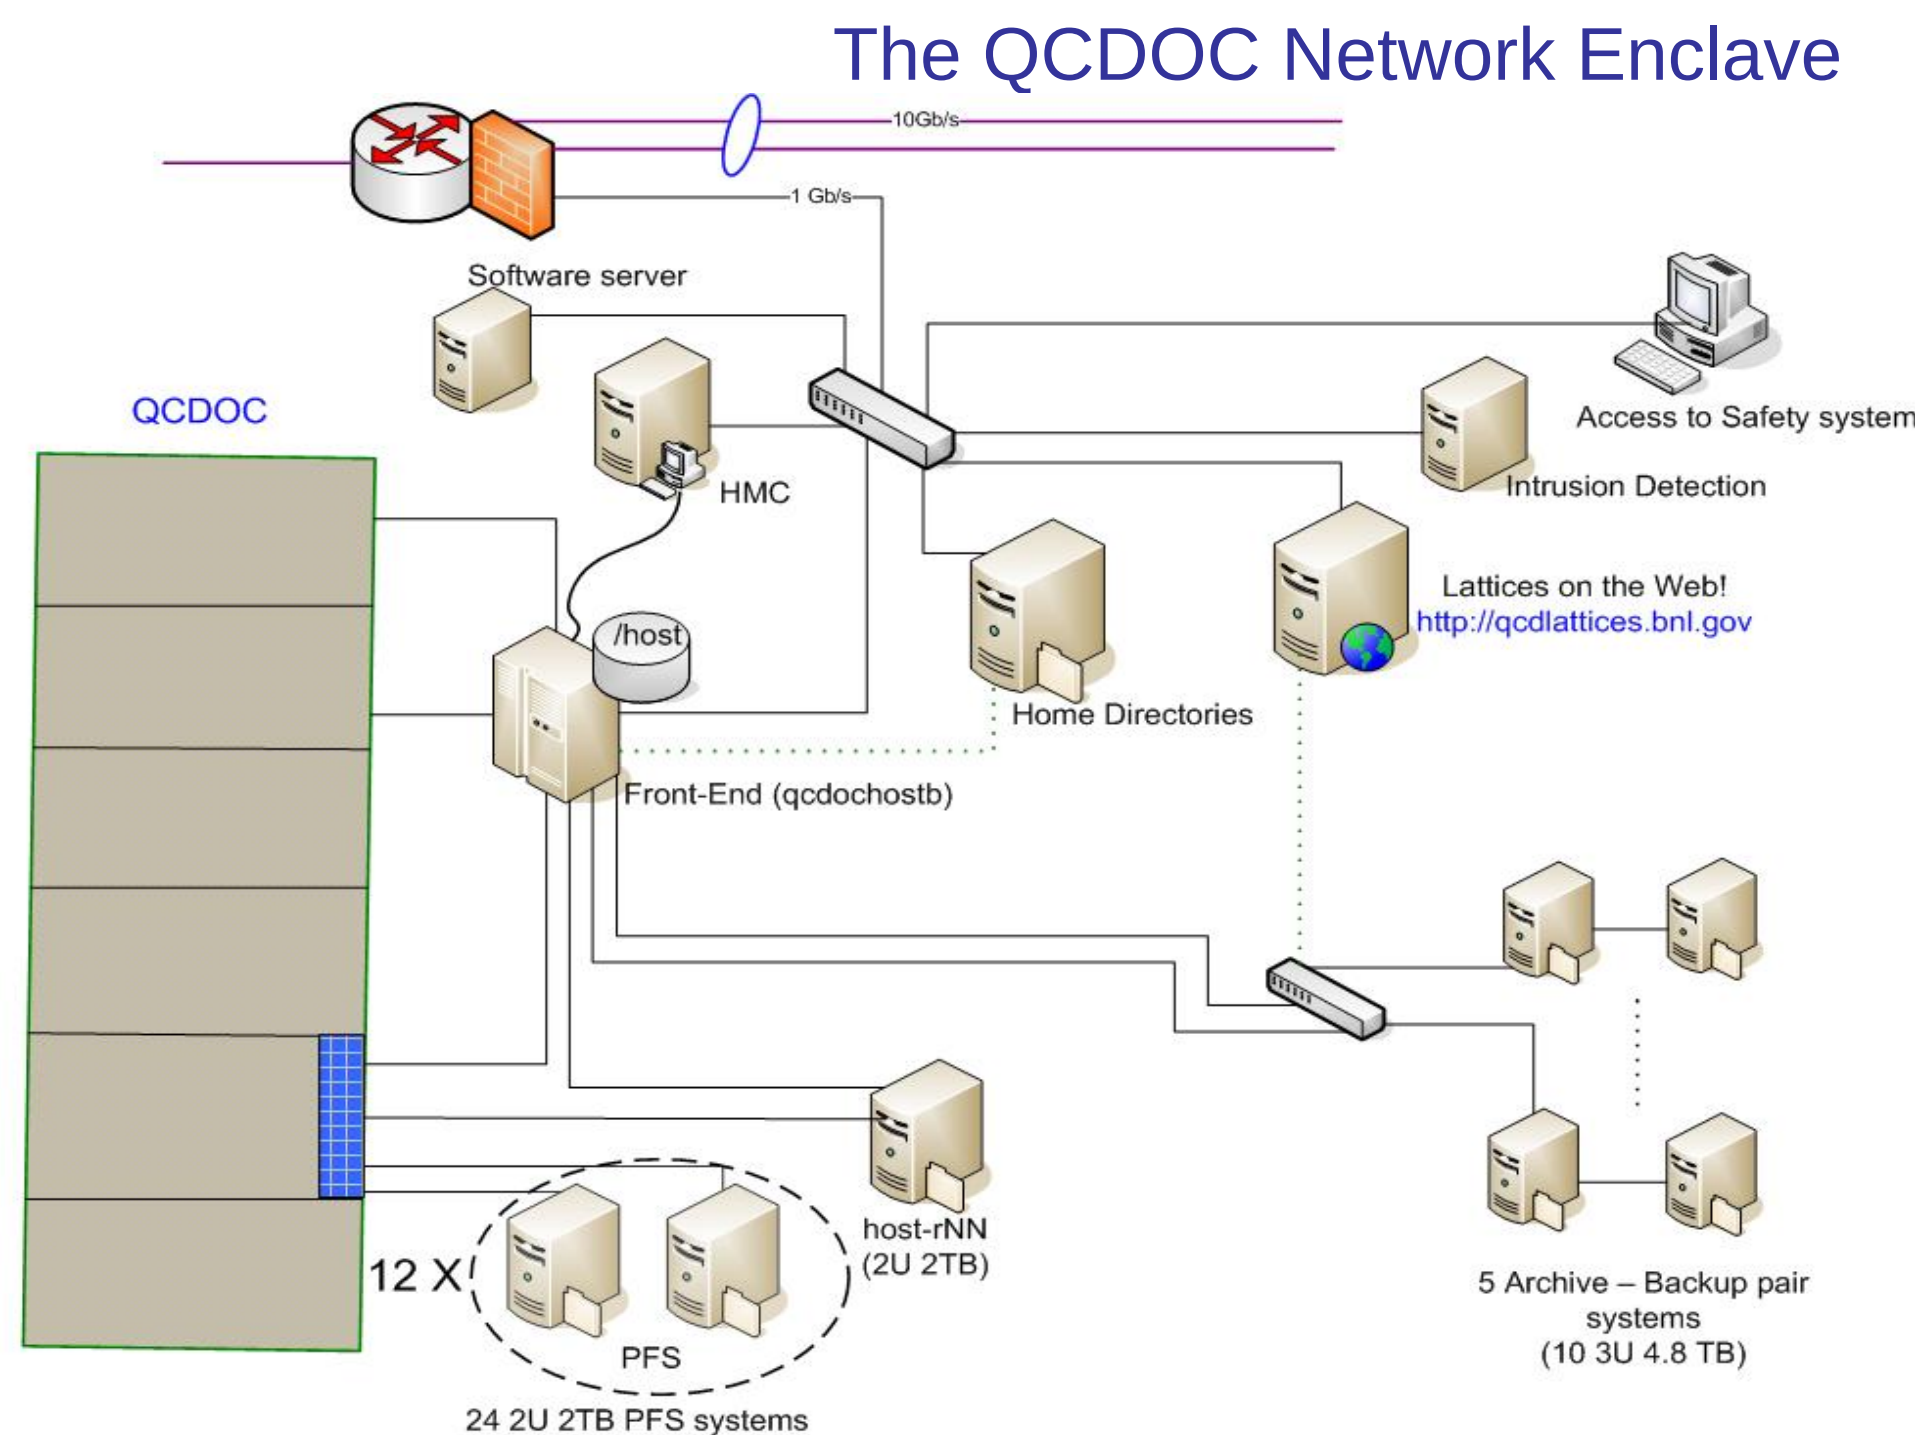

Host Computer

- The front-end host is an IBM P655 server with 8CPUs, 8GB of memory running AIX.
- Multiple network interfaces provide access to machine partitions.
- Users cross-compile code and manage allocated machine partitions using host software (qdaemon, qcsh).

Disk Systems: /host

- A custom NFS client is part of the node kernel supporting two mount points (open/read/write/close).
- A disk file system that is globally shared by all partition nodes and accessible by the front-end system.
- Usually a disk on the front-end systems.

Disk Systems: /pfs

- The Parallel File System (pfs) is used for high throughput IO from machine partition nodes.

- Provided by 2U rack-mounted LINUX NAS Servers. 2 RAID-5 file servers per machine rack (one per crate), total disk space 48TB.

- Similar to cluster “scratch-disk” on every node.

- Each NAS is shared by several nodes. Each node is using a unique directory, ex. */R24/C0/B0/D21/A1/*

- Temporary data staging, not backed up.

- Accessible by the front-end systems (ex. */pfs/r16c0/R16/C0/…*)

- Tunned NFS, RAID controller and kernel parameters to optimize IO performance (same I/O tunning across the Atlantic).

Local Storage and transfers of Lattice Configurations

- 10 4.8TB ANACAPA file servers make up five archive/backup pairs. The five archive servers are auto-mounted on the front-end host while the backup servers simply mirror the archives.

- Transferring files to BNL (or Jlab) may be a 2-hop process or use ssh tunneling (dedicated qcdoc ssh gateways at BNL).

- Transferring files to FNAL, a kerberized utility must be used (rcp, scp) after obtaining a kerberos ticket (kinit).

Lattice Configurations on the web

- The purpose is to make a number of generated lattices easily available to the community with various degrees of password protection.

- Set up web site inside qcdoc network enclave (reverse-proxy to the outside world): http://qcdlattices.bnl.gov

- Configurations can be downloaded via http. Provide tools to easily download “clusters” of lattices (wget wrapper).

- Configurations are in the QCD Archive Format.

Resource Monitoring (Servers)

- Nagios monitors our servers (front-end, archive, pfs, host and home file servers and ssh gateways) for usage CPU loads, services (ssh, NFS), disk space etc.

QCDOC Systems - high priority (qcd-important)

Host	Services
qcdochosta.qcdoc.bnl.gov	DISK:/ DISK:/data0 DISK:/host DISK:/host-r12 DISK:/host1 DISK:/opt DISK:/pfs/acc5 DISK:/pfs/acc8 DISK:/pfs/r04 DISK:/pfs/r05 DISK:/pfs/r06 DISK:/pfs/r07 DISK:/pfs/r08 DISK:/pfs/r09 DISK:/pfs/r10 DISK:/pfs/r DISK:/pfs/r12 DISK:/pfs/r13 DISK:/pfs/r14 DISK:/pfs/r15 DISK:/qcdoc DISK:/qdata00 DISK:/space DISK:/tmp DISK:/usr DISK:/var LOAD NFS server:10.20.1.10 NFS server:pfs-acc5 NFS server:pfs-acc8 NFS server:pfs-r0 NFS server:pfs-r06 NFS server:pfs-r07 NFS server:pfs-r08 NFS server:pfs-r09 NFS server:pfs-r10 NFS server:p r12 NFS server:pfs-r13 NFS server:pfs-r14 NFS server:pfs-r15 NFS server:qcdraida NFS server:qcdraidb SSH
qcdochostb.qcdoc.bnl.gov	DISK:/ DISK:/archive/a0 DISK:/archive/a1 DISK:/archive/a2 DISK:/archive/a3 DISK:/archive/a4 DISK:/host DIS DISK:/host-r18c0 DISK:/host-r19c0 DISK:/host-r20c0 DISK:/host-r22-23 DISK:/host-r24-27 DISK:/host1 DISK: DISK:/pfs/r16c0 DISK:/pfs/r16c0a DISK:/pfs/r16c1 DISK:/pfs/r16c1a DISK:/pfs/r17c0 DISK:/pfs/r17c1 DISK:/pf DISK:/pfs/r19c0 DISK:/pfs/r19c1 DISK:/pfs/r20c0 DISK:/pfs/r20c1 DISK:/pfs/r21c0 DISK:/pfs/r21c1 DISK:/pfs/ DISK:/pfs/r23c0 DISK:/pfs/r23c1 DISK:/pfs/r24c0 DISK:/pfs/r24c1 DISK:/pfs/r25c0 DISK:/pfs/r25c1 DISK:/pfs/ DISK:/pfs/r27c0 DISK:/pfs/r27c1 DISK:/qcdoc DISK:/tmp DISK:/usr DISK:/var LOAD NFS server:doearchive0 NFS server:doearchive1 NFS server:doearchive2 NFS server:doearchive3 NFS server:doearchive4 NFS server:h server:host-r24-27 NFS server:pfs-acc7 NFS server:pfs-r16c0 NFS server:pfs-r16c0a NFS server:pfs-r16c1 NFS server:pfs-r16c1a NFS server:pfs-r17c0 NFS server:pfs-r17c r18c0 NFS server:pfs-r18c1 NFS server:pfs-r19c0 NFS server:pfs-r19c1 NFS server:pfs-r20c0 NFS server:pfs-r20c1 NFS server:pfs-r21c0 NFS server:pfs-r21c1 N r22c0 NFS server:pfs-r22c1 NFS server:pfs-r23c0 NFS server:pfs-r23c1 NFS server:pfs-r24c0 NFS server:pfs-r24c1 NFS server:pfs-r25c0 NFS server:pfs-r25c1 N r26c0 NFS server:pfs-r26c1 NFS server:pfs-r27c0 NFS server:pfs-r27c1 NFS server:qcdraida NFS server:qcdraidc SSH
qcdraida.qcdoc.bnl.gov	DISK:/ DISK:/boot LOAD SSH
qcdraidb.qcdoc.bnl.gov	DISK:/ DISK:/boot LOAD SSH

QCDOC Batch System

- Goal: increase machine usage and productivity.

- Standard Batch features:
  - File Staging
  - Accounting
  - Job Status Notification
  - ACLs
  - Resource Limits

- Added Features in scripts:
  - Partition Allocation
  - Partition Startup
  - Idle Job checks
  - Partition resets and Powercycling
  - qhdwcheck wrapper
  - Job re-runs and Error Limits
  - Error and Event Logging, User Notification

QCDOC Batch System

- PBS (Torque) with MAUI scheduler.

- Restricted machine partitions:
  - Each mapped to a PBS queue.
  - Simple FIFO scheduling
  - ACL limits user access
  - No limits on resource usage
  - Partition can allocated using the web allocator

- ACC Single MBd partitions:
  - Queues with walltime limits (OneHr, FourHr, EightHr and SixteenHr) on four MBds (slot4-7)
  - Queued jobs may transfer to empty (not running) queues.
  - Queues can be marked as unavailable if there are hardware problems

- Latest version of PBS scripts:
  /qcdoc/local/batch/v3
- Web site: http://www3.bnl.gov/qcdoc/pbs

Job Accounting

Monitor *qdaemon* and *qcsh* processes on the front-end host:
http://www3.bnl.gov/qcdoc/status

QCDOC Current Status

(Water-Cooled Racks Only)

Partition (Click for partition details)	Allocated? (Click to Allocate)	Alloc. User	Run User (Click for User Jobs)	Run Time [Days-]Hrs:Mins:Secs
rack16/14mb	True	drenner	drenner	3-12:28:58
rack16/crate0/slot0	False	-	-	-
rack16/crate1/slot0	True	mawhinne	-	-
rack17/crate0	True	chulwoo	chulwoo	08:44:50
rack17/crate1	True	chulwoo	chulwoo	9-15:30:48
rack18	True	chulwoo	chulwoo	06:50:32
rack19	True	edwardsr	edwardsr	2-09:15:12
rack20-21	True	sdcohen	sdcohen	4-01:11:23
rack22-23	True	drenner	drenner	3-17:52:12
rack24-27	True	drenner	drenner	17:33:02

Current Status List All Running Jobs List All Done Jobs List All Jobs

Error Accounting and Daughterboard Tracking

- QOS reports:
  - List of daughterboards in machine partition (location files).
  - List of generated errors (Dbd, Wire and Ethernet).
- *qhdwcheck* wrapper stores error counters in a database.
- User reported problems via CTS.
- Web Fronts (jsp) by Joe Aronson (SUNY SB):
  http://www3.bnl.gov/qcdoc/errors

QCDOC Hardware Monitoring (Safety System)

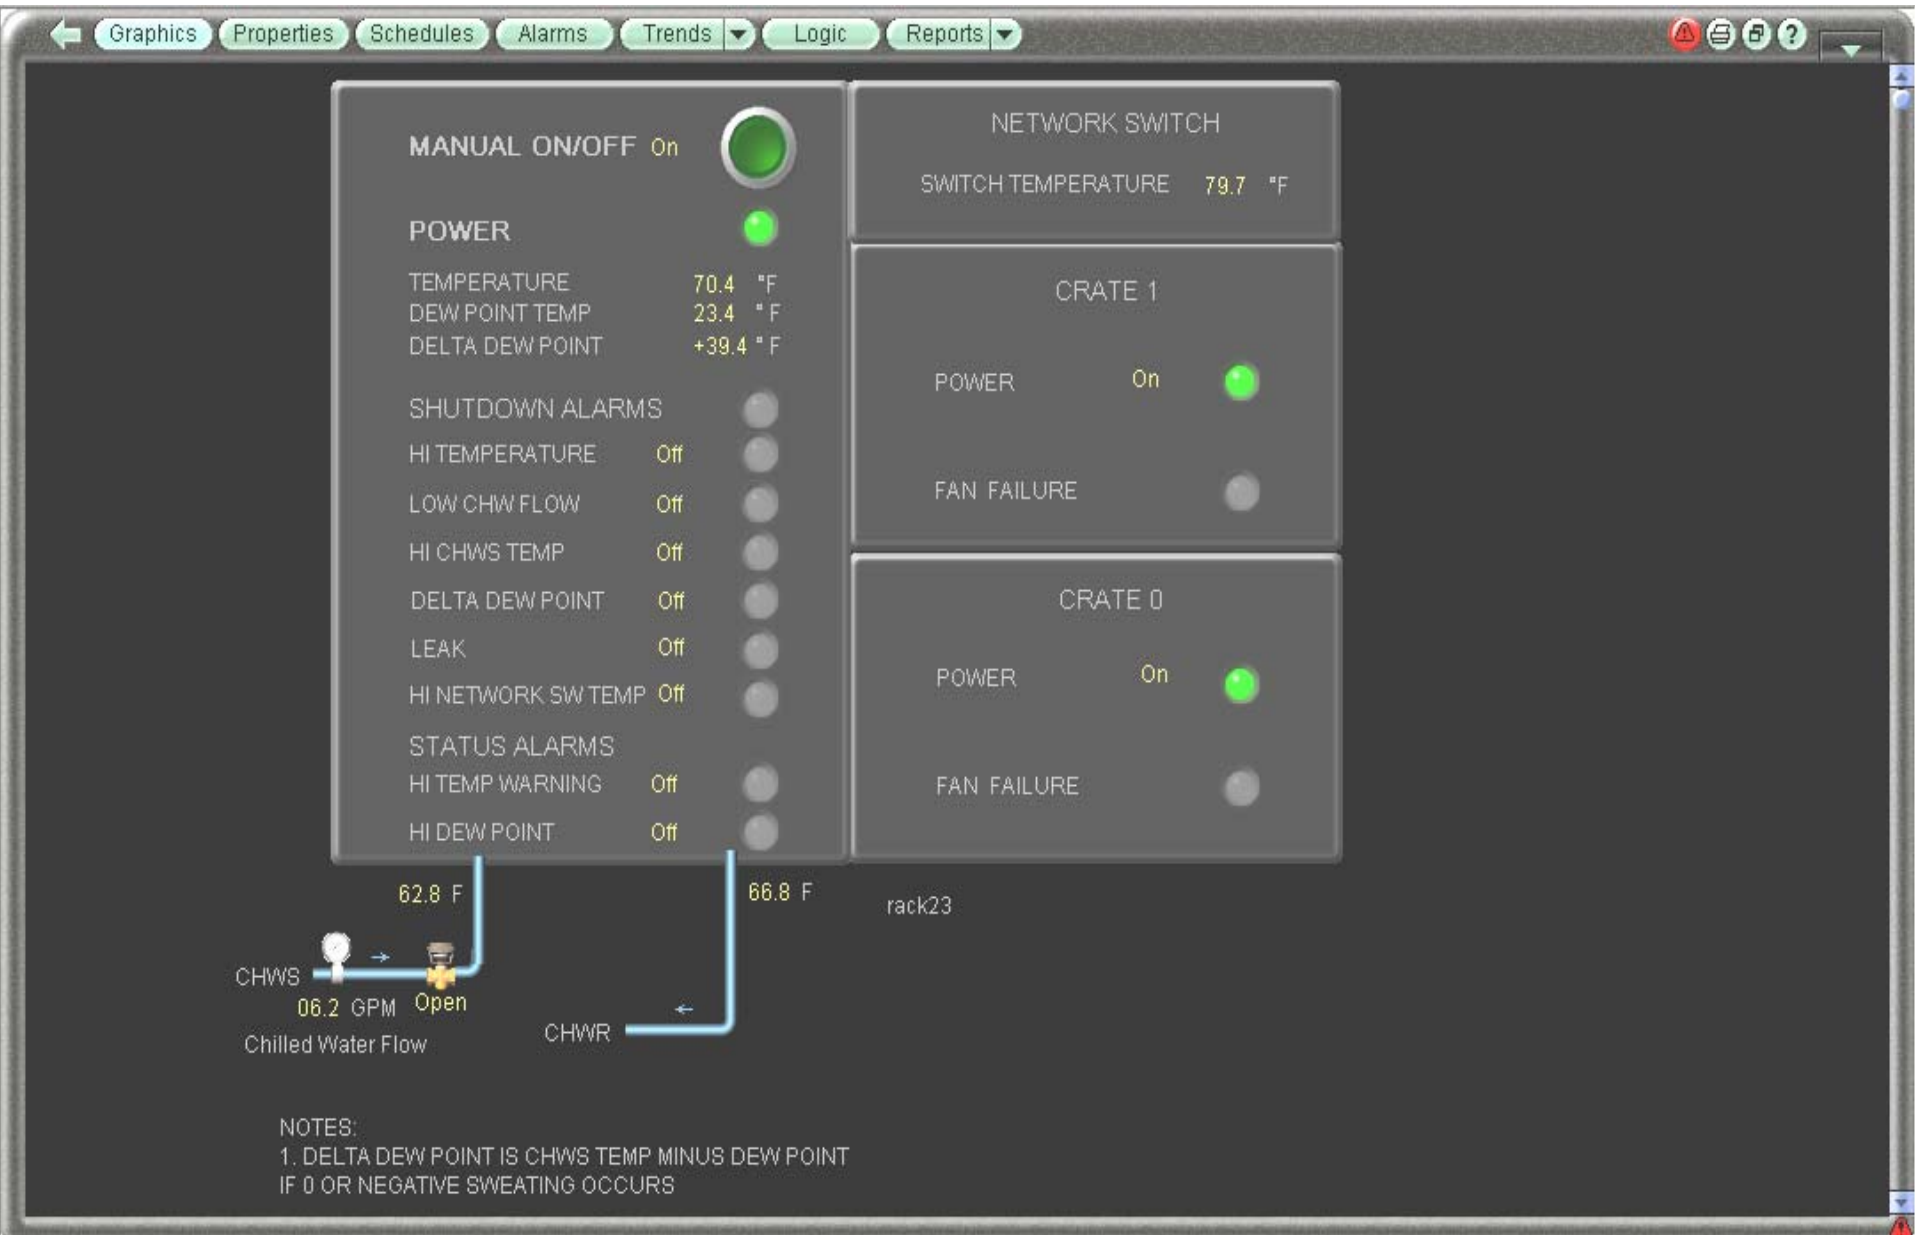

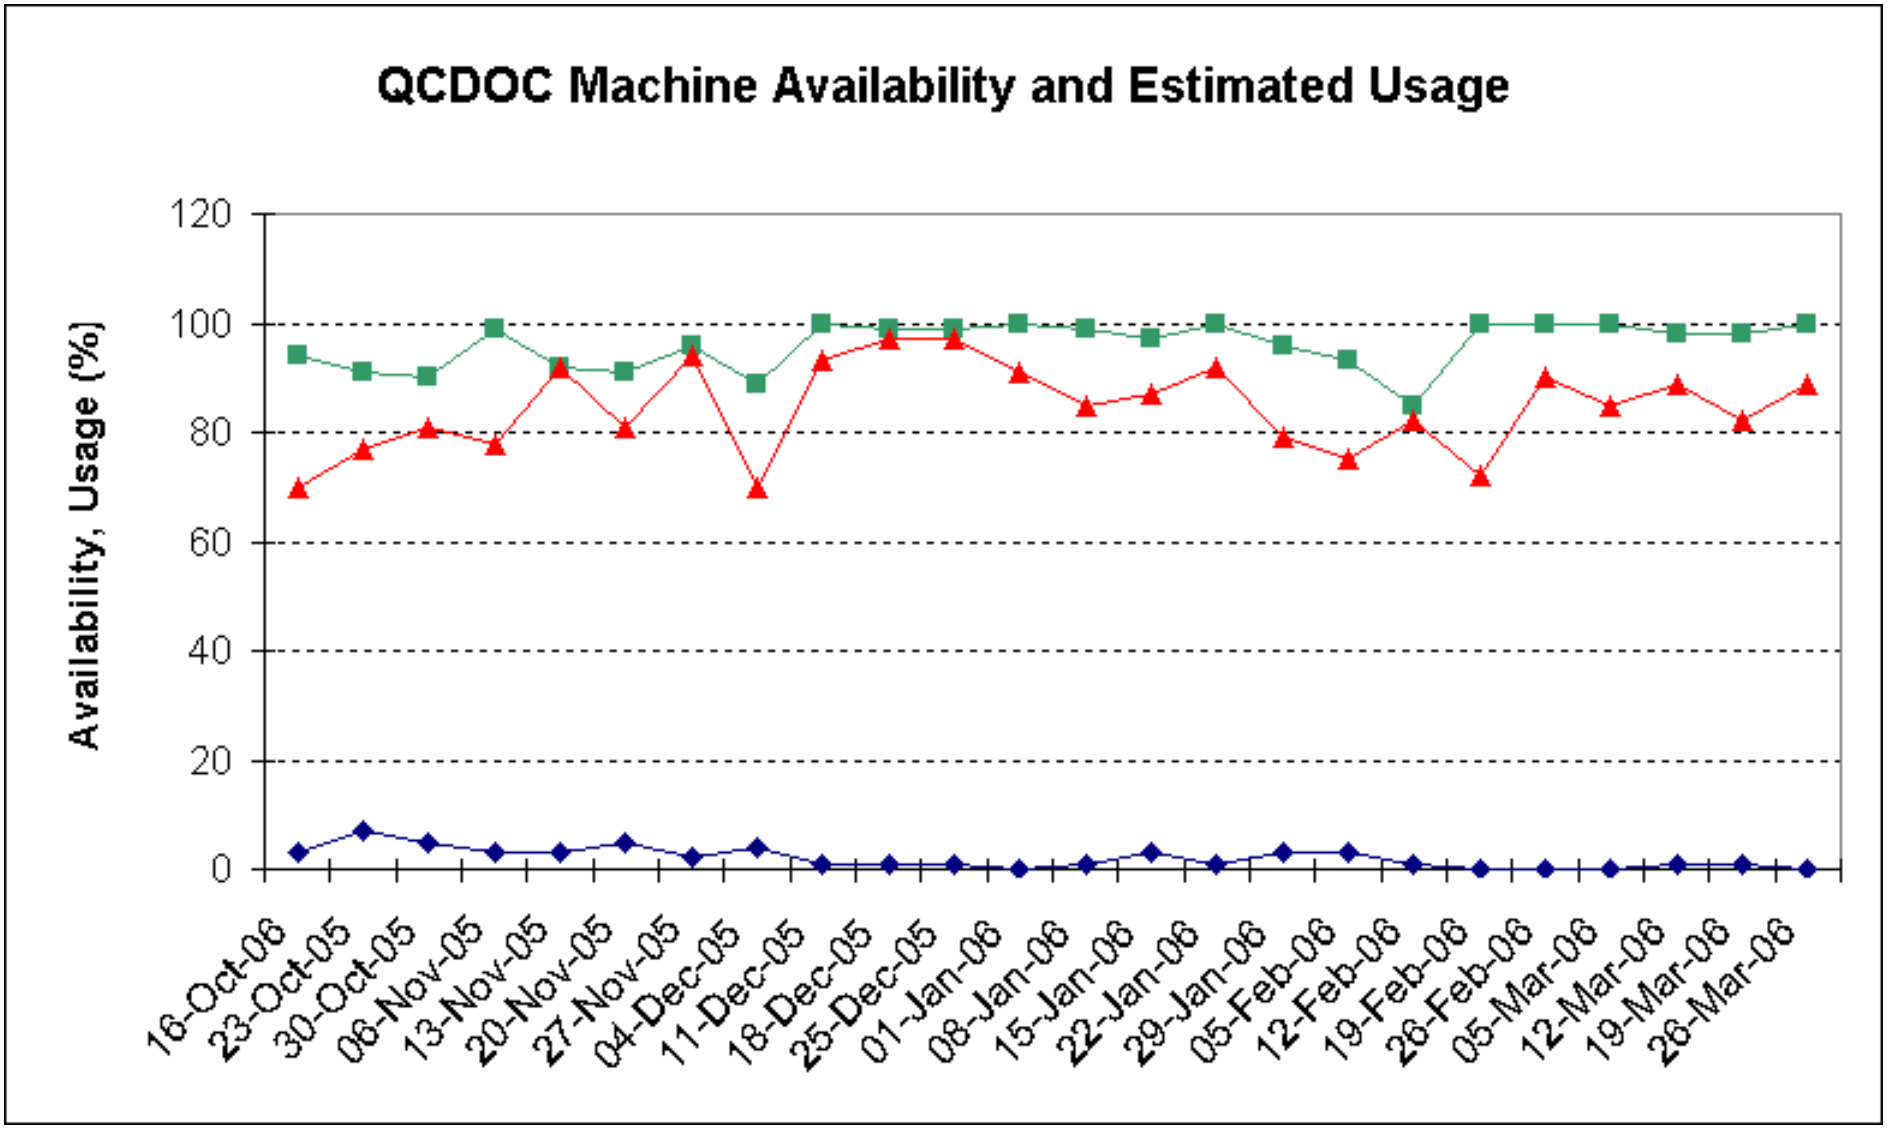
QCDOC Machine Availability and Estimated Usage
Availability, Usage (%)
120
100
80
60
40
20
0
16-Oct-06
23-Oct-05
30-Oct-05
06-Nov-05
13-Nov-05
20-Nov-05
27-Nov-05
04-Dec-05
11-Dec-05
18-Dec-05
25-Dec-05
01-Jan-06
08-Jan-06
15-Jan-06
22-Jan-06
29-Jan-06
05-Feb-06
12-Feb-06
19-Feb-06
26-Feb-06
05-Mar-06
12-Mar-06
19-Mar-06
26-Mar-06

Summary/Future Plans

- The USDOE QCDOC machine at BNL is in full production.
- Largest partition is a 4-rack machine (4096-node).
- Half of the machine (6 racks) is used by MILC producing lattice configurations at a steady rate.
- All partition nodes are running at 400MHz

- Hardware issues:
  - Lost Ethernet Contact (infreq. during booting or I/O). May recover with a software reset or may require power cycle.
  - Serial Communication Errors

- Improve our (common) runtime environment
- *qrun* wrapper
- Interactive PBS queues.
- Improve and use hardware error tracking

Conclusions

- QCDOC Web Site at BNL: http://www.bnl.gov/lqcd/
- Web Allocator: https://rbc.bnl.gov/
- Call Tracking System (CTS): https://qcdoc.phys.columbia.edu/cts
- Batch system: http://www3.bnl.gov/qcdoc/batch/pbs
- QCDOC people at BNL (Led by Bob Mawhinney)

Ed McFadden
Eric Blum
Christopher Channing
Andy Como
Paul Poleski
Joe Depace
Don Gates
Efstratios Efstathiadis
Chulwoo Jung
Enno Scholz


BROOKHAVEN
NATIONAL LABORATORY
QCD SUPERCOMPUTER CENTER
QCDSP
QCDOC