

Wrangler supercomputer speeds through big data

March 17 2016



Niall Gaffney, Director of Data Intensive Computing, Texas Advanced Computing Center. Credit: TACC

Handling big data can sometimes feel like driving on an unpaved road for researchers with a need for speed and supercomputers.

"When you're in the world of [data](#), there are rocks and bumps in the way, and a lot of things that you have to take care of," said Niall Gaffney, a former Hubble Space Telescope scientist who now heads the Data Intensive Computing group at the Texas Advanced Computing Center (TACC).

Gaffney led the effort to bring online a new kind of supercomputer, called Wrangler. Like the old Western cowboys who tamed wild horses, Wrangler tames beasts of big data, such as computing problems that involve analyzing thousands of files that need to be quickly opened, examined and cross-correlated.

Wrangler fills a gap in the supercomputing resources of XSEDE, the Extreme Science and Engineering Discovery Environment, supported by the National Science Foundation (NSF). XSEDE is a collection of advanced digital resources that scientists can easily use to share and analyze the massive datasets being produced in nearly every field of research today. In 2013, NSF awarded TACC and its academic partners Indiana University and the University of Chicago \$11.2 million to build and operate Wrangler, a supercomputer to handle data-intensive high performance computing.

Wrangler was designed to work closely with the Stampede supercomputer, the 10th most powerful in the world according to the bi-annual Top500 list, and the flagship of TACC at The University of Texas at Austin (UT Austin). Stampede has computed over six million jobs for open science since it came online in 2013.

"We kept a lot of what was good with systems like Stampede," said Gaffney, "but added new things to it like a very large flash storage

system, a very large distributed spinning disc storage system, and high speed network access. This allows people who have data problems that weren't being fulfilled by systems like Stampede and Lonestar to be able to do those in ways that they never could before."

Gaffney made the analogy that supercomputers like Stampede are like racing sports cars, with fantastic compute engines optimized for going fast on smooth, well-defined race-tracks. Wrangler, on the other hand, is built like a rally car to go fast on unpaved, bumpy roads with muddy gravel.

"If you take a Ferrari off-road you may want to change the way that the suspension is done," Gaffney said. "You want to change the way that the entire car is put together, even though it uses the same components, to build something suitable for people who have a different job."

At the heart of Wrangler lie 600 terabytes of flash memory shared via PCI interconnect across Wrangler's over 3,000 Haswell compute cores. "All parts of the system can access the same storage," Gaffney said. "They can work in parallel together on the data that are stored inside this high-speed storage system to get larger results they couldn't get otherwise."

This massive amount of flash storage comes from DSSD, a startup co-founded by Andy Bechtolsheim of Sun Microsystems fame and acquired in May of 2015 by EMC. Bechtolsheim's influence at TACC goes back to the 'Magnum' Infiniband network switch he led design on for the now-decommissioned Ranger supercomputer, the predecessor to Stampede.

What's new is that DSSD took a shortcut between the CPU and the data. "The connection from the brain of the computer goes directly to the storage system. There's no translation in between," Gaffney said. "It actually allows people to compute directly with some of the fastest

storage that you can get your hands on, with no bottlenecks in between."

Speeding up the gene analysis pipeline

Gaffney recalled the hang-up scientists had with code called OrthoMCL, which combs through DNA sequences to find common genetic ancestry in seemingly unrelated species. The problem was that OrthoMCL let loose databases wild as a bucking bronco.

"It generates a very large database and then runs computational programs outside and has to interact with this database," said biologist Rebecca Young of the Department of Integrative Biology and the Center for Computational Biology and Bioinformatics at UT Austin. She added, "That's not what Lonestar and Stampede and some of the other TACC resources were set up for."

Young recounted how at first, using OrthoMCL with online resources, she was only able to pull out 350 comparable genes across 10 species. "When I run OrthoMCL on Wrangler, I'm able to get almost 2,000 genes that are comparable across the species," Young said. "This is an enormous improvement from what is already available. What we're looking to do with OrthoMCL is to allow us to make an increasing number of comparisons across species when we're looking at these very divergent, these very ancient species separated by 450 million years of evolution."

"We were able to go through all of these work cases in anywhere between 15 minutes and 6 hours," Gaffney said. "This is a game changer."

Gaffney added that getting results quickly lets scientists explore new and deeper questions by working with larger collections of data and driving previously unattainable discoveries.

Tuning energy efficiency in buildings

Computer scientist Joshua New with the Oak Ridge National Laboratory (ORNL) hopes to take advantage of Wrangler's ability to tame [big data](#). New is the principal investigator of the Autotune project, which creates a software version of a building and calibrates the model with over 3,000 different data inputs from sources like utility bills to generate useful information such as what an optimal energy-efficient retrofit might be.

"Wrangler has enough horsepower that we can run some very large studies and get meaningful results in a single run," New said. He currently uses the Titan supercomputer of ORNL to run 500,000 simulations and write 45 TB of data to disk in 68 minutes. He said he wants to scale out his parametric studies to simulate all 125.1 million buildings in the U.S.

"I think that Wrangler fills a specific niche for us in that we're turning our analysis into an end-to-end workflow, where we define what parameters we want to vary," New said. "It creates the sampling matrix. It creates the input files. It does the computationally challenging task of running all the simulations in parallel. It creates the output. Then we run our artificial intelligence and statistic techniques to analyze that data on the back end. Doing that from beginning to end as a solid workflow on Wrangler is something that we're very excited about."

When Gaffney talks about storage on Wrangler, he's talking about a lot of data storage—a 10 petabyte Lustre-based file system hosted at TACC and replicated at Indiana University. "We want to preserve data," Gaffney said. "The system for Wrangler has been set up for making data a first-class citizen amongst what people do for research, allowing one to hold onto data and curate, share, and work with people with it. Those are the founding tenants of what we wanted to do with Wrangler."

Shedding light on dark energy

"Data is really the biggest challenge with our project," said UT Austin astronomer Steve Finkelstein. His NSF-funded project is called HETDEX, the Hobby-Eberly Telescope Dark Energy Experiment. It's the largest survey of galaxies ever attempted. Scientists expect HETDEX to map over a million galaxies in three dimensions, in the process discovering thousands of new galaxies. The main goal is to study dark energy, a mysterious force pushing galaxies apart.

"Every single night that we observe—and we plan to observe more or less every single night for at least three years—we're going to make 200 GB of data," Finkelstein said. It'll measure the spectra of 34,000 points of skylight every six minutes.

"On Wrangler is our pipeline," Finkelstein said. "It's going to live there. As the data comes in, it's going to have a little routine that basically looks for new data, and as it comes in every six minutes or so it will process it. By the end of the night it will actually be able to take all the data together to find new galaxies."

Human origins buried in fossil data

Another example of a new HPC user Wrangler enables is an NSF-funded science initiative called PaleoCore. It hopes to take advantage of Wrangler's swiftness with databases to build a repository for scientists to dig through geospatially-aware data on all fossils related to human origins. This would combine older digital collections in formats like Excel worksheets and SQL databases with newer ways of gathering data such as real-time fossil GPS information collected from iPhones or iPads.

"We're looking at big opportunities in linked [open data](#)," PaleoCore principal investigator Denne Reed said. Reed is an associate professor in the Department of Anthropology at UT Austin.

Linked open data allows for queries to get meaning from the relationships of seemingly disparate pieces of data. "Wrangler is the type of platform that enables that," Reed said. "It enables us to store large amounts of data, both in terms of photo imagery, satellite imagery and related things that go along with geospatial data. Then also, it allows us to start looking at ways to effectively link those data with other data repositories in real time."

Data analytics for science

Wrangler's shared memory supports data analytics on the Hadoop and Apache Spark frameworks. "Hadoop is a big buzzword in all of data science at this point," Gaffney said. "We have all of that and are able to configure the system to be able to essentially be like the Google Search engines are today in data centers. The big difference is that we are servicing a few people at a time, as opposed to Google."

Users bring data in and out of Wrangler in one of the fastest ways possible. Wrangler connects to Internet2, an optical network which provides 100 gigabytes per second worth of throughput to most of the other academic institutions around the country.

What's more, TACC has tools and techniques to transfer their data in parallel. "It's sort of like being at the supermarket," explained Gaffney. "If there's only one lane open, it is just as fast as one person checking you out. But if you go in and have 15 lanes open, you can spread that traffic across and get more people through in less time."

A new user community for supercomputers

Biologists, astronomers, energy efficiency experts, and paleontologists are just a small slice of the new user community Wrangler aims to attract.

Wrangler is also more web-enabled than typically found in [high performance computing](#). A web portal allows users to manage the system and gives the ability to use web interfaces such as VNC, RStudio, and Jupyter Notebooks to support more desktop-like user interactions with the system.

"We need these bigger systems for science," Gaffney said. "We need more kinds of systems. And we need more kinds of users. That's where we're pushing towards with these sort of portals. This is going to be the new face, I believe, for many of these systems that we're moving forward with now. Much more web-driven, much more graphical, much less command line driven. "

"The NSF shares with TACC great pride in Wrangler's continuing delivery of world-leading technical throughput performance as an operational resource available to the open science community in specific characteristics most responsive to advance data-focused research," said Robert Chadduck, the program officer overseeing the NSF award.

Wrangler is primed to lead the way in computing the bumpy world of data-intensive science research. "There are some great systems and great researchers out there who are doing groundbreaking and very important work on data, to change the way we live and to change the world," Gaffney said. "Wrangler is pushing forth on the sharing of these results, so that everybody can see what's going on."

Provided by University of Texas at Austin

Citation: Wrangler supercomputer speeds through big data (2016, March 17) retrieved 27 April 2024 from <https://phys.org/news/2016-03-wrangler-supercomputer-big.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.