DELL Technologies

intel

A Supercomputing Condo

San Diego Supercomputer Center empowers academic researchers with a shared Intel-based high performance computing system from Dell Technologies.

SDSC SAN DIEGO SUPERCOMPUTER CENTER

Scientific research | United States

Business needs

The San Diego Supercomputer Center provides HPC resources to faculty and students on the University of California San Diego campus via its Triton Shared Computing Cluster.

Solutions at a glance

- Dell EMC PowerEdge servers
- Intel® Xeon® processors
- Lustre parallel file system storage
- InfiniBand and Ethernet interconnects

Business results

- Making HPC accessible to hundreds of researchers
- Supporting widely diverse workloads in a shared HPC environment
- Accelerating scientific discovery in a wide range of academic disciplines
- Enabling hybrid HPC services with cloud connections

The Triton Shared Computing Cluster is used by hundreds of researchers in more than

35 labs and groups

The multi-user cluster is rated to deliver HPC performance of more than

100 teraflops

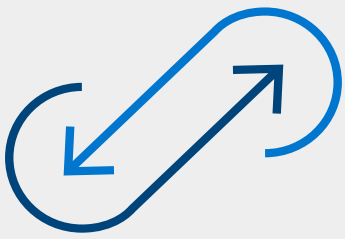

Powering groundbreaking research

The San Diego Supercomputer Center serves as a national resource that provides high performance computing resources to tens of thousands of researchers across the United States, including scientists and engineers working in academic and industrial domains. The center, known informally as SDSC, has gained a great deal of attention for its petascale supercomputers funded by the National Science Foundation, including its Comet supercomputer and its new Expanse system, which is slated to enter production in late 2020.

As part of its mission, the SDSC also provides a research cluster serving the computational science needs of faculty and students on the University of California San Diego campus. This midsize HPC system, known as the Triton Shared Computing Cluster, provides essential capabilities to power the groundbreaking research that takes place every day on the UC San Diego campus.

A shared computing cluster

The Triton Shared Computing Cluster is UC San Diego's primary HPC system for research. The system is operated under a "condo cluster" model, in which researchers use their contract and grant funding to buy compute nodes for the cluster.

This is a common model for universities these days, according to Ron Hawkins, program manager for the Triton Shared Computing Cluster. It works well when researchers have contract and grant funds or faculty startup funds that they can use to buy equipment. It allows them to contribute a relatively small fraction of the resources used in the system to obtain access to a much larger pool of resources.

"Under the condo model, researchers get access to the cluster proportional to the number of resources they contribute," Hawkins says. "But the beauty is, they could buy one compute node and then be able to run jobs that take 10 compute nodes or 20 compute nodes. So they get access to a much larger resource than they could afford just for their lab or their group."

The Triton Shared Computing Cluster has about 400 compute nodes based on the x86 processor architecture developed by Intel and about 300 GPUs. The system, launched in 2013, is a highly heterogeneous cluster that has grown organically over time as researchers bought additional nodes for the system.

In 2017, after evaluating proposals from multiple technology vendors, SDSC moved to standardize the system on Dell EMC PowerEdge servers with Intel® Xeon® processors. Since then, the standard server components in the system are the Dell EMC PowerEdge C6400 four-node compute chassis and the Dell EMC PowerEdge R740 server, which is used for one- and two-node requirements.

Powering diverse research workloads

The Triton Shared Computing Cluster serves as the go-to HPC resource for more than 35 labs or groups on the UC San Diego campus. That equates to hundreds of system users running very diverse workloads.

In the realm of the hard sciences, the system supports applications for genomics, biomedical research, engineering, computational chemistry, biology, geophysics, earthquake simulations, climate research and more.

"Our science users run the gamut, from biomedical researchers looking at origins and therapies for pediatric brain disease, to causes and treatments for neurological disease in aging brains, to new materials for lithium ion rechargeable batteries, to chemistry research in protein structures," Hawkins says.

The Triton Shared Computing Cluster also runs data-intensive applications used by economists, political scientists, business faculty members and others. Many of these researchers are now using advanced data analytics tools and machine and deep learning techniques to analyze large datasets.

"The use of computational methods is broadening into virtually every scientific domain now," Hawkins notes. "And that's partly driven by the big data phenomenon and, in general, the eagerness to apply computational methods, including machine learning and neural networks, to all types of research."

Understanding earthquakes

SDSC's supercomputing clusters support many research projects focused on earthquakes, including studies of seismic faults, simulations of catastrophic events and predictions about what happens when the earth begins to shake.

Some of this work is propelled forward with the support of seismic-simulation software developed in a collaborative effort by SDSC, the University of California San Diego and Intel. This initiative yielded extremely fast computer-based simulations to predict ground motions and improve our ability to save lives and minimize property damage from earthquakes.

SDSC's HPC systems also support a wide range of earthquake research funded by the U.S. Geological Survey, the National Science Foundation and the Southern California Earthquake Center. In one of these projects, researchers used SDSC's Comet supercomputer to advance dynamic earthquake rupture models of faults in a seismic zone in southern California.

The study provides seismologists and geologists with a new understanding of a complex set of faults in the region that has the potential to impact the lives of millions of people in the United States and Mexico, according to SDSC representatives. The center notes that some of the findings point to the possibility of a multi-fault earthquake in Southern California, which could have dire consequences.

In looking at the larger scientific context, the lead author of the study — Christodoulos Kyriakopoulos, a research geophysicist at the University of California Riverside — notes that this research has contributed to a better understanding of multi-fault ruptures, which could lead to better assessments of the earthquake hazards.

"If we know how faults interact during earthquake ruptures, we can be better prepared for future large earthquakes — in particular how several fault segments could interact during an earthquake to enhance or interrupt major ruptures," he says.

Looking ahead: Triton Stratus

The HPC resources at SDSC never stand still. The organization is always looking ahead to system expansions and new supercomputers. That's the case with the Triton Shared Computing Cluster. SDSC is now laying the groundwork for a major enhancement of the system, called Triton Stratus.

When it goes into service in 2020, Triton Stratus will provide researchers with improved facilities for using emerging computing paradigms and tools, including interactive and portal-based computing, and scaling them to commercial cloud computing resources. Triton Stratus, funded by the National Science Foundation, will establish a pool of new compute nodes dedicated to Jupyter Notebooks and RStudio, tools that are used to implement computational and data analysis functions and workflows.

About the San Diego Supercomputer Center

- SDSC is a leader in all aspects of advanced computing, including data integration and storage, performance modeling, data mining and predictive analytics, and software development.
- SDSC provides resources, services and expertise to the national research community, including academia, industry and government.
- SDSC supports hundreds of multidisciplinary programs spanning a wide variety of domains, from astrophysics and bioinformatics to environmental sciences and health IT.
- SDSC is a Dell Technologies HPC & AI Center of Excellence.

DELL Technologies

The project will also explore and deploy techniques for packaging tested notebooks and sending them or bursting them to commercial cloud services for greater scale or throughput.

“We see Triton Stratus operating under a hybrid on-premises cloud model,” Hawkins says. “We have a number of science users who are already using commercial clouds in varying ways, for data sharing primarily. They would like to have more integration with what they are doing on the local cluster.”

With Triton Stratus, researchers have the flexibility to leverage a mix of on-premises and cloud-based HPC resources.

“In the case of the interactive computing, they can use the on-premises part of the cluster for development and testing — for example, to develop and test scripts for Jupyter Notebooks,” Hawkins says. “And then when they have those fully tested and they want to package them up and run multiple instances in a high-throughput computing mode capacity, they can use the cloud — to get more elastic resources for that short duration when they need to scale for high-throughput.”

Working with Dell Technologies

In working with Dell Technologies on the Triton Shared Computing Cluster and the Triton Stratus expansion, SDSC is building on a well-established relationship. The HPC specialists at SDSC worked closely with Dell Technologies to design and deliver the Comet system, which has nearly 2,000 Dell EMC PowerEdge servers, and the new Expanse supercomputer, which will offer 93,000 compute cores and 5 petaflops of HPC performance.

“Dell has been a great partner for us for the Triton Shared Computing Cluster and other HPC systems at SDSC,” Hawkins says. “And now we are expanding capabilities with Dell’s partnership and assistance. We are appreciative of that, and excited about the new capabilities that we are going to be providing to our researchers in the near future.”

SDSC SAN DIEGO SUPERCOMPUTER CENTER

Intel Inside®. Powerful Solution Outside.