DataDirect NETWORKS

# Case Study

www.datadirectnet.com

S2A Technology

SDSC

## SAN DIEGO SUPERCOMPUTER CENTER

San Diego Supercomputer Center Provides Extreme Computing using DataDirect Networks Storage System

+1.800.TERABYTE (837.2298) – sales@datadirectnet.com

DataDirect NETWORKS

# Case Study

www.datadirectnet.com

## San Diego Supercomputer Center Provides Extreme Computing using DataDirect Networks Storage System

### Application

San Diego Supercomputer Center operates multiple large cluster computers, enabling global, breakthrough science and engineering discoveries. It is also the data-intensive site lead for TeraGrid, a large-scale, multi-site grid infrastructure for open scientific research.

### Challenge

To expand SDSC's storage infrastructure using the best-in-class solution for large capacities, high availability, high reliability, heterogeneous platform support, scalability, and high-performance I/O within budget.

### Solution

DataDirect Networks S2A storage solution outperformed competitive offerings across SDSC's requirements, providing superior read/write performance, reliability, data integrity, and availability for massive data ingest, storage and access, archiving, and compute-intensive parallel processing, with excellent scalability.

S2A TECHNOLOGY

Founded in 1985, the San Diego Supercomputer Center (SDSC) enables international science and engineering discoveries through advances in computational science and high performance computing. From incredibly detailed simulations of the effects of "the Big One" – a 7.7-magnitude earthquake on Southern California's San Andreas Fault – to discovering the structure of autism-related proteins, and zeroing in on the causes of Parkinson's disease and Alzheimer's disease, or simulating the origin of the universe's first star, the resources of SDSC are helping scientists move virtual worlds closer to reality.

Over the years, SDSC has served more than 10,000 researchers at 300 academic, government, and industrial institutions in the United States and around the world. Today, these scientists and engineers rely more than ever on the availability of globally accessible data "cyberinfrastructure" tools, or advanced computing resources available between remote locations, to drive increasingly complex, large-scale, and cooperative scientific endeavors.

> "We love the capabilities of the DataDirect Networks solution; especially its ability to run RAID 6 at top speed with no performance hit in reads or writes."
>
> — Bryan Banister, Manager of Storage Systems and Production Servers at SDSC

### SDSC's Extreme Computing Cyberinfrastructure

SDSC operates powerful high-end computing resources including DataStar, a 15.6 TeraFLOPS IBM Power4+ supercomputer with total aggregate memory of 7.3 terabytes. In addition, SDSC was the first academic institution to install an IBM BlueGene/L eServer, called Intimidata by SDSC staff, which now provides 17.2 TeraFLOPS of compute power with over 6,000 processors. Both DataStar and Intimidata are ranked among the top supercomputers in the world, according to Top500.org's list of the Top 500 supercomputer sites, and are used for large-scale, compute-intensive scientific research applications.

SDSC also serves as the data-intensive site lead in the National Science Foundation (NSF)-funded TeraGrid, a multiyear effort to build and deploy the world's first large-scale multi-site production grid infrastructure for open scientific research. SDSC hosts a 4.4-TeraFLOPS IA-64 Linux cluster, 1.4 petabytes of online disk storage with more than 25 petabytes of archival storage, and 220 terabytes of storage on IBM's General Parallel File System (GPFS) mounted across the TeraGrid, which is connected to the other national TeraGrid partners by a 20-Gbps cross-country backbone. A key component of TeraGrid is SDSC's Storage Resource Broker (SRB), a metadata resource that provides transparent access to data stored in network-accessible data servers and archives.

### Extreme Computing Demands Best-in-Class Storage

Bryan Banister, Manager of Storage Systems and Production Servers at SDSC, described TeraShake – a highly compute intensive, cooperative project between SDSC and the Southern California Earthquake Center. "TeraShake simulates the propagation of seismic waves from a magnitude 7.7 earthquake on the San Andreas Fault. The simulation takes 240 processors 5 days to complete, and generates 47TB of data per simulation, with 10-20TB of data transferred to archive daily."

And while these data sizes are massive, Banister realizes that both data sizes and the resulting storage requirements will continue to grow rapidly. The cooperative team of the Earthquake Center and SDSC are continually working to optimize and enhance both I/O management and checkpointing for the TeraShake simulation code, and hope to double the output resolution within a year. He calculated, "A doubling of resolution will require 1,000 processors working for up to 20 days, with resulting data in excess of 1PB for one simulation." At this rate, he anticipates that his archival data needs will grow to more than 100PB by 2011.

Extreme computing tasks like TeraShake, with their resultant massive file sizes, put a heavy burden on the Center's storage system. When the Center recently tripled the size of its Intimidata cluster computer, it also had to expand its storage infrastructure to eliminate bottlenecks between its massive compute resources and its storage.

DataDirect NETWORKS

# Case Study

www.datadirectnet.com

## Discovering the Origin of the Universe

By compressing 100,000 years of star evolution into weeks of simulations, SDSC has helped astrophysicists gain a better understanding of the origin of our universe. Using vast amounts of data and the power of SDSC's supercomputers, scientists can gauge the progression of star formation.

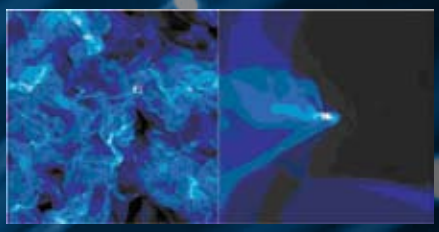

**Simulations of star formation using SDSC's DataStar supercomputer shed new light on how stars are born. These simulations complement voluminous and growing observational data.**

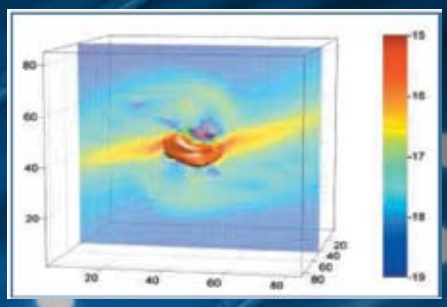

**Capturing the complexities of star formation, this simulation shows the core outflow in a simulation of fragmenting filaments in a turbulent cloud.**

Banister and his team had several mandatory requirements for the new system, with heavy focus on reliability and availability. Their checklist of mandatory features, as well as capacity and performance per terabyte, was matched against the offerings of nine storage vendors. After evaluating the field of choices, Banister and his team chose DataDirect Networks' S2A (Silicon Storage Architecture) system because it outperformed the other solutions across all of SDSC's requirements. They have installed nearly 1PB of DataDirect Networks storage in a compact footprint.

## DataDirect Networks S2A Solution

Banister explained that their storage system needed to provide enough aggregate throughput to handle multiple simultaneous tasks at once: the concurrent ingest of new data and retrieval of existing data for processing during a simulation, along with backup and archiving functions, including the staging of data being restored from tape.

"The performance of the S2A is excellent," Banister confirmed, highlighting the S2A's ability to deliver sustained high-throughput reads and writes to a shared storage pool as multiple systems and users concurrently access content for processing while new content is being ingested, along with backup activities going on in parallel. DataDirect Networks integrates multiple zero-latency, full-access host ports, each able to simultaneously access the entire pool of storage at full speed.

Given their history of using large numbers of SATA drives for large data sets, such as 13TB of celestial images for the National Virtual Observatory, the team was concerned that they would have data loss due to a double disk failure during a rebuild operation. "SDSC is a premiere scientific computing center. You can't lose people's data," he emphasized. However, the cost of enterprise-class drives was prohibitive.

Uptime was also crucial. According to Banister, "If our storage system isn't available, then the compute systems don't run. The cost of being offline easily exceeds thousands of dollars per hour just for the system, and that doesn't include loss of research time, manpower, or the domino effect that delays subsequent projects."

The S2A system includes DataDirect's SATAssure intelligent SATA drive management system, which protects against silent data corruption and double disk failures in addition to supporting partial rebuilds. The S2A delivers enterprise-class data protection to affordable SATA drives without sacrificing performance or uptime, even during drive recovery events.

> "We had some concerns about the reliability and performance of SATA, but with the S2A, it has proven to be as reliable as Fibre Channel under heavy workloads and duty cycles," Banister commented.

Other "must haves" were rock-solid data integrity, including parity calculations on both reads and writes, and the ability to support the heterogeneous systems in use by SDSC. The Center uses two different file systems: for archiving, backup, and remote disaster recovery, in addition to the use of IBM's GPFS parallel file system with TeraGrid.

"We love the RAID 6 capabilities. Data integrity is a very important feature, and the fact that the S2A does parity calculations on reads as well as writes helped sway our decision to select DataDirect Networks. That wasn't available from the other vendors, who only do it on writes."

DirectRAID is a unique S2A feature, providing hardware-based high performance and predictability for applications, including real-time, full-speed double-parity protection, which enables the S2A's hitless performance even in the face of disk failures.

He expanded on the S2A's reliability, adding, "The S2A's fault-tolerant architecture with zero-time failover has been an important asset. It allows us to keep the system online even in the event of an appliance failure."

As Banister observed in relation to the TeraShake project, SDSC's storage needs will continue to grow at an astounding rate. They also plan to implement faster interconnects and add more compute nodes. He expects the DataDirect Networks solution to provide the scalability he needs as he looks to expand beyond a single storage unit. "I don't want to allocate a whole array to a single project. I want to allocate only what's needed across multiple arrays."

"If our archival data grows to more than 100PB within the next five years as I expect, we need a storage solution that can support our requirements and ensure that our system stays online. We believe the DataDirect Networks S2A system is designed to grow with us."

S2A TECHNOLOGY

DataDirect Networks is the leading provider of open, scalable storage systems for performance and capacity driven applications. DataDirect's S2A (Silicon Storage Architecture) appliance enables modern applications such as video streaming, content delivery, modeling and simulation, backup and archiving, cluster and supercomputing, and real-time collaborative workflows, that are driving the explosive demand for storage performance and capacity. DataDirect's S2A technology and solutions solve today's most challenging storage requirements, including providing shared, high-speed access to a common pool of data, minimizing data center footprints and storage costs for massive archives, reducing simulation computational times, and capturing and serving massive amounts of digital content.

Major corporations, supercomputing centers and rich media organizations, including AOL, Ascent Media, Boeing, CINECA, CGGVeritas, CNN, Disney, Federal Reserve Board, Ford, Hess, Kodak Gallery, Lawrence Livermore National Laboratories, NASA Ames, RIOT, Sandia National Laboratories, Sony, Technical University Dresden, Technicolor, Time Warner, Thomson, Trinity College Dublin and Universal, utilize DataDirect high performance, high capacity solutions.

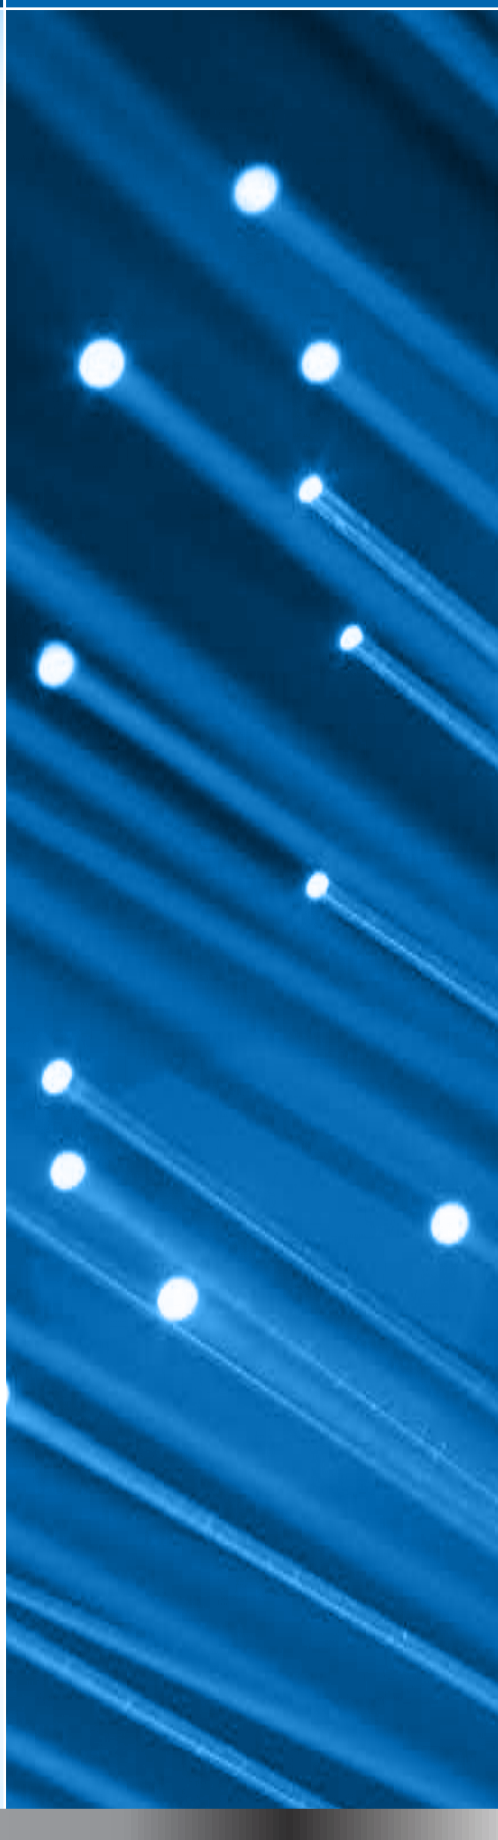

DataDirect NETWORKS

Performance. Capacity. Innovation.

9351 Deering Avenue . Chatsworth . California 91311
phone +1.800.TERABYTE (837.2298) . fax +1.818.700.7601
sales@datadirectnet.com
www.datadirectnet.com