

A supercomputer for the 'long tail' of science

October 20 2015



SDSC's new Comet supercomputer, a petascale cluster funded by the NSF, is designed to run as many as 10,000 research jobs simultaneously. Credit: Ben Tolo, SDSC

The San Diego Supercomputer Center (SDSC) at the University of

California, San Diego this week formally launched "Comet," a new petascale supercomputer designed to transform scientific research by expanding computational access among a larger number of researchers and across a wider range of domains.

"The launch of Comet marks yet another stage in SDSC's leadership in the national cyberinfrastructure ecosystem," said Jim Kurose, assistant director of the National Science Foundation for Computer and Information Science and Engineering (CISE), during remarks Oct. 14 at the SDSC event. "Through this launch and the extraordinary computing capabilities of SDSC, the Center will continue to expand the frontiers of science and engineering, allowing researchers to open new windows into phenomena as vast as the universe and as small as nanoparticles."

The result of an NSF award valued at roughly \$24 million, including hardware and operating funds, Comet is designed to meet the emerging requirements often referred to as the "long tail" of science—the idea that the large number of modest-sized computationally based research projects represent, in aggregate, a tremendous amount of research that can yield scientific advances and discovery.

Comet joins SDSC's Gordon supercomputer as another key resource within the NSF's XSEDE (eXtreme Science and Engineering Discovery Environment) computer resource-sharing system, which comprises the most advanced collection of integrated digital resources and services in the world.

In his remarks, Kurose noted that cyberinfrastructure—defined as a dynamic ecosystem consisting of advanced computing systems, data, software and, most importantly, people, all linked by high-speed networks—has increasingly become a critical component of discovery and innovation.

"It is essential for accelerating the pace of discovery and innovation in all fields of inquiry...and for the progress of science, engineering, national competitiveness, health and security," he said.

SDSC used the formal launch of Comet to also celebrate 30 years as a national resource for advanced computation.

"SDSC has been at the forefront of cyberinfrastructure since the center's beginning in 1985, as one of NSF's original [supercomputer](#) centers," Kurose said. "It has been a leader of the data science revolution as well. The work done at SDSC has not only provided insights into fundamental science but it has also included a strong dedication to K-12 educational outreach."

How SDSC's "Comet" Supercomputer Serves Science and Society

Comet is configured to help transform advanced computing by expanding access and capacity not only among research domains that typically rely on HPC—such as chemistry and biophysics—but among domains which are relatively new to supercomputers, such as genomics, finance and the social sciences. Some of the domains already being served by Comet include:

Astrophysics: Supercomputers can greatly accelerate timescales for researching the origins of the universe.

Neurosciences, Brain Research: SDSC's Neuroscience Gateways project will contribute to the national BRAIN Initiative announced by the Obama Administration to deepen our understanding of the human brain.

Social Sciences: Sociologists and political scientists are analyzing newly

accessible data sets to study censorship of the press, factors that affect participation in the political process, and the properties of social networks.

Molecular Science: Studying the properties of lipids, proteins, nucleic acids, and small molecules can advance our understanding of biophysical processes at the atomic scale, leading to new drug designs and reducing disease.

DNA Nanostructures: Conducting nanoscale biomolecular research could lead to low-cost DNA sequencing technologies, and in turn create targeted drug delivery systems and help explain the molecular causes of disease.

Alternative Energy Solutions/New Materials Research: Finding new and more efficient solutions to energy harvesting, nanoporous membranes for water desalinization, solar thermal fuels and more.

Fluid Turbulent Physics: Supercomputers can create highly detailed simulations to track ocean currents or improve industry methods related to the discharge of pollutants or oil flow in pipelines.

Climate Change/Environmental Sciences: Modeling atmospheric aerosols, identified as influencing the chemical composition and radiative balance of the troposphere, has direct implications for our climate and public health.

Seismic Research/Disaster Prevention: Keys to hazard management for major earthquakes, hurricanes, and wildfires include the ability to predict a wide range of possibilities. Supercomputer-generated simulations are used to inform decision-making strategies.

The Tree of Life: Biologists construct phylogenetic trees to capture the

evolutionary relationship between species, and help us better understand the functions and interactions of genes, the origin and spread of diseases, the co-evolution of hosts and parasites and migration of human populations.

Key Features of Comet:

- ~2 petaflops of overall peak performance—one million billion operations or calculations per second.
- Dell compute nodes using next-generation Intel Xeon processors, 27 racks of compute nodes totaling 1,944 nodes or 46,656 cores.
- 128 gigabytes of dynamic RAM and 320 GB of flash memory per standard compute node.
- 72 nodes per rack with full bisection InfiniBand FDR interconnect in each rack, and a 4:1 bisection cross-rack interconnect.
- Additional GPU and large-memory (1.5 Terabytes) nodes for applications such as visualization, molecular dynamics simulations, or de novo genome assembly.
- 7 petabytes of Lustre-based high-performance storage from Aeon, and 6 petabytes of durable storage for data reliability.
- First XSEDE production system to support high-performance virtualization.

Provided by National Science Foundation

Citation: A supercomputer for the 'long tail' of science (2015, October 20) retrieved 20 April 2024 from <https://phys.org/news/2015-10-supercomputer-tail-science.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.