

High Performance Computing Palmetto Cluster

Clemson's high performance computing resources includes a "condominium" style cluster, known as Palmetto, developed to serve the university's wide-ranging research needs. Ranked #81 in the world (TOP500 list, November 2013), the system currently benchmarks at 396.7 teraflops. The system currently has 1,978 compute nodes with 19,814 total cores. The nodes are a mixture of vendors such as Dell, Sun, IBM, and HP. The system also contains 5 large memory nodes with 512GB of memory and 1 node with 2TB of memory. All nodes are connected via 1Gbps Ethernet.

The system has a combination of 10Gbps Myrinet and 56Gbps Infiniband interconnections. Recent additions to the cluster contain 20 nodes with dual Nvidia Tesla M2075 GPUs and 197 nodes with dual Nvidia Tesla K20m GPUs. The system is directly connected to a Brocade MLXe-32 switch that will provide 100Gbps bandwidth to Internet2 and 40Gbps bandwidth to the Clemson campus. The system is monitored 24x7 using Nagios and CA Spectrum. Storage for the cluster consists of 2 DDN 6620 arrays with 120 TB of raw disk storage attached to 2 Sun X4600 nodes running the SAM-QFS filesystem, and 32 Dell R510 nodes with 220TB of raw disk storage on each running the OrangeFS parallel filesystem.

Developed by Clemson Computing and Information Technology (CCIT) in collaboration with faculty researchers across the university, the Palmetto Cluster provides a shared platform that optimizes resources for the benefit of all users. Named for South Carolina's state tree, the Palmetto Cluster is designed to suit many different research applications, with a large number of powerful multi-core nodes, each with a significant amount of memory.

The Palmetto Cluster is a fully staff-supported high-performance research computing center, which is located in its own secured facility at Clemson University's research park. (CCIT) maintains this community resource, jointly funded by Clemson University and research grants secured by participating faculty that provides a cost-effective shared computing infrastructure for faculty and students.

The Palmetto Cluster is made possible through a partnership between faculty and information technology administrators. A shared funding formula makes Palmetto available to the entire Clemson community on a first come, first served basis; faculty researchers who require predictable computational availability may purchase nodes with reserved priority. The university provides additional nodes, hardware infrastructure, and system administration support. This gives individual researchers high performance computing capacity previously unavailable at Clemson, and allows the university to compete for the research grants and contracts that are vital to Clemson's goals as a major research university.

If you need guaranteed availability and have funding for purchase, the Cluster Condominium Program provides a very cost-competitive alternative to purchasing, operating and managing your own, stand-alone cluster. Condominium owners have access to additional Palmetto computational capacity as needed.

How the Cluster Condominium Program Works:

- Faculty may purchase any number of compute nodes. Once installed, the nodes become part of the cluster but the faculty member retains ownership of the purchased equipment.
- Faculty and their research collaborators always have first priority on their own equipment or the equivalent in CPU time and storage. In addition, owners have access to compute cycles available on all the other nodes in the cluster on a first-come, first-served basis.
- Unused compute cycles in the condominium are available for use by the rest of the Clemson community.
- Faculty with storage needs above the 50 GB individual quota and 50TB temporary scratch space may purchase high throughput storage by the TB.
- Clemson Computing and Information Technology provides: hardware, operating system, and system administration support; equipment room space, racks, power, and cooling; a high speed Myrinet I/O interface for each node; and research scientist support staff to assist faculty in getting their application up and running.

Clemson University subsidizes these costs in order to build High Performance Computing as a strategic university resource; in addition, vendors have provided substantial hardware discounts since the use is for university research. Any similar, non-subsidized services provided by CCIT are therefore likely to be billed at a higher cost.