

UC San Diego

UC San Diego News Center

November 13, 2014 | By Jan Zverina

SDSC to Start Loading the Nodes for its New “Comet” Supercomputer

New Petascale System Slated for Launch in Early 2015



Comet, a new petascale supercomputer designed to transform advanced scientific computing by expanding access and capacity among traditional as well as non-traditional research domains, will soon be taking shape at the San Diego Supercomputer Center (SDSC) at the University of California, San Diego.

The result of a National Science Foundation (NSF) award currently valued at \$21.6 million including hardware and operating funds, *Comet* will be capable of an overall peak performance of two petaflops, or two quadrillion operations per second. *Comet* will join [SDSC's Gordon](#)

supercomputer as another key resource within the NSF's XSEDE (Extreme Science and Engineering Discovery Environment) program, which comprises the most advanced collection of integrated digital resources and services in the world.

Researchers can apply for time on *Comet* and other resources via XSEDE. *Comet*'s production startup is scheduled for early 2015, to be followed by a formal launch event in the spring.

Gateway to Discovery

“*Comet* is really all about providing high-performance computing to a much larger research community – what we call ‘HPC for the 99 percent’ – and serving as a gateway to discovery,” said SDSC Director Michael Norman, the project’s principal investigator. “*Comet* has been specifically configured to meet the needs of underserved researchers in domains that have not traditionally relied on supercomputers to help solve problems, as opposed to the way such systems have historically been used.”

Comet was designed to provide a solution for, the emerging research requirements often referred to as the ‘long tail’ of science, which describes the idea that the large number of modest-sized computationally-based research projects still represents, in aggregate, a tremendous amount of research and resulting scientific impact and advance.

“For example, in 2012, 99 percent of jobs run on NSF’s HPC resources used less than 2,048 cores and those jobs consumed more than half of all core-hours,” said Norman. “So one of our key strategies for *Comet* has been to support modest-scale users across the entire spectrum of NSF communities, while also welcoming those research communities that are not typically users of more traditional HPC systems, such as genomics, the social sciences, and economics.”

Key Features of Comet

- Dell-integrated cluster based on the Intel® Xeon® Processor E5-2600 v3 family (two processors per node.)
- Estimated overall peak performance of two petaflops – two quadrillion operations per second
- Designed to optimize capacity for modest-scale jobs: Each 72-node rack (1,728 cores) features full bisection InfiniBand FDR interconnect from Mellanox, with a 4:1 bisection interconnect across the racks. Total node count is 1,944 or 46,656 cores.
- Total 253 TB DDR4 RAM and 620 TB of flash memory.
- Four large-memory nodes (1.5 TB of memory per), plus 36 GPU nodes to accommodate applications such as visualizations, molecular dynamics simulations, or de novo genome assembly.
- 7.6 PB of Lustre-based high-performance storage; plus 6 PB of durable storage for data reliability.

- 100 Gbps connectivity to Internet2 and ESNet, allowing users to rapidly move data to SDSC for analysis and data sharing and return data to their institutions for local use.
- Comet is the first XSEDE production system to support high-performance Single Root I/O Virtualization at the multi-node cluster level.

Build-out Begins

Comet will be the successor to SDSC's *Trestles* computer cluster, to be decommissioned when *Comet* comes online.

"Think of *Comet* as our *Trestles* cluster on steroids," said SDSC Deputy Director Richard Moore, a co-PI of the *Comet* project. "It will have all of the features that made *Trestles* popular with users, but with significantly more capacity to appeal to a broader base of researchers, while providing ease-of-access and minimal wait times."

Comet will be a Dell-integrated cluster using Intel's Xeon® Processor E5-2600 v3 family, with two processors per node and 12 cores per processor running at 2.5GHz. Each compute node will have 128 GB (gigabytes) of traditional DRAM, and 320 GB of local flash memory. Since *Comet* is designed to optimize capacity for modest-scale jobs, each rack of 72 nodes (1,728 cores) will have a full bisection InfiniBand FDR interconnect from Mellanox, with a 4:1 over-subscription across the racks. There will be a total of 27 racks of these compute nodes, totaling 1,944 nodes or 46,656 cores.

In addition, *Comet* will include four large-memory nodes, each with four processors and 1.5 TB of memory, as well as 36 GPU nodes, each with four NVIDIA GPUs (graphic processing units). The GPU and large-memory nodes will target specific applications such as visualizations, molecular dynamics simulations, *orde novogenome* assembly.

Comet users will also have access to 7.6 PB (petabytes) of Lustre-based high-performance storage, with 200 GBps bandwidth to the cluster. It is based on an evolution of SDSC's *Data Oasis* storage system, with Aeon Computing as the primary storage vendor. The system will feature a new 100 Gbps (Gigabit per second) connectivity to Internet2 and ESNet, allowing users to rapidly move data to SDSC for analysis and data sharing, and to return data to their institutions for local use.

'Secret Sauce'

By the summer of 2015, *Comet* will be the first XSEDE production system to support high-performance virtualization at the multi-node cluster level. *Comet's* use of Single Root I/O Virtualization (SR-IOV) means researchers can use their own software environment, as they do with cloud computing, but can achieve the high performance they expect from a supercomputer.

“We are pioneering the area of virtualized clusters, specifically with SR-IOV,” said Philip Papadopoulos, SDSC’s chief technical officer. “This will allow virtual sub-clusters to run applications over InfiniBand at near-native speeds, representing a huge step forward in HPC virtualization. In fact the new ‘secret sauce’ in *Comet* is virtualization for customized software stacks, which will lower the entry barrier for a wide range of researchers.”

“The variety of hardware and support for complex, customized software environments will be of particular benefit to Science Gateway developers,” said Nancy Wilkins-Diehr, co-PI of the XSEDE program and SDSC’s associate director. “We now have more than 30 such Science Gateways running on XSEDE, each designed to address the computational needs of a particular community such as computational chemistry, atmospheric science or the social sciences.”

SDSC team members plan to work closely with communities and enable them to develop the customized software stacks that meet their needs by defining virtual clusters. With significant advances in SR-IOV, virtual clusters will be able to attain near native hardware performance in both latency and bandwidth, making them suitable for MPI-style parallel computing.

SDSC’s *Comet* program is funded under NSF grant number ACI 1341698. Other SDSC staff includes Distinguished Scientist Wayne Pfeiffer, in addition to Diane Baxter, Amit Majumdar, Mahidhar Tatineni, and Robert Sinkovits, and Rick Wagner. Geoffrey Fox, Distinguished Professor of Computer Science and Informatics at Indiana University and PI of the NSF’s FutureGrid project, is a strategic partner in the project.

MEDIA CONTACT

Jan Zverina, 858-534-5111, jzverina@sdsc.edu

Warren R. Froelich, 858-822-3622, froelich@sdsc.edu

UC San Diego’s [Studio Ten 300](#) offers radio and television connections for media interviews with our faculty, which can be coordinated via studio@ucsd.edu. To connect with a UC San Diego faculty expert on relevant issues and trending news stories, visit <https://ucsdnews.ucsd.edu/media-resources/faculty-experts>.