



## **Lawrence Livermore National Laboratory**

Terascale Simulation Facility Employs  
S2A Storage System to Deliver a Scalable,  
High Performance, High Capacity, Open-  
system Solution for High Speed Scientific  
Simulation, Data Analysis and Visualization

www.datadirectnet.com

## Challenge

BlueGene/L needed a storage solution that would meet its demanding bandwidth, capacity and reliability requirements. Further, BlueGene/L required cost-effective scalable storage solution with support for access from multiple computing resources in the BlueGene/L simulation environment by means of a non-proprietary solution.

## Application

Installed 28 DataDirect Networks' S2A8500s with over 1 PetaByte (PB) of SATA storage capacity supporting 65,536 compute nodes (131,072 processors) and 1,024 I/O nodes. The S2A8500 solution delivered sustained performance over 24 GB/s.

## Solution

DataDirect Networks' S2A (Silicon Storage Appliance) Storage System enabled extreme application performance and scaling with GB/second performance. S2A supported a variety of LLNL applications and provided open solution choice for IBM's BlueGene supercomputer and CFS's Lustre File System. Further, the S2A reduced complexity providing ease of scalability and quicker time to management. This reduced complexity decreased storage, facility and service costs.

BlueGene/L, Number 1 on the TOP500 list of supercomputers, is a revolutionary, low-power per FLOP/sec machine that delivers extraordinary computing power for scientific simulations and programmatic work. Located in the Terascale Simulation Facility at LLNL, BlueGene/L is running a wide variety of scientific simulations in support of the Advanced Strategic Computing (ASC) and Stockpile Stewardship programs. These simulations include classical molecular dynamics, ab initio quantum molecular dynamics, dislocation dynamics, and turbulent hydrodynamics applications delivered at speeds over 100 teraFLOP/s to address materials aging, performance of high explosives, carbon graphite transitions, super-ionic water, turbulence and instability physics, plastic strength of materials all relating to LLNL's core Stockpile Stewardship mission. From a computer science and engineering perspective, BlueGene/L is also used to explore the potential of system-on-a-chip technologies to achieve extreme scalability and performance while minimizing floor space and electrical power consumption.

BlueGene/L is fully assembled and installed for classified service in support of the National Nuclear Security Administration's stockpile science mission of Los Alamos, Lawrence Livermore, and Sandia National Laboratories. The machine recently achieved a sustained world-record speed of 280.6 teraFLOP/s on the Linpack benchmark and over 207.3 teraFLOP/s on quantum molecular dynamics code "Qbox," demonstrating the machine's huge potential for rapid "time to solution" for applications in molecular dynamics and materials science.

To meet those performance needs, LLNL selected DataDirect Networks' Silicon Storage Appliance 8500 (S2A8500) for use with BlueGene/L in this critical national security work. The Terascale Simulation Facility purchased and installed 28 S2A8500s with over 1 PB of SATA storage capacity.

"The deployment of BlueGene/L into Lawrence Livermore's scalable scientific simulation environment required a commensurate quantum jump in I/O scalability, performance, and reliability. Our mission drivers require a storage architecture that enables an extremely cost effective, scalable, and balanced I/O architecture that can be quickly and reliably deployed and operated while reducing complexity," said Mark Seager, Assistant Department Head for Advanced Technology at Livermore. "We needed a solution which was optimized for the unrelenting demands of scientific computing at the BlueGene/L scale."

DataDirect also supports other unclassified LLNL simulation environments, such as the two previous generations of the world's fastest Linux clusters:

Thunder at 23 teraFLOP/s, MCR at 11.1 teraFLOP/s and ALC at 9.6 teraFLOP/s. This massive, world-class simulation environment is powered by more than 200 of DataDirect's S2A storage controllers serving 7 PB of DataDirect storage.

## BlueGene/L Storage Solution

BlueGene/L boasts a peak speed of over 360 teraFLOP/s and a total memory of 32 TB, using only 1.8 megawatts of power and 2,500 square feet of floor space. The full system has 65,536 dual-processor compute nodes (131,072 processors), 1,024 I/O nodes and is connected to the storage area and general access networks by 1,024 Gigabit Ethernet links.

For its storage infrastructure, BlueGene/L runs a NAS solution referred to as a BlueGene/L I/O Scalable Unit. At the heart of this I/O Scalable Unit is DataDirect Networks' S2A8500 Couplet (dual controller) configured with 32 TB of useable capacity, employing 148 250 GB SATA disk drives in a 8+P+S RAID 3 configuration.

The I/O Scalable Unit incorporates Cluster File System's Lustre File System Object Storage System (OSS). Lustre is closely coupled with DataDirect Networks' S2A8500. Each S2A8500 is paired with eight dual Xeon processor nodes running the Lustre OSS.

The entire I/O performance of BlueGene/L is supported by only 28 S2A8500 storage systems and 224 Lustre OSS nodes. This NAS solution delivers sustained performance over 24 GB/s.

BlueGene/L is attacking large scale simulation, visualization and modeling problems with parallel applications. These parallel applications benefit greatly from the ability to minimize I/O cycles when dumping large amounts of data from system memory to storage. In many cases, storage becomes the limiting factor in dumping large amounts of data from cache to disk. The S2A8500 with its large sustained bandwidth enables the I/O Scalable Unit to





write data at very high speed allowing BlueGene/L to dump system memory quickly so it can go back to processing. This very high write performance increases application productivity by limiting the amount of time BlueGene/L spends processing I/O.

## DataDirect Networks' Technology Advantage

DataDirect Networks storage system is based on the Silicon Storage Appliance. The S2A delivers consolidated RAID 3, data protection, switching, storage management, data management, and storage performance functionality in one simple, easy to use intelligent device. This concentrated power delivers sustained read or write performance of up to 2.8 GB/s (when using the S2A9550), supporting up to 840 Terabytes (TB) per system.

The S2A's industry leading performance is enabled by a parallel, non-blocking architecture with 11.2 GBytes/second internal bandwidth. The internal performance is powered by a lightning fast cache surrounded by parallel and redundant host and disk side high-speed data-mover ASICs and software. This consolidated functionality integrated with a parallel processing architecture delivers high performance, high capacity storage solutions that are easy to install, use, scale and manage. The S2A's unique architecture and consolidated functionality removes the complexity, congestion, latency and contention of generic SAN RAID systems, delivering very high performance and capacity with dramatically lower total cost of ownership.

The S2A's innovative architecture maintains performance even in the event of multiple drive rebuilds or failures. With integrated hardware RAID 6, parity checking of all read I/Os and automated data protection, the S2A delivers enterprise-class reliability with no loss in performance under the most demanding compute-intensive conditions.

The S2A8500 and S2A9500 Series scale from 10 TB to 840 TB in a single storage system with Fibre Channel and/or SATA drives. They support a 720 TB package that fits in just two racks (a little over two floor tiles). This scalability and cost effective density makes it ideal for NAS and SAN storage, parallel or shared file system storage applications.

All S2A Series products support Fibre Channel and InfiniBand fabrics, but the S2A9500 Series support both natively. With 8 Fibre Channel-4 and/or 4 InfiniBand 4X host ports, the S2A9500 Series can connect natively to servers in either network at the same time. This interconnect delivers the flexibility to move from one network to another without losing an organizations storage investment.

With 7 generations of deployment experience and an open solution architecture, the S2A offers broad support for cluster, collaborative and high performance computing applications, servers, networks, operating systems and clustered file systems.

The S2A storage solutions support NAS or SAN storage, parallel or shared file system storage, production or nearline storage.

Customers using DataDirect's Silicon Storage Appliances include Sandia National Laboratories, ORNL, PCS, NCSA, NOAA, NASA, TACC, NERSC, MAUI SC, SDSC, Argonne National Laboratory, U.S. Army Research Lab, U.S. Air Force Research Lab, CEA, AWE, Dresden TUD, and many others. DataDirect Networks' successes also include the number one supercomputer in Europe, "Tera10" at CEA.

## S2A Storage Solutions Specifically Designed for Supercomputing Clusters and High Performance Computing Applications

The S2A storage system with its consolidated SAN functionality and parallel and non-blocking architecture was specifically designed to eliminate storage bottlenecks and enable scalable, open solution storage architectures that elevate the performance of supercomputing clusters and high performance computing applications.

Whether parallel or shared NAS cluster file systems with Fibre Channel, GigE or InfiniBand fabrics, the S2A storage system provides scalable, cost-effective performance for thousands of client computers and I/O servers with low latency and Quality of Service.

The S2A storage solution brings the benefits of scalability, storage consolidation and ease of management by directly supporting thousands of compute clients and terabytes of storage from a single storage system.

The S2A storage system supports, parallel NAS and shared SAN clusters and a broad range of compute systems including, IBM, Dell, Cray, SGI, Bull, LNXI, HP and others.

S2A8500 and BlueGene/L I/O Scalable Unit



DataDirect Networks is the leading provider of scalable storage systems for performance and capacity drive applications. DataDirect's S2A (Silicon Storage Appliance) architecture 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, Autodesk, Boeing, CNN, Disney, Federal Reserve Board, FedEx, Ford, Hess, Kodak Gallery, Lawrence Livermore National Laboratories, NASA Ames, RIOT, Sandia National Laboratories, Sony, Technicolor, Time Warner, Thomson, Universal, and Veritas DGC, utilize DataDirect Networks 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