

SDSC Expands Its DataStar Supercomputer To Support Researchers' Extreme Data-Intensive Needs

DataStar Now Offers 15.6 Teraflop Capacity and More Than 2048 Processors for Unprecedented Capability at SDSC

September 21, 2005

Ashley Wood

As part of a continuing effort to serve the broad community of science and engineering researchers and educators, the San Diego Supercomputer Center at UCSD has expanded the capacity and capability of its new DataStar supercomputer. Through the addition of 96 8-way IBM Power 4+ p655 compute nodes, SDSC's users will now have access to one of the largest computers available to the open academic community in the nation. The newly expanded DataStar will provide users 50 percent more capacity at SDSC, which will help meet the heavy demand for the center's compute time. In addition, DataStar's memory and parallel file system will almost double in size, giving users the ability to output more data in research areas such as astronomy, geosciences, fluid dynamics and others.

SDSC is one of the premier TeraGrid sites, with particular responsibility for data-intensive computing. TeraGrid - built over the past 4 years - is the world's largest, most comprehensive distributed cyberinfrastructure for open scientific research. The DataStar expansion will greatly improve overall performance, making it a more powerful tool for all of its users.

"DataStar is considered the premier environment for users whose codes are both compute-intensive and data-intensive," said Dr. Fran Berman, director of SDSC. "With the DataStar expansion, we will be able to support both a greater number of codes, and more "heroic" codes including larger-scale simulations, deeper analyses, and more complex application models. We are delighted to be able to provide this service for the user community."

"The DataStar Expansion brings us over the threshold, allowing us to simulate turbulence at world-class grid resolution and rivaling work on the Earth Simulator," said P.K. Yeung, professor of Aerospace Engineering at Georgia Tech and SDSC user. "In the future, this will allow the details of turbulent mixing and associated scaling laws to be computed more definitively than ever before."

With 2172 IBM Power 4+ processors now available in the 8-way p655 nodes, as well as an additional 11 32-way p690 nodes, the expanded DataStar can offer users lightening-fast capability for the most extreme needs of data-intensive applications. This increase will give SDSC users access to an additional five million processor hours for a total of 15 million processor hours that the center offers free to users each year.

The system now offers 15.6 Teraflops of compute capability - an increase of more than 50 percent.

The aggregate memory as well as the size of DataStar's parallel file system will almost double as a part of this expansion. The expanded system will give SDSC users 7.3 terabytes of aggregate memory and 115 terabytes of parallel file system disk storage to help them run larger computations and store more data.

With 2172 IBM Power 4+ processors now available in the 8-way p655 nodes, as well as an additional 11 32-way p690 nodes, the expanded DataStar can offer users lightening-fast capability for the most extreme needs

of data-intensive applications. This increase will give SDSC users access to an additional five million processor hours for a total of 15 million processor hours that the center offers free to users each year.

The system now offers 15.6 Teraflops of compute capability - an increase of more than 50 percent.

The aggregate memory as well as the size of DataStar's parallel file system will almost double as a part of this expansion. The expanded system will give SDSC users 7.3 terabytes of aggregate memory and 115 terabytes of parallel file system disk storage to help them run larger computations and store more data.

"Data-intensive computing has rapidly become a principal mode for scientific exploration. This expansion to SDSC's DataStar system demonstrates SDSC's and NSF's recognition that the best tools must be made available to the science and engineering communities," said Jose Munoz, deputy director of the office of Cyberinfrastructure at the National Science Foundation (NSF). "SDSC has been a leader in data-intensive computing and this upgrade maintains that leadership. Capabilities such as those being made available through this expansion, coupled with SDSC's excellent scientific and support staff, are paramount for continued US leadership in science, engineering and education."

The 96 new IBM p665s nodes provide faster CPUs and double the memory than the originally installed nodes. Each new node has 1.7 GHz of processing power and 32 GB of memory.

"This remarkable system is the result of deep and ongoing collaboration between IBM and the San Diego Supercomputer Center, now empowering its community of researchers even further with increased performance and ability," said Dave Turek, vice president of deep computing, IBM. "Increases in performance such as this one are a result of the scalability benefits of IBM's systems for high performance computing."

Media Contacts: Greg Lund, SDSC Communications, (858) 534-8314 Ashley Wood, SDSC Communications, (858) 534-8363