

December 08, 2020 | By Jan Zverina

## SDSC's 'Expanse' Supercomputer Formally Enters Production


**National Science Foundation-funded system now available to academic & industry researchers**

The San Diego Supercomputer Center (SDSC) at UC San Diego announced that its new *Expanse* supercomputer formally entered service for researchers following a program review by the National Science Foundation (NSF), which awarded SDSC a grant in mid-2019 to build the innovative system.

At over twice the performance of *Comet*, SDSC's current petascale supercomputer, *Expanse* supports SDSC's theme of 'Computing without Boundaries' with powerful CPUs, GPUs, and a data-centric architecture that supports a wide range of scientific workloads including experimental facilities, edge computing, and public clouds.

"The name of our new system says it all," said SDSC Director Michael Norman, the principal investigator (PI) for *Expanse* and a computational astrophysicist. "With innovations in cloud integration and other features such as composable systems, as well as continued support for science gateways and distributed computing via [the Open Science Grid \(OSG\)](#), *Expanse* will allow researchers to push the boundaries of computing and substantially reduce their times to discovery."

A key innovation of *Expanse* is its ability to support composable systems, which can be described as the integration of computing elements such as a combination of CPU, GPU, and other resources into scientific workflows that may include data acquisition and processing,



*The all-new Expanse supercomputer at the San Diego Supercomputer Center (SDSC) on the University of California San Diego campus. Funded by the National Science Foundation, the new system features cloud integration and is now available to researchers across a broad swath of science domains. Image: Owen Stanley, SDSC/UC San Diego*

machine learning, and traditional simulation. *Expanse* also supports integration with the public cloud providers, leveraging high-speed networks to ease data movement to/from the cloud, and a familiar scheduler-based approach.

The new system has been in early-user testing for the past month, with researchers from various domains and institutions running actual research projects to validate *Expanse*'s overall performance, capabilities, and reliability (see *sidebar*).

Like *Comet*, which is slated to conclude operations as an NSF resource on March 2021 after six years of service, *Expanse* is designed for modest-scale jobs of just one core to several hundred cores. This includes high-throughput computing jobs via integration with the OSG, which can have tens of thousands of single-core jobs. Those modest-scale jobs are often referred to as the 'long tail' of science. Virtually every discipline, from multi-messenger astronomy, genomics, and the social sciences, as well as more traditional ones such as earth sciences and biology, depend upon these medium-scale, innovative systems for much of their productive computing.

"*Comet*'s focus on reliability, throughput, and usability has made it one of the most successful resources for the national research community, supporting tens of thousands of users across all domains," said SDSC Deputy Director Shawn Strande, a co-PI and project manager for *Expanse*. "Our approach with *Expanse* was to assess the emerging needs of the community and then work with our key partners including Dell, AMD, NVIDIA, Mellanox, and Aeon, to design a system that meets or exceeds those needs."

*Expanse*'s standard compute nodes are each powered by two 64-core AMD EPYC 7742 processors and contain 256 GB of DDR4 memory, while each GPU node contains four NVIDIA V100s connected via NVLINK, and dual 20-core Intel Xeon 6248 CPUs. *Expanse* also has four 2 TB large memory nodes. The entire system, integrated by Dell, is organized into 13 SDSC Scalable Compute Units (SSCUs), comprising 56 standard nodes and four GPU nodes, and connected with 100 GB/s HDR InfiniBand.

Remarkably, *Expanse* delivers over 90,000 compute cores in a footprint of only 14 racks. Direct liquid cooling (DLC) to the compute nodes provides high core count processors with a cooling solution that improves system reliability and contributes to SDSC's energy efficient data center.

Every *Expanse* node has access to a 12 PB InfiniBand-based Lustre parallel file system (provided by Aeon Computing) that delivers over 140 GB/s. Local NVMe on each node gives users a fast scratch file system that dramatically improves I/O performance for many applications. In 2021, a Ceph-based file system will be added to *Expanse* to support complex

workflows, data sharing, and staging to/from external sources. The *Expanse* cluster is managed using the Bright Computing HPC Cluster management system, and the SLURM workload manager for job scheduling.

A video of SDSC's presentation on *Expanse* during last month's International Conference for High-Performance Computing, Networking, Storage, and Analysis (SC20) [can be found here](#). Other details and full specifications [can be found here](#).

*Expanse* will serve as a key resource within the NSF's Extreme Science and Engineering Discovery Environment (XSEDE), which comprises the most advanced collection of integrated digital resources and services in the world. The NSF award for *Expanse* runs from October 1, 2020 to September 30, 2025 and is valued at \$10 million for acquisition and deployment, plus an estimated \$12.5 million for operations and maintenance.

---

#### MEDIA CONTACT

**Jan Zverina**, 858-534-5111, [jzverina@sdsc.edu](mailto:jzverina@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](mailto: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>.