Report on QCDOC

All Hands' Meeting
US Lattice QCD Collaboration Meeting
FNAL, May 14-15, 2009

Stratos Efstathiadis
BNL

OUTLINE

- Introduction
- Available Hardware
- Machine Monitoring and Usage
- User Environment, File systems, Batch System
- User Support

QCDOC: QCD on Chip

- Optimized for LQCD calculations.

- Collaboration: Columbia University, UKQCD, Riken-BNL Research Center, SciDAC, IBM Research.

- Designed on optimizing performance/cost.

- 32-bit PowerPC 500MHz with a 64-bit FPU (1 Gflops) with good performance/Watt ratio.

- First supercomputer built using IBM's System-on-Chip technology.

- Three Large 12K-nodes machines (water cooled)
  - ❑ USDOE (BNL)
  - ❑ RBRC (BNL)
  - ❑ UKQCD (Edinburgh)

Packaging

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

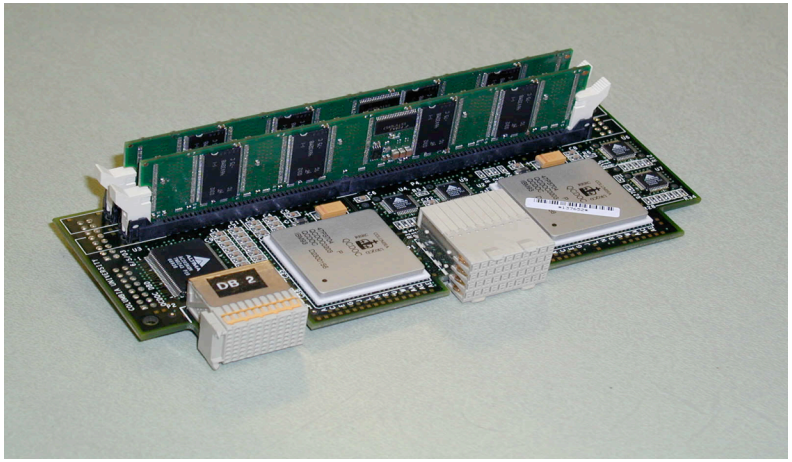

A daughterboard with two independent nodes and the vertically mounted DDR SDRAMs (128MB at BNL)

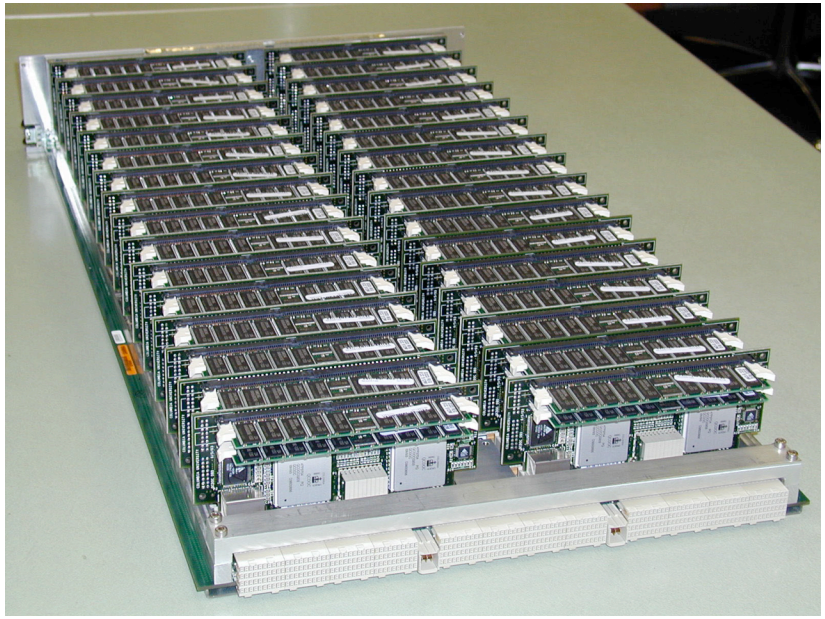

A single motherboard. Two rows of 16 daughterboard with 2 nodes each provide a total of 64 nodes. 14.5in x 27 in

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

Available Hardware

12 water cooled racks (12288 nodes)

Air Cooled Crates (1024 nodes)

Single Slot Back Plane (SSBP8 and 9)

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
acc6	True	chulwoo	chulwoo	05:26:33
rack16/crate0	True	cschmidt	cschmidt	5-07:29:07
rack16/crate1	True	cschmidt	-	-
rack17/crate0	True	petreczk	petreczk	2-17:03:06
rack17/crate1	True	cschmidt	cschmidt	2-16:39:33
rack18-19	True	petreczk	-	-
rack20	True	gfleming	gfleming	2-01:49:26
rack21	True	gfleming	gfleming	9-07:45:02
rack22	True	gfleming	gfleming	7-09:00:39
rack23	True	gfleming	gfleming	7-06:16:24
rack24-25	True	chulwoo	chulwoo	22:47:17
rack26-27	True	rzhou	rzhou	15:39:25
ssbp8-9	True	qiliu	qiliu	3-05:23:43

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

http://www3.bnl.gov/qcdoc/status/

Racks 16, 17	4 x 512-nodes	PI: Peter Petreczky
Racks 18, 19	1 x 2048 partition	PI: Peter Petreczky
Racks 20, 21, 22, 23	1x 4096 partition 4 x 1024	PI: Bob Mawhinney MILC (2 months) PI: Peter Petreczky PI: G. Fleming (01/01/2009)
Racks 24, 25, 26, 27	1 x 4096 partition	PI: Bob Mawhinney
	2 x 2048	PI: S. Sharpe

QCDOC Machine Availability and Estimated Usage
USDOE Water-Cooled Racks

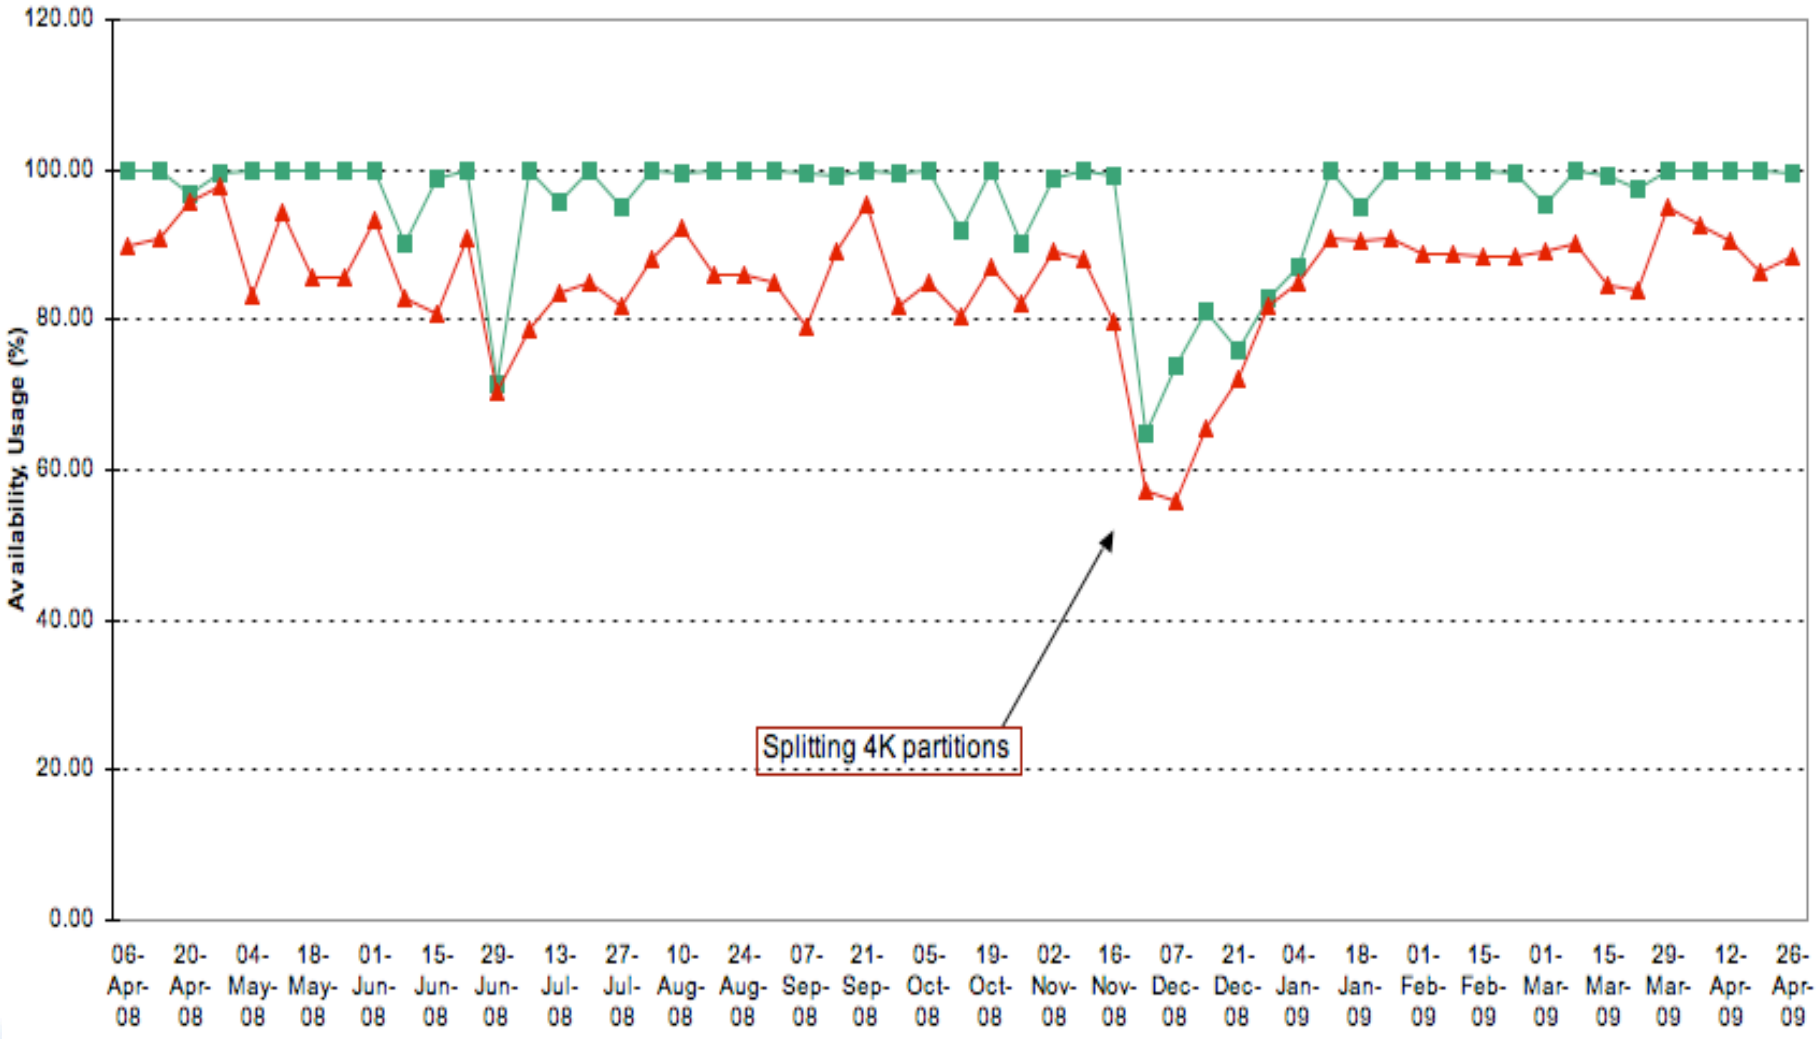

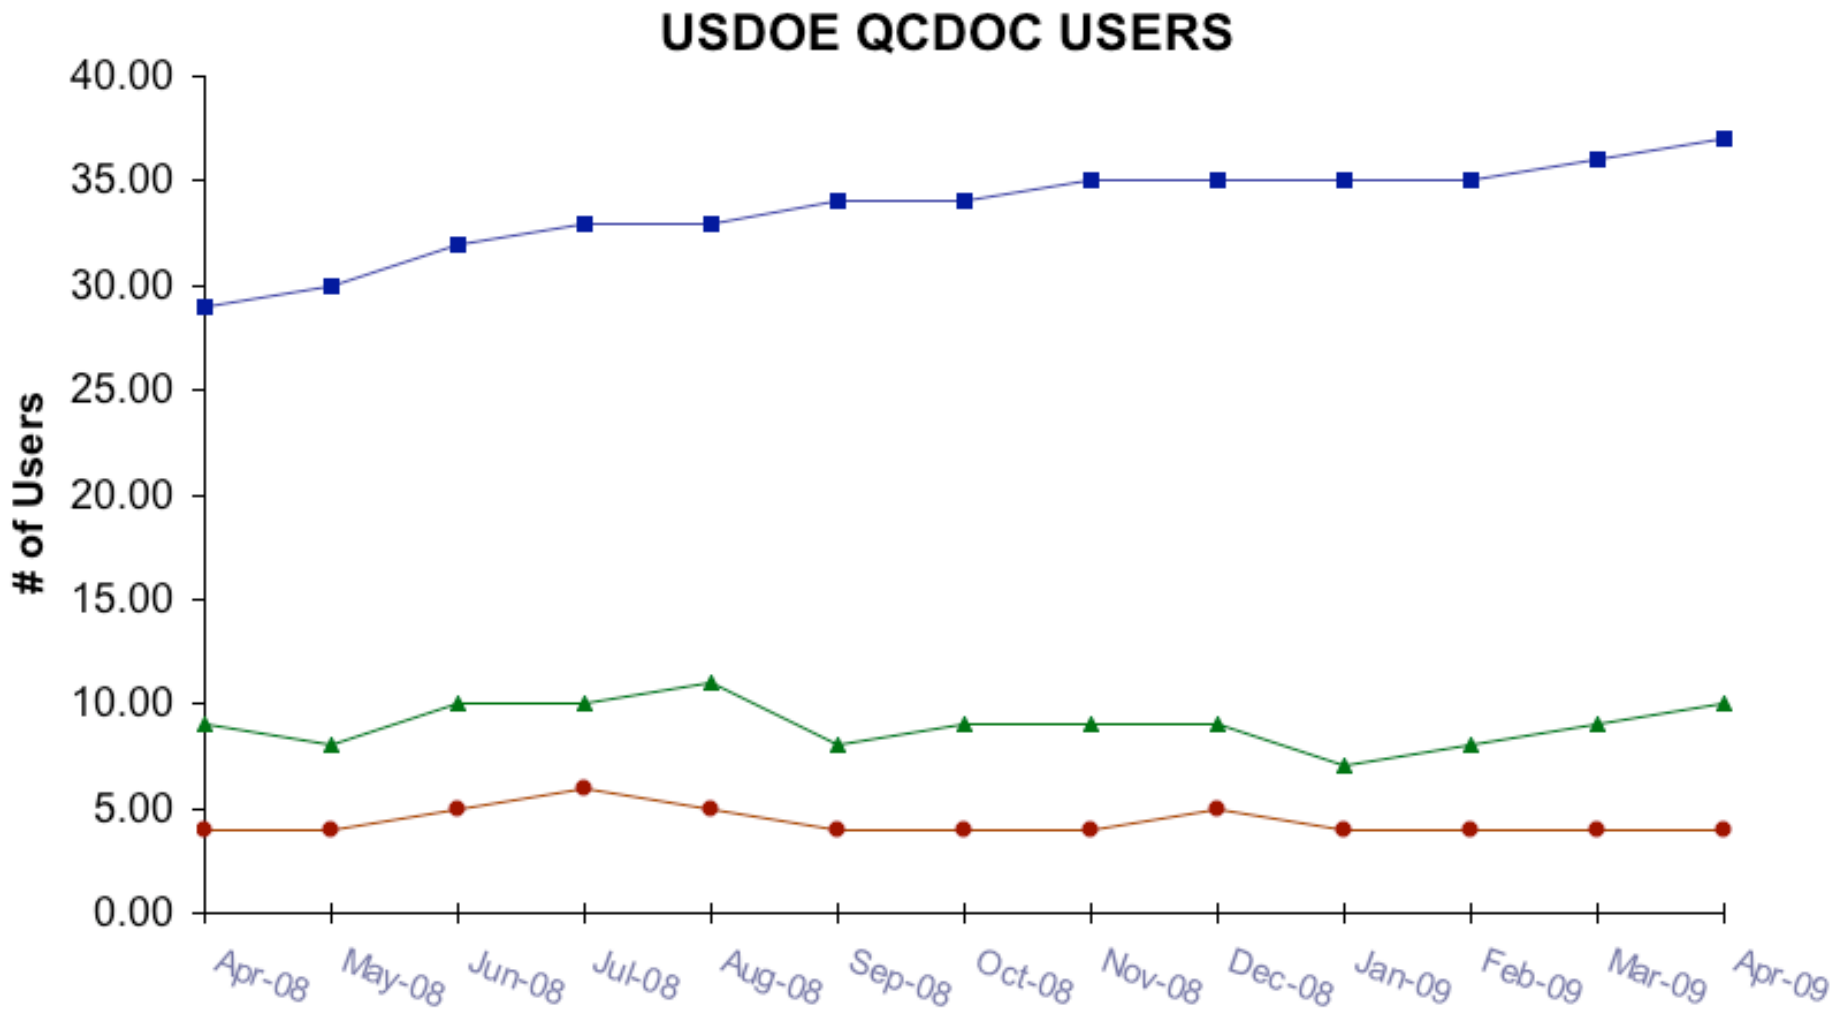

All users (since 01/2006)

Current users (WC, ACC, SSBP)

Users of Water-cooled racks

User Environment

- LQCD Computing Web Site at BNL:
  - http://lqcd.bnl.gov/comp/

- Two-factor authentication is required to access the QCDOC ssh gateways
  - ssh.qcdoc.bnl.gov (outside the BNL network)
  - ssh.qcdoc.bnl.local (inside)

  Two-factor auth. is also required to access the front-end server
  - qcdochostb.qcdoc.bnl.gov

- QDCOC user accounts now under ***centrify***

Runtime Environment

Setup Script:

source $CRE_HOME/bin/setup.(c)sh

- General purpose variables:
  - ex. $PATH, $http_proxy, etc.
- File System variables and Utilities
  - ex. $QCACHE_USER, $QCACHE_PROJECT, etc.
- Cross-compiler, Linker, Assembler etc.
  - ex. $QCC, $QCXX, $QAS, etc.
- SciDAC and Third-Party Software Env. Variables
  - ex. $**PKG**_HOME, $HOST_**PKG**_HOME, (LIBS, CFLAGS, LDFLAGS)

    PKG: QIO, QLA, QMP, LIBXML2, etc

File Systems available to QCDOC Compute Nodes

- A custom NFS client is part of the node kernel supporting two mount points (open/read/write/close).

- The Host File System
  - Globally shared by all compute nodes in a partition.
  - Provided by a disk on the front-end host (or NFS mounted on the front-end)
  - Not backed up.

- The Parallel File System (PFS)
  - Similar to cluster “scratch-disk” on every node.
  - Each node uses a unique directory, ex. */R24/C0/B0/D21/A1/*
  - Temporary data staging, not backed up.

File Systems available to QCDOC Compute Nodes

• Host and PFS file systems are provided by 2U rack-mounted LINUX NAS Servers.
  2 RAID-5 PFS file servers per machine rack (one per crate), total disk space 48TB.

• Host and PFS systems are mounted on the front-end host.

•The machine.txt file determines what Host and PFS systems are used by compute nodes in a partition.

• Env. Variable $QDATA points to the Host filesystem of a given partition ($QMACHINE):

$QDATA=/host/$QMACHINE/$USER

• For PFS systems there is a mapping between Compute Nodes and PFS directories (the layout file).

File Management Utilities

- **The Layout File**: Mapping of Compute Nodes to PFS directories

QCSH:> source $CRE_HOME/bin/qlayout.qcsh <qlayout_file>

- QIO utility wrappers:
  - qsplit: splits a single QIO file into part files
  - qscatter: moves part files into pfs systems

  - qgather: gathers part files from PFS directories into a single dir.
  - qunsplit: merges part files into a single file.
    (comes in three versions: qunsplitILDG, qunsplitSCIDAC and qunsplitDWF).

- File Management has been integrated in PBS.

- http://lqcd.bnl.gov/comp/CRE_filemanagement .html

Local Storage and File Transfers between sites

- 10 4.8TB ANACAPA file servers make up five archive/backup disk pairs.

- The five archive servers are mounted on the front-end host:
  /archive/a0 (a1, a2, a3, a4)

- Related Env. Variables:
  $QCACHE_USER=/cache/users/$USER
  $QCACHE_PROJECT=/cache/projects/<Project_Name>

- 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 requires a kerberized utility, such as *rcp*, *fscp*.

QCDOC Batch System

- Torque
- Each partition is mapped to a PBS queue (rack16/crate0, rack26-27, etc.)
- Queues with walltime limits (OneHr, FourHr, EightHr and SixteenHr) on four ACC7 MBds.
- Interactive queues (I1, I2) on ACC7 MBbs with one hour limit.
- PBS scripts (latest version at $QBATCH_HOME):
  - allocate and start up partitions
  - QIO file splitting/unsplitting
  - Check for ‘stopped’ jobs
  - Reset and powercycle racks
  - Checks for preset error limits
  - Error accounting
  - Job status notifications
  - etc.

USDOE QCDOC Batch System Status

Queue	Walltime	Running Job	Runtime	User	Queued Jobs	State
ssbp8-9	--	19477.qcdochostb	00:01:12	qiliu	-	R
rack26-27	--	19503.qcdochostb	00:06:18	rzhou	19504.qcdochostb	R
rack24-25	--	19473.qcdochostb	02:02:26	chulwoo	19498.qcdochostb 19497.qcdochostb 19494.qcdochostb	R
rack23	--	19158.qcdochostb	04:07:43	gfleming	19401.qcdochostb	R
rack22	--	19398.qcdochostb	00:15:53	gfleming	19399.qcdochostb	R
rack21	--	18981.qcdochostb	00:13:13	gfleming	19400.qcdochostb	R
rack20	--	19488.qcdochostb	00:06:08	gfleming	19489.qcdochostb	R
rack18-19	--	-	-	-	-	R
rack17/crate1	--	19476.qcdochostb	00:50:35	cschmidt	-	R
rack17/crate0	--	19483.qcdochostb	00:42:19	petreczk	-	R
rack16/crate1	--	-	-	-	-	R
rack16/crate0	--	19431.qcdochostb	01:18:10	cschmidt	-	R
acc6	--	19455.qcdochostb	00:04:10	chulwoo	19456.qcdochostb	R

http://lqcd.bnl.gov/comp/batchStatus.html

Monitoring and Accounting

- Safety System
  - Monitors water-cooled racks (chilled water temperature and flow, air temperature, humidity, power status, etc).
  - Web interface and SOAP interface for remote access (scripts: ***powerstatus***, ***powercycle***, ***poweroff***, etc.)

- Nagios
  - monitors services (nfs, ssh, etc.), load, disk space on servers (front-end, file servers, ssh gateways etc.).

- DaughterBoard Location tracking
  - based on QOS location files

- Error Accounting
  - Error counters are stored in a DB
  - Web front to the DB

- Job Tracking
  - Monitoring *qdaemon* processes on front-end.
  - Batch System logs.

User Support

- QCDOC Computing Web Site at BNL: http://lqcd.bnl.gov/comp
- Reporting Problems
  - Call Tracking System (CTS)
  - Web Front: https://qcdoc.phys.columbia.edu/cts
  - A CTS account is required.
  - Maintained by Zhihua Dong at CU
- Level of Support
  - 5X10
  - Increased Automation (powercycling scripts, PBS, etc).
- Users Mailing List (announce only)
  - qcdoc-doe-users-l@lists.bnl.gov
  - To subscribe: http://lists.bnl.gov/mailman/listinfo/qcdoc-doe-users-l

User Support

QCDOC Team at BNL (Led by Bob Mawhinney)

- Management
  - Eric Blum
    - BNL Site Mgr for the LQCD Computing Project
    - BCF Mgr
- Software
  - Efstratios Efstathiadis
  - Chulwoo Jung
  - Oliver Witzel (replaced Enno Scholtz)
- Hardware
  - Marty Gormezano (replaced Ed Brosnan 05/01/09)
  - Joe Depace
  - Robert Riccobono (replaced Don Gates 05/01/09)

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