CASE STUDY: SLAC NATIONAL ACCELERATOR LABORATORY

Samsung Solid State Drives Speed Scientific Discovery

OVERVIEW

Customer Need

SLAC National Accelerator Laboratory wanted to maximize the use of its powerful Linac Coherent Light Source (LCLS) – an X-ray laser that creates stop-motion images of atoms and molecules in action, resulting in greater understanding of fundamental processes, like photosynthesis. SLAC wanted to increase the rate at which the system could acquire data, so more scientists could make greater use of the technology.

Samsung Solution

SLAC implemented Samsung Solid State Drive (SSD) technology to store and retrieve the data from its online cache. Unlike traditional Hard Disk Drive (HDD) technology, the Samsung 256GB 2.5-inch SSD 840 Pro Series does not use any mechanical parts, such as motor-driven spindles. Instead, the data are stored in the form of electrons within the arrays of tiny transistors in a chip, allowing the technology to operate at much faster speeds.

Results

The Samsung SSD technology allows the SLAC X-ray laser to acquire data at up to 8 gigabytes per second of data, compared to just 2 gigabytes per second before. In the short-term, this increased productivity will allow more experiments to be conducted, while allaying concerns of information being lost because of the slow read-write process of HDD. In addition, the greater capacity will allow SLAC to introduce larger detectors and other technologies to improve the quality and granularity of the images.

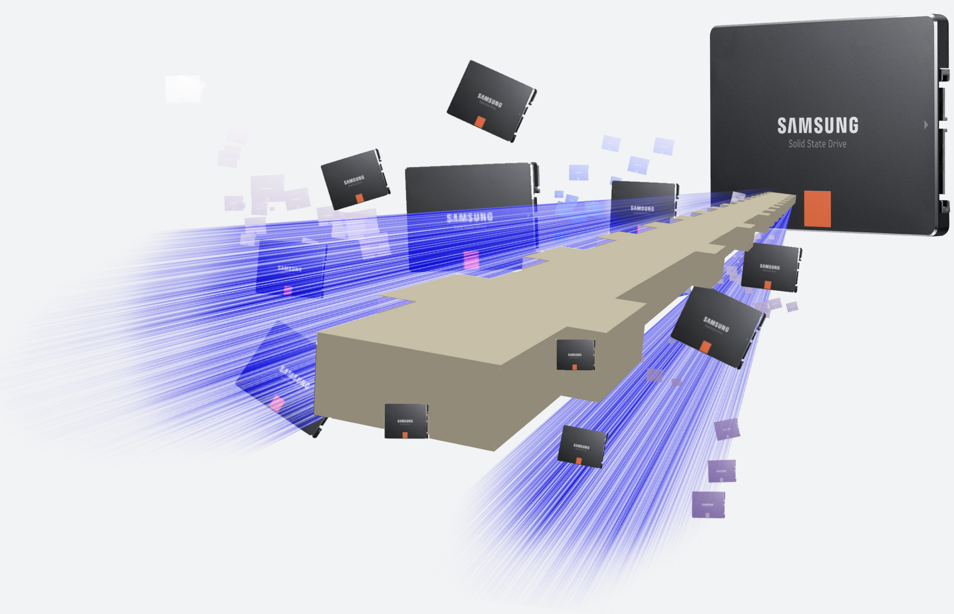

THE CUSTOMER

SLAC National Accelerator Laboratory

SLAC National Accelerator Laboratory, which is operated by Stanford University on behalf of the U.S. Department of Energy, has spent the last half-century on the forefront of technological innovation.

Since SLAC opened in 1962, six scientists have received Nobel Prizes for work done at the lab, and more than 1,000 scientific papers are published annually based on research that is conducted there. SLAC created the first website in North America, built the world's largest particle accelerator, and has discovered some of the fundamental building blocks of matter.

The facility is spread across 151 buildings on 426 acres in Menlo Park, in the renowned Silicon Valley. The lab has some 1,700 full-time employees, 3,400 facility users each year, and 124 graduate students who are conducting work in particle physics and other advanced scientific realms.

Scientists from around the world come to use SLAC's cutting-edge technology, including SLAC's revolutionary X-ray laser, which provides remarkable insight into the intimate details of atoms and chemical reactions. The laser, which shoots a beam narrower than a human hair, is the latest of the landmark technology that has been implemented at SLAC.

These well-known researchers, along with the lab's own staff scientists, have discovered new drugs for healing, new materials for electronics, and new ways to produce clean energy. SLAC regularly makes headlines around the world for the scientific mysteries it unlocks, such as solving a 150-million-year-old riddle about the "dinobird"—famous fossils that link dinosaurs and birds.

THE CUSTOMER NEED

Fast, Reliable Storage for Data Acquisition

SLAC is home to the linear accelerator (or linac), a two-mile-long structure that is buried 30 feet below the ground. Called "the world's straightest object," the linac greatly increases the velocity of charged subatomic particles so they can be used for cutting-edge physics experiments. In April 2009, the linac was applied to a new kind of laser called the Linac Coherent Light Source (LCLS).

The LCLS creates powerful X-ray pulses, each of which are millions of times brighter than medical X-rays. These pulses are analogous to flashes of an extremely fast stroke light. Scientists use the pulses to take stop-motion pictures of atoms and molecules in action. The technology allows a "shutter speed" of less than 1/10th of a trillionth of a second. These images can provide insight about fundamental processes, like photosynthesis, that was not possible before.

LCLS acquires, on average, 100 terabytes of data per week. An experiment at one of the LCLS instruments produces about 10 million X-ray images over the course of 24 hours. LCLS currently manages about 10 petabytes of data.

In the LCLS data acquisition system, the contributions from the various detectors are sent via multicast packets to a set of nodes with local storage. These nodes, which constitute the so-called online cache, assemble the data into events (one event corresponds to one pulse of the LCLS X-ray beam) and write the event to disk. The events are distributed across the online cache nodes in a round-robin fashion in order to provide scalability.

The event data are then transferred to another storage layer where the users can access the data for extracting the science for their experiment. These concurrent read operations can create intervals during which write operations are stalling, causing the buffers to fill up and generating drops in the network traffic. By replacing the HDDs with SSDs, the LCLS data acquisition was able to eliminate these concurrency effects and, consequently, to increase its throughput by a factor of four.

THE SAMSUNG SOLUTION

SSDs Provide Read/Write Buffer

To improve the efficiency of the online cache, SLAC implemented Samsung Solid State Drive (SSD) technology. The LCLS system uses 100 256GB 2.5-inch SSD 840 Pro.

SSDs store and retrieve information in a fundamentally different way than traditional HDD. HDDs read and write data using motor-driven spindles, actuator arms, and other mechanical parts that place a recording head over rapidly spinning magnetic-coated disks.

In contrast, SSDs have no moving parts, instead storing information in the form of electrons within the arrays of tiny transistors in a chip. Because SSDs do not require mechanical movements, the technology accesses stored information in microseconds, up to 100 times faster than even the speediest HDD.

"SSD is normally used by people who need a lot of I/O (input/output) operations," says Amedeo Perazzo, who leads the Photon Controls and Data Systems Department at SLAC. Larger-capacity SSDs tend to have more I/O channels working in parallel, which allows them to handle a bigger stream of data, and thus perform more work faster.

"That isn't the reason we use SSDs, though," Perazzo said. "For us, the SDD acts as a buffer between the data acquisition system and the user system. Solid state technology was the best way to separate the read operations from the write operations."

He notes that the use of HDD mechanical spindles for the online cache, caused the read operations to slow down the write operations badly enough to lose some of the multicast packets, a technique for one-to-many communication over an IP infrastructure in a network. "The standard spindles weren't fast enough to give the writers access to the disk," Perazzo said.

SSDs have other advantages in addition to speed: They operate silently, since they have no moving parts. They consume about half as much power as HDD. And they can store just as much information.

The amount of data the SDD stores depends on the number of flash memory chips it has. Although a chip is only the size of a postage stamp, it can hold a huge amount of information—approximately 300,000 pages of Microsoft Word documents, 5,000 photos, 1,000 MP3 music files, or a dozen movies.

QUICK PROFILE

As used by SLAC National Accelerator Laboratory

SAMSUNG 840 PRO SERIES SOLID STATE DRIVES

Form Factor: 2.5-inch

Capacity: Available in 128GB, 256GB & 512GB

Host Interface: Serial ATA interface of 6.0Gbps; compliant with ATA/ATAPI-8 Standard

Sequential Read Speed: Up to 540MB/s for 256GB & 512GB models; up to 530MB/s for 128GB model

Sequential Write Speed: Up to 520MB/s for 256GB & 512GB models; up to 390MB/s for 128GB model

Random Read Speed: Up to 100K IOPS for 256GB & 512GB models; up to 97K IOPS for 128GB model

Random Write Speed: Up to 90K IOPS

Power Consumption: 0.15W

Encryption: AES 256-bit Full Disk Encryption; Class0 Self Encryption Drive; user can set HDD password in BIOS setup mode

Operating Systems: Windows Vista or Later

Environmental Specs: Operating Temperature of 32°F to 140°F

Weight: 0.15lb

Warranty: 5 years

THE RESULTS

Eliminate Data Loss and Maximize "Beam Time"

The Samsung SSD technology dramatically increased the rate at which SLAC's LCLS system could process data, accommodating the demands of cutting edge scientific research that attempts to understand things as complicated as dark matter and the origins of the universe itself.

"We were able to go from 2 gigabytes per second to 8 gigabytes per second without the need to add new nodes and without introducing dead time in the system or losing multicast packets, which means we can make better use of machine time," Perazzo said. "The cost of operating the machine is the same, so if you can triple or quadruple the rate, you get a proportional increase in productivity."

With spindles, the read operations in the online cache slowed down enough that the experimenters were afraid they might lose some of the multicast packets used by the data acquisition system to move the science data from the instruments to the online cache. "The data that we collect are extremely valuable and very expensive," he said. "Our goal is to have a lossless network."

Since data can be acquired faster, more institutions will be able to use the facilities, allowing for more ground-breaking scientific experiments. "Users are hungry to access the data," he said. "The time on the LCLS is precious."

The LCLS experiments depend on speed; scientists need to review the data in real time so they can adjust the parameters of their experiments, allowing them to get maximum benefit from scarce "beam time."

In particle physics, scientists attempt to immediately and aggressively filter out the data so they can single out particular, and rare, types of events they want to focus on. In LCLS, though, all events may contain, in principle, essential scientific information.

LCLS stores data at a rate and scale comparable to experiments at the world's most powerful particle collider, the Large Hadron Collider in Switzerland.

Implementing the faster SSD cache was also necessary to accommodate the many advances that the lab plans, such as larger, more sensitive detectors, more complex experiments, multiple simultaneous experiments, and a possible increase in the beam pulse rate.

"With the HDD system, we weren't able to scale as well," Perazzo said. "We know bigger detectors are coming, which will allow images with a finer granularity. The result is higher resolution in the processed data which can give us, for example, a better understanding of how certain processes happen in nature."

In an era of Big Data, SLAC's X-ray laser has emerged as a prototypical example of how to handle massive amounts of information, and the SSD system will help usher in the lab's next era of great scientific discovery.

"In the near future we will be taking a factor of 20 times more data than we are taking today," Perazzo said. "Eventually, we expect to use the SSD technology in our offline systems as well, so that the analysis of the data will be able to keep up with the data acquisition rates."

> "We were able to go from 2 gigabytes per second to 8 gigabytes per second...which means we can make better use of machine time,"
>
> – Amedeo Perazzo, SLAC National Accelerator Laboratory