


Cypress is a campus high-performance computing powerhouse

Faith Dawson

fdawson@tulane.edu



Tulane's Cypress supercomputer processes calculations much faster than a traditional computer and is available for use by any Tulane faculty member who is the principal investigator of a research group, as well as members of those research groups. (Photo by Ryan Rivet)

Cypress is Tulane University's supercomputer that is capable of processing massive quantities of data for faculty and their research groups.

According to Carl Baribault, computational scientist in Information Technology and a Cypress user, a supercomputer offers high performance computing (HPC) using parallel processing, which allows it to process calculations much faster than a

traditional computer. Baribault and his colleague, Hideki Fujioka, have hosted Cypress workshops for interested researchers.

[Cypress accounts](#) are available to any Tulanian, but they must be associated with a Tulane faculty member who is the principal investigator of a research group. Some knowledge of basic Linux commands is required in order to create and manipulate files and directories.

Data analysis from medical and scientific fields is a common use, according to [Baribault](#) and [Fujioka](#). But any digital data could work with HPC.

Cypress works with large sets of data that need to be analyzed at the same time. It has 124 compute nodes; each node is itself a compute server that includes two onboard central processing units (CPU) with 10 compute cores each, for a total of 20 compute cores per node. Also, each node has twin Xeon Phi co-processors attached via peripheral bus; each co-processor has 60 hyperthreaded compute cores. Thus, each compute node on its own is much more powerful than the computing capability that is typically available on an investigator's individual workstation, which has, for instance, eight compute cores. Moreover, the combined theoretical computing speed of Cypress's compute nodes and their attached Xeon Phi co-processors is 369 teraflops (trillions of floating-point operations — "FLOPs" — per second) and slightly more than 69 teraflops without the use of co-processors, as compared to less than 100 gigaflops (billions of floating-point operations per second) with a standard CPU. Also, all of Cypress's compute nodes have between 64 and 256 gigabytes of random-access memory (RAM), more than the 32 gigabytes at a typical workstation.

[Tiong Aw](#), microbiologist and assistant professor at the School of Public Health and Tropical Medicine, and Katie Vigil, microbiologist and Tulane PhD student, are using Cypress in identifying [water-borne pathogens](#) in marine mammals.

"Viruses are abundant in marine ecosystems, and some can cause diseases in marine mammals," Aw said. "Therefore, it is important to monitor viruses in marine mammals and their surrounding environments."

"Our lab uses Cypress to process and analyze a large amount of viral metagenomic sequencing data associated with marine mammals," Vigil said. "As next-generation sequencing (NGS) becomes more rapid, accurate and affordable, it is important to have a supercomputer such as Cypress on site to process large files (up to gigabytes) that may take days, even weeks to process on your local computer."

Cypress use and related consultation services are free for Tulane faculty and their research group members.

“I was a novice at using a Linux operating system and had to learn coding to be able to operate Cypress,” Vigil added. “With the help of Carl Baribault, I quickly learned how to run commands using (Linux-based software) Ubuntu (on our workstation and CentOS) on Cypress to process my NGS data.”

Cypress made its debut at Tulane in 2014 and has been ranked among the top 500 supercomputers in the United States. To work with Cypress, [click here for more information](#).