



T2K - Tsukuba 95 TFLOPS Cluster System at CCS

System Specifications



T2K-Tsukuba, is the latest supercomputer system installed at the Center for Computational Sciences, University of Tsukuba, and was installed in June 2008. It is a very large scale PC cluster with 10368 cores with a peak performance of 95.4 TFLOPS. To execute a wide variety of scientific computations, each computation node is configured as a fat-style PC server with 16 cores and 32 GB main memory. A large scale shared file system is provided by Lustre with 800 TB of user space under RAID6 physical file system, which is connected to all computation nodes via Infiniband to provide 16 GB/s of I/O bandwidth.

The operating system is RedHat Enterprise Linux WS5 and the entire cluster system is managed under ACE (Appro Cluster Engine), MVAPICH2 by Ohio State University, and Grid Engine batch job scheduler.

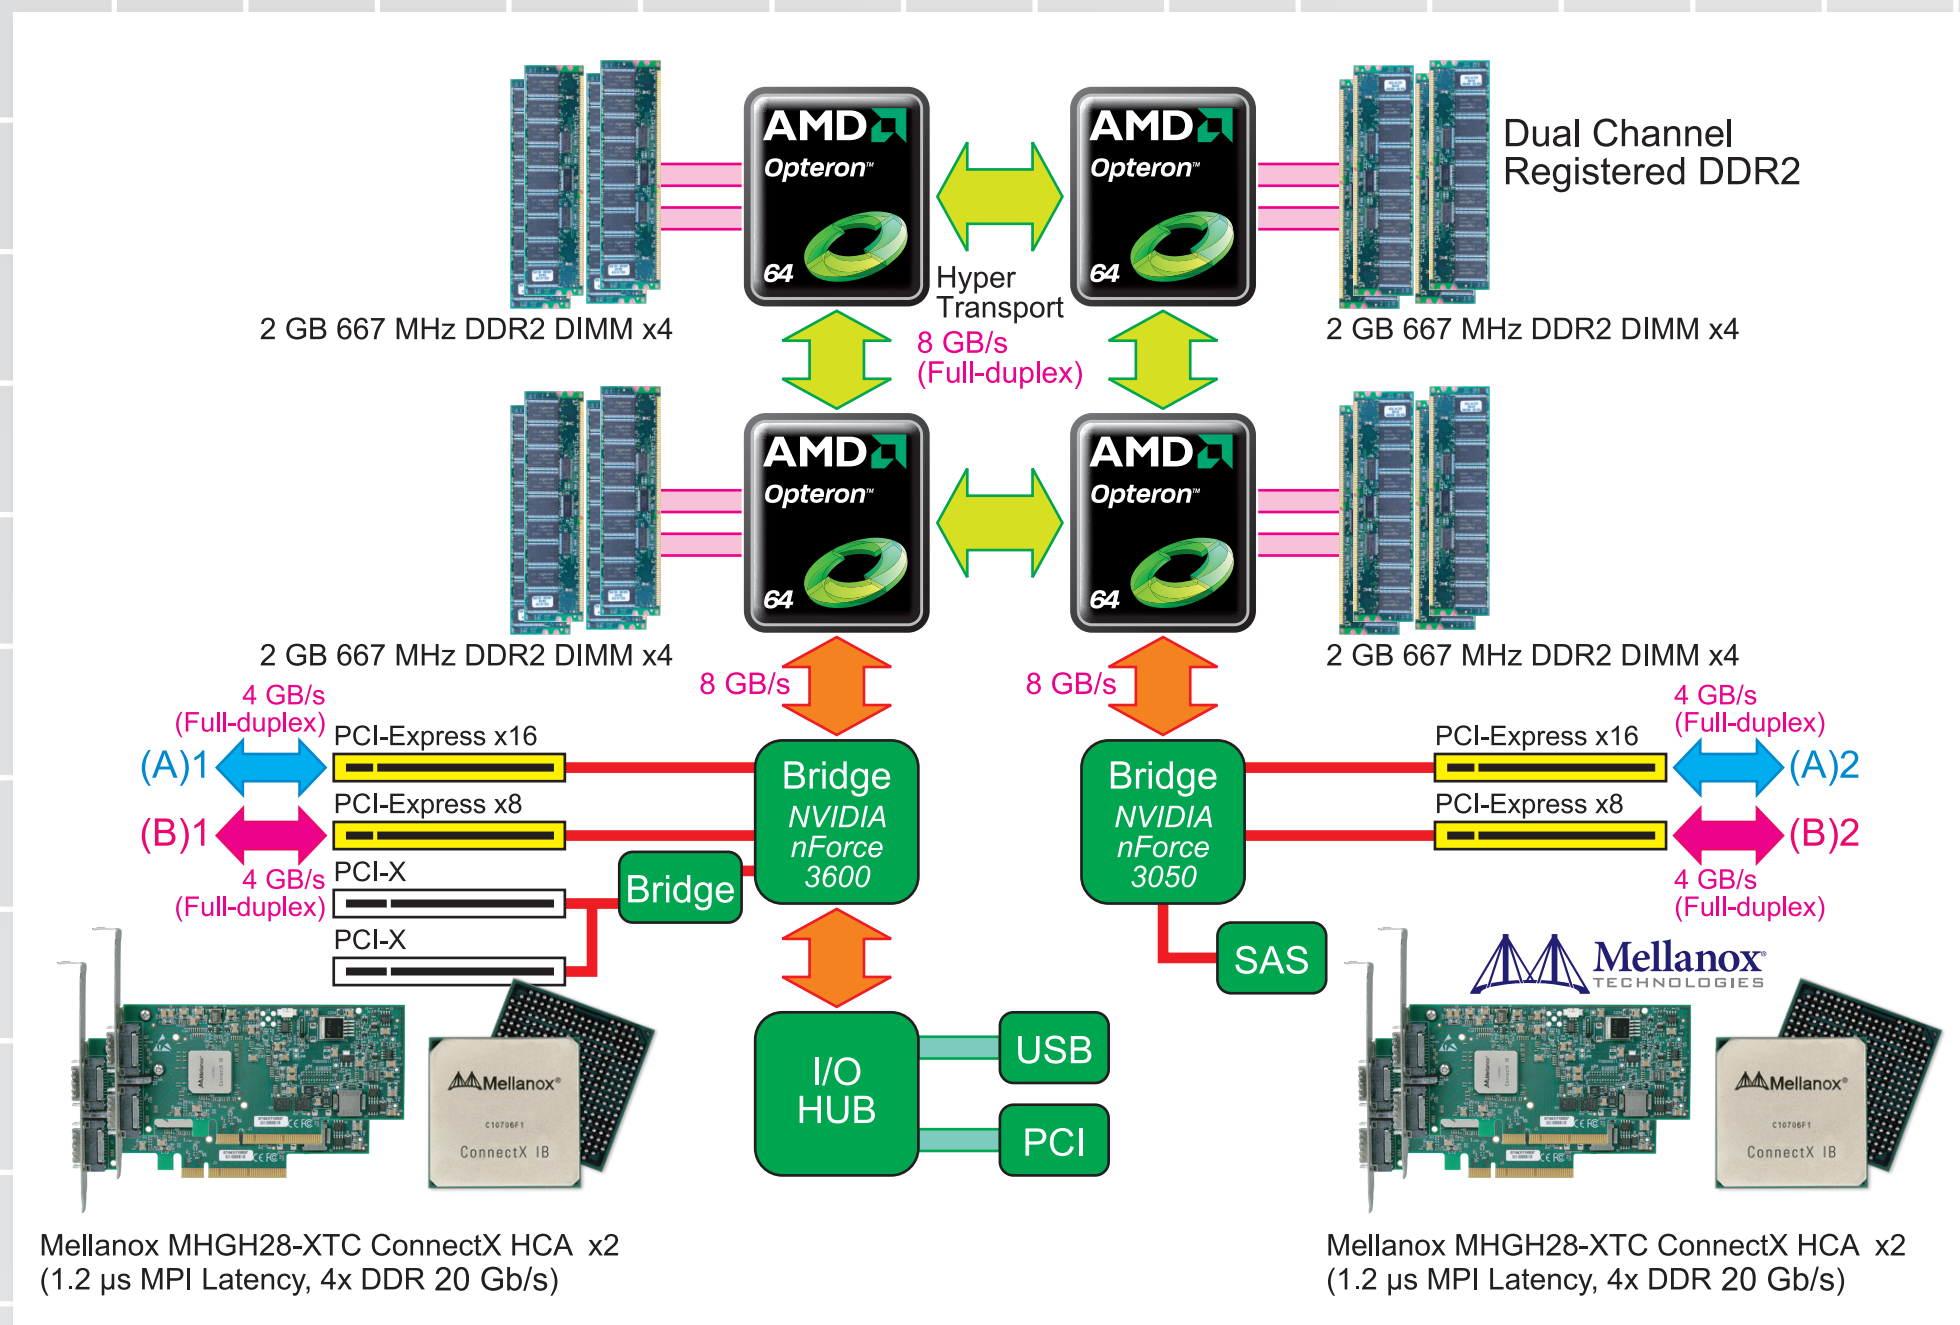
T2K-Tsukuba started its full operation from October 2008, and currently utilized for more than 30 large scale computing projects. The system is ranked 20th in TOP500 list as of June 2008. As of November 2008, T2K-Tsukuba is the largest single-system cluster in Japan.

Alliance and Vendors

T2K-Tsukuba is designed and procured under the T2K Open Supercomputer Alliance in cooperation with three Japanese universities: the University of Tsukuba, the University of Tokyo and Kyoto University.

T2K-Tsukuba is a multi-vendor system; Appro XtremeServer-X3 as computation node, Mellanox MHGH28-XTC ConnectX as Infiniband HCA, Flextronics 24-port IB switch as interconnection switch, DDN S2A9550 RAID6 as shared file server, and the entire system is integrated and maintained by Cray.

Computation Node



Each computation node is equipped with a 4-socket quad-core Opteron (AMD Barcelona, 2.3 GHz) in a 16-core shared memory configuration with a peak performance of 147.2 GFLOPS. To support such a high-performance node, a quad-rail 4xDDR Infiniband HCA is employed with 16 GB/s (bidirection) of communication performance.

Interconnection Network

Each port of quad-rail Infiniband HCA (Mellanox ConnectX 4xDDR) is connected to a full-bisection bandwidth Fat-Tree under an automatically configured fault-tolerant interconnection network. It consists of only 24-port switches and is configured as a 3-stage network. The network consists of only 616 switches and 8554 cables. This powerful interconnection network supports very large scale parallel applications with a peak performance of 95 TFLOPS.

