The ASCI Q System: 30 TeraOPS Capability at Los Alamos National Laboratory

Figure 1 : The first sections of the ASCI Q 30-TeraOPS supercomputer being installed at Los Alamos National Laboratory are now up and running.

The Q supercomputing system at Los Alamos National Laboratory (LANL) is the most recent component of the Advanced Simulation and Computing (ASCI) program, a collaboration between the U. S. Department of Energy's National Nuclear Security Administration and the Sandia, Lawrence Livermore, and Los Alamos national laboratories. ASCI's mission is to create and use leading-edge capabilities in simulation and computational modeling. In an era without nuclear testing, these computational goals are vital for maintaining the safety and reliability of the nation's aging nuclear stockpile.

ASCI Q Hardware

The Q system, when complete, will include 3 segments, each providing 10 TeraOPS capability. The three segments will be able to operate independently or as a single system. One-third of the final system has been available to users for classified ASCI codes since August 2002 (with a smaller initial system available since April). This portion of the system, known as QA, comprises 1024 AlphaServer ES45 SMPs from Hewlett Packard (HP), each with 4 Alpha 21264 EV-68 processors. Each of these 4,096 CPUs has 1.25-GHz capability, creating an aggregate 10 TeraOPS.

On QA, the Linpack benchmark ran at 7.727 TeraOPS. This is 75.48% of the 10.24-TeraOPS theoretical peak of the system.

An identical segment, QB, is now being tested with unclassified scientific runs, but will soon be available for secure computing. Los Alamos has an option to purchase the third 10 TeraOPS system from HP.

The final Q system will provide 30 TeraOPS capability:

- 3072 AlphaServer ES45s from Hewlett Packard (formerly Compaq)
- 12,288 EV-68 1.25-GHz CPUs with 16-MB cache
- 33 Terabytes (TB) memory
- Gigabit fiber-channel disk drives providing 664 TB of global storage
- Dual controller accessible 72 GB drives arranged in 1536 5+1 RAID5 storage arrays, interconnected through fiber-channel switches to 384 file server nodes.

The Network - Tying together 3072 SMPs

Very integral to Q is the Quadrics (QSW) dual-rail switch interconnect, which uses a fat-tree configuration. The final switch system will include 6144 QSW PCI adapters and six 1024-way QSW federated switches, providing high bandwidth (250 Mbytes/s/rail) and low latency (~5 us). The Quadrics network enables high-performance file serving within the segments. A 6th level Quadrics network will connect the 3 segments.

The new Nicholas C. Metropolis Center for Modeling and Simulation houses Q, the ASCI 30-TeraOPS supercomputer at Los Alamos National Laboratory.

Performance on QA

Even at one-third of the final capability, performance on ASCI Q is impressive. Several ASCI codes have scaled to the full 4096 processors, and many applications have experienced significant performance increases (5-8 times faster) over previous ASCI systems. LANL will run its December 2002 ASCI Milepost calculation on the QA segment.

Supporting Q - Facilities

Q is housed in the new 303,000 sq ft Nicholas C. Metropolis Center for Modeling and Simulation. The Metropolis Center includes a 43,500 sq ft unobstructed computer room and office space for about 300 staff. In addition, it has the facilities to support air or water cooling of computers and 7.1 MW of power, expandable to 30 MW. The final 30T Q system will occupy 20,000 sq ft and use about 3 MW power. The final system will comprise about 900 cabinets for the 3072 AlphaServer ES45 SMPS and related peripherals. Cable trays 1.8 miles in length will hold about 204 miles of cable under the floor.

Supporting Q - Staff

A team of about 50 Los Alamos and HP employees supports the Q system. The work of this team involves extensive systems integration, tying together system management, networking, security, distributed resource management, data storage, applications support, development of parallel tools, user consultation, documentation, problem tracking, usage monitoring, operations, and facilities management. In addition to the Q system segments described here, this team manages several other Q-like clusters, providing additional resources to users.

For more information about ASCI Q, contact: John Morrison, CCN Division Leader (jfm@lanl.gov or 505-667-6164), James Peery, LANL Deputy Associate Director for ASCI (jspeery@lanl.gov or 505-667-0940), Manuel Vigil, Q Project Leader (mbv@lanl.gov or 505-665-1960), Ray Miller, Q Project team member (rdm@lanl.gov or 505-665-3222), or Cheryl Wampler, CCN-7 Group Leader (clw@lanl.gov or 505-667-0147)

ASCI Q is a DOE/NNSA/ASCI/HP (formerly Compaq) Partnership, operated by the Computing, Communications and Networking (CCN) Division at Los Alamos National Laboratory. http://www.lanl.gov/asci

LALP-02-0230