# The Next Generation of Computer Storage

WORK recently began at Lawrence Livermore National Laboratory and sister Department of Energy laboratories on the Accelerated Strategic Computing Initiative (ASCI), one of the largest supercomputing projects of all time. A major component of DOE's science-based Stockpile Stewardship and Management Program, ASCI's computational modeling and simulation capabilities will be used to assess the safety, security, and reliability of our nuclear stockpile.

ASCI's systems will soon be performing a trillion (tera or $10^{12}$) floating-point operations per second (flops) requiring memories of tens of trillions of bytes, which is well beyond the range of existing supercomputers. By 2004, ASCI systems will be performing in the 100-teraflops range. These machines will require systems that may be called on to store a quintillion or $10^{18}$ bytes (an exabyte), which is over ten thousand times beyond the capability of today's supercomputing storage systems. In addition, the transfer rates between these massive processing and storage systems will have to be on the order of tens to hundreds of billions of bytes per second. Achieving a balance between supercomputer processing and memory capacities and storage capabilities is critical not only to the success of ASCI but also to other high-end applications in science modeling, data collection, and multimedia.

Recognizing this coming need and the long-term effort required to achieve this balance, the National Storage Laboratory (NSL) was established in 1992 to develop, demonstrate, and commercialize technology for storage systems that serve even the most demanding supercomputers and high-speed networks. The NSL consisted of an advanced storage hardware testbed at Livermore and distributed software development partners. It involved more than 20 participants from industry, the Department of Energy, other federal laboratories, universities, and National Science Foundation supercomputer centers. The NSL collaboration was based on the premise that no single organization has the ability to confront all of the system-level issues that must be resolved in a timely manner for significant advancement in high-performance storage system technology. Lawrence Livermore and its sister DOE laboratories play leadership roles in developing high-performance storage systems because of their long history of development and innovation in high-end computing—of which storage is a critical component—in order to meet their national defense and scientific missions.

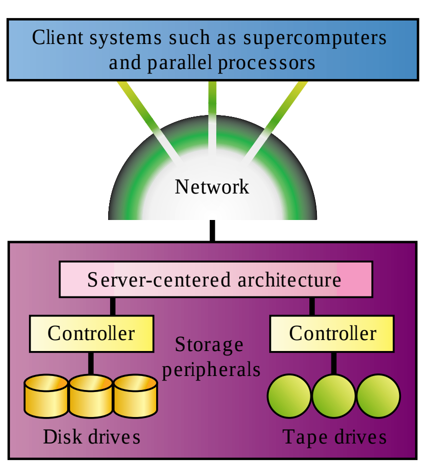


With typical large-scale data storage systems today, general-purpose computers act as storage servers connecting storage units with client systems. As storage rates and capacities increase, the storage servers increase in size and cost, and bottlenecks occur in the transfer of stored data to client systems.

## High-Performance Storage Systems

The High-Performance Storage System (HPSS) software development project grew out of NSL work. A major requirement for HPSS was that it be "scalable" in several dimensions—to allow huge capacities and transfer rates and to support many distributed systems and users. The system also had to be reliable, secure, and portable to many computing platforms and manageable by a small staff.

Work completed by the NSL had shown that HPSS could only be successful if it were based on a network-centered design. Typically, large-scale data storage has been handled by general-purpose computers acting as storage servers that connect to storage units such as disks and tapes (see figure below, left). The servers act as intermediaries in passing data to client systems like workstations or supercomputers on their network. As requirements for storage device data rates and capacities increase, the storage server must handle even more data faster. As data rates increase for storage devices and communications links, the size of the server must also increase to provide the required capacity and total data throughput bandwidth. These high data rates and capacity demands tend to drive the storage server into the mainframe class, which can be expensive to purchase and maintain and can have scalability limits.

If the storage software system and storage devices are instead distributed over a network, control of the storage system can be separated from the flow of data (see figure, p. 23). The bottleneck is removed, allowing more rapid data transmission and scalability of performance and capacity. Workstation-class systems used as storage servers provide the high-performance required and reduce the cost for storage server hardware in the bargain.

## Focus on the Network

Operating on a high-performance network, the High-Performance Storage System uses a variety of cooperating distributed servers to control the management and movement of data stored on devices attached directly to the network. HPSS is designed to allow data to be transferred directly from one or more





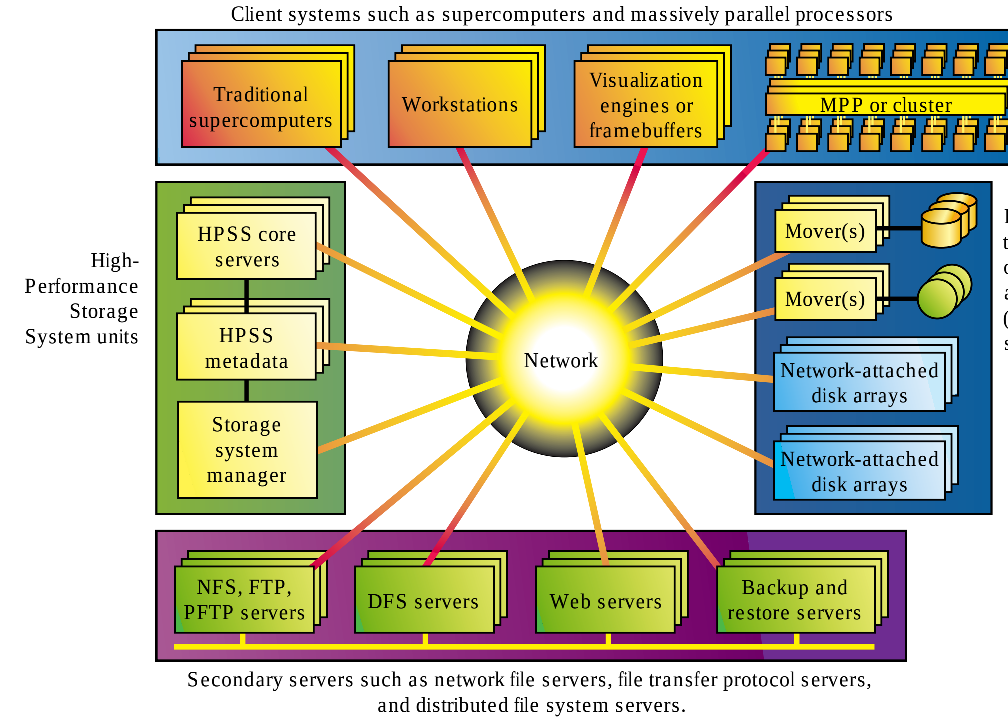


The High-Performance Storage System eliminates the bottlenecks and scalability limits of server-centered architecture by connecting the storage units directly to a high-speed network. This network-centered architecture allows the data to flow directly between client systems and storage units, bypassing the storage server. Throughput is scalable to and beyond the rate of a gigabyte (a billion bytes) per-second.

disk or tape controllers to a client once an HPSS server has established a transfer session. Its interfaces support parallel or sequential access to storage devices by clients executing parallel or sequential applications. HPSS can even manage data transfers in a situation where the number of data sources and destinations are different. Parallel data transfer is vital in situations that demand fast access to very large files and to reach the high data transfer rates of present and future supercomputers.

All aspects of HPSS are scalable so that the storage system can grow incrementally as user needs increase. The parallel nature of HPSS is one key to its scalability. For example, if a system has a storage device that can deliver 100 megabytes (100 million bytes) per second but a gigabyte (a billion bytes) per second is needed, then 10 devices in parallel, controlled by HPSS software, can be used to "scale up" to the new requirement. With this design, HPSS will be able to handle almost unlimited storage capacity, data transfer rates of billions of bytes per second and beyond, virtually unlimited file sizes, millions of naming directories, and hundreds to thousands of simultaneous clients.

HPSS uses several mechanisms to ensure data reliability and integrity. An important one is the use of transactions, which are groups of operations that either take place together or not at all. The problem with distributed servers working together on a common job is that one server may fail or not be able to do its part. Transactions assure that all servers successfully complete their job or the function is aborted. Although transactional integrity is common in relational data management systems, it is new in storage systems.

HPSS was designed to support a range of supercomputing client platforms, operate on many vendors' platforms, and use industry-standard storage hardware. The basic infrastructure of HPSS is the Open Software Foundation's Distributed Computing Environment because of its wide adoption among vendors and its almost universal acceptance by the computer industry. The HPSS code is also available to vendors and users for transferring HPSS to new platforms.

The principal HPSS development partners are IBM Worldwide Government Industry and four national laboratories—Lawrence Livermore, Los Alamos, Oak Ridge, and Sandia. There have been two releases of HPSS thus far, and IBM is marketing the system commercially. HPSS has already been adopted by the California Institute of Technology/Jet Propulsion Laboratory, Cornell Theory Center, Fermi National Accelerator Laboratory, Maui High-Performance Computer Center, NASA Langley Research Center, San Diego Supercomputer Center, and the University of Washington, as well as by the participating Department of Energy laboratories.

In combination with computers that can produce and manipulate huge amounts of data at ever-increasing rates, HPSS's scalable, parallel, network-based design gives users the capability to solve problems that could not be tackled before. As computing capacity and memory grow, so will HPSS evolve to meet the demand.

*—Katie Walter*

**Key Words:** computer network, hierarchical storage management, large-scale computer storage, parallel computing, supercomputing.

*For further information, contact Dick Watson (510) 422-9216 (dwatson@llnl.gov) or visit the HPSS Internet home page at http://www.sdsc.edu/hpss/.*

# A Powerful New Tool to Detect Clandestine Nuclear Tests

Experiments with the Laboratory's new method of detecting clandestine nuclear tests were conducted on the rocky Rainier Mesa at the Nevada Test Site during periods of low atmospheric pressure, mainly at the beginning of storms, so that tracer gases could rise toward the surface through natural faults and fractures.

WHEN President Clinton and other world leaders signed the landmark Comprehensive Nuclear Test Ban Treaty last September, they served notice that any signatory nation trying to conceal an underground nuclear test would have to elude a vigorous international verification program armed with the latest monitoring technologies. Thanks to the work of a multidisciplinary Lawrence Livermore team, the international community now has a powerful new forensic tool to help enforce the treaty by detecting even deeply buried clandestine nuclear tests.

Under the terms of the treaty, which bans all nuclear weapons test explosions, a system of verification and inspection will be administered by the Comprehensive Test Ban Treaty Organization in Vienna, Austria.

Lawrence Livermore scientists have long played an important role in providing monitoring technologies in support of nuclear treaty verification and on-site inspection. The latest Livermore technology is based on the discovery that minute amounts of rare, radioactive gases generated in underground nuclear detonations will migrate toward the surface along natural fault lines and earth fissures.

Livermore geophysicist Charles Carrigan led the team that included physicists Ray Heinle, Bryant Hudson, and John Nitao and geophysicist Jay Zucca. With the help of results from earlier studies, they theorized that highly sensitive instruments might detect telltale radioactive gases rising during periods of barometric low pressure through natural fissures in the ground above the blast. To test the hypothesis, the team obtained two gases, 0.2 kilograms (7 ounces) of helium-3 and 50 kilograms (110 pounds) of sulfur hexafluoride, as tracers. These nonradioactive gases are ideal tracers because they are present in very low quantities in the natural environment.

As the photo on p. 25 shows, the bottles containing the gases were placed with a 1.3-kiloton charge of chemical explosives into a mined cavity that was 15 meters (50 feet) in diameter and 5 meters (17 feet) high. The cavity was located 400 meters (1,300 feet) below the surface, two to three times deeper than that required for a similar sized underground nuclear test. A somewhat shallower detonation, says Carrigan, might have produced a collapse crater or extensive fractures connecting the cavity with the surface, both telltale signs of an underground explosion. Hence, clandestine tests would very likely be conducted at the greater depth to avoid easy detection of treaty violations.

## Simulating a Nuclear Test

The detonation, known as the Non-Proliferation Experiment, occurred on September 22, 1993, in the rocky Rainier Mesa of the Nevada Test Site, where some of the nation's nuclear tests





were conducted until a testing moratorium went into effect in 1992. The chemical explosion simulated a 1-kiloton underground nuclear detonation, which, as expected, did not produce any visible new cracks in the Earth.

Over the year and a half following the blast, team members, including technical support personnel from Test Site contractors EG&G and REECo, collected nearly 200 samples of subsoil gases for measurement. At some sampling stations, sampling tubes were driven into the ground to depths of 1.5 to 5 meters (5 to 16 feet) along fractures and faults. At other stations, tubes were simply placed beneath plastic sheeting that was spread on the ground to trap rising soil gases and to limit atmospheric infiltration (see photo, p. 26).

The first positive finding came 50 days after the explosion, when sulfur hexafluoride was detected in fractures along a fault. Interestingly, the much lighter helium-3 showed up 375 days—more than a year—following the explosion. Both gases were first detected along the same natural fissure within 550 meters (1,800 feet) of the blast site.

Over the course of the extended sampling period, virtually all the samples yielding concentrations of the two tracers appeared along natural faults and fractures in the mesa during periods of low atmospheric pressure, mainly at the beginning of storms. The low pressure accompanying storms, says Carrigan, makes it possible for the gases to move toward the surface along the faults. Although over the course of a year the number of low-pressure days equal the number of high-pressure days, the gases are eventually drawn upward. "There's a ratcheting effect," he explains. "The gases don't go back down as much as they go up." (See the simulation on p. 26.)

Carrigan notes that it is counterintuitive that helium-3 takes so much longer to make its way up natural fissures than sulfur hexafluoride, which is 50 times heavier. Computer models developed at Livermore showed that this result occurred because most of the heavier sulfur hexafluoride gas moved directly up the rock fractures. In contrast, the helium-3 diffused readily into the porous walls of the rocks as it slowly moved upward toward the soil surface. Critical to determining why helium-3 behaved as it did was Bryant Hudson's analysis of helium-3 in Livermore's noble gas laboratory, where he used mass spectrometry to measure the presence of helium-3 in soil-gas samples down to parts per trillion.

## Modeling the Detonation

Carrigan and Nitao modeled the experiment using a porous-flow simulation software called NUFT (Non-Isothermal Unsaturated Flow and Transport) developed at LLNL by Nitao. In attempting to make the simulation as realistic as possible, the team used actual barometric pressure variation data from the Rainier Mesa weather station. The simulation showed the two gases moving at different rates toward the surface following the detonation. The calculated arrival times at the surface for both tracers were in excellent agreement with the data.

Given the good agreement between the computer model and the observations, the team then used NUFT to simulate the gases released from an underground 1-kiloton nuclear test under atmospheric conditions similar to those that followed the 1993 Non-Proliferation Experiment. The software was used to predict the arrival of detectable concentrations of the rare gases argon-37 and xenon-133 at 50 and 80 days, respectively, after the detonation.

These two isotopes are ideal indicators of nuclear explosions because they are not produced naturally in significant quantities; thus, background levels are extremely low. Also, their short half-lives of 34.8 days and 5.2 days can be used to infer how recently an event had occurred. Other, more long-lived isotopes might still be present in the environment from decades-old tests and would tend to muddy the conclusions of investigators trying to determine whether a clandestine test had recently occurred.

The successful confirmation of the experiment by computer simulation implies that sampling of soil gases for rare, explosion-produced radioactive tracer gases at the surface near a suspected underground test can be an extremely sensitive way


A bottle of sulfur hexafluoride gas is separated from the explosives in a mined test cavity to prevent thermal decomposition of the tracer gas during detonation.

A crew of scientists from Livermore and the Nevada Test Site collect soil-gas samples from tubes inserted to a depth of 5 meters (16 feet) in soils that cover rock containing geologic faults and fractures. The soil gases were detected following a contained, 1-kiloton, underground chemical explosion 400 meters (1,300 feet) beneath Nevada's Rainier Mesa.

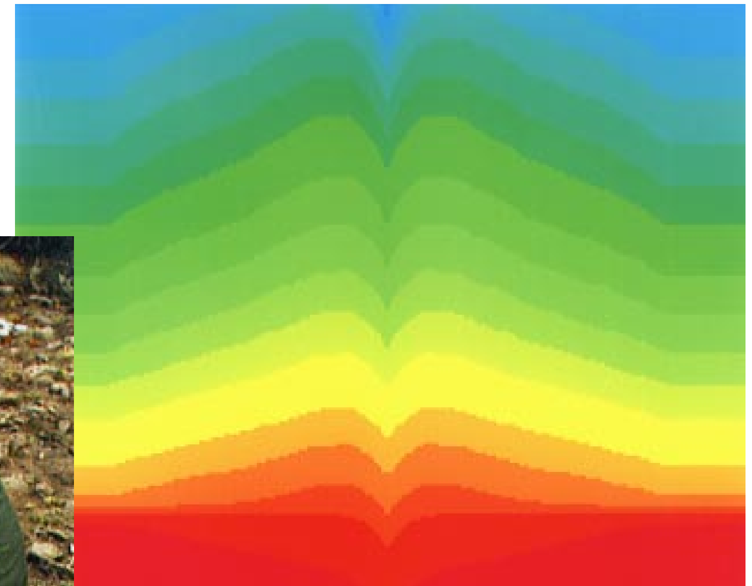

Using Livermore's NUFT (Non-Isothermal Unsaturated Flow and Transport) simulation software, the team was able to model gases moving toward the surface following detonation. Shown is a "rainbow" simulation of barometric "ratcheting" of trace gas in the porous walls of a 300-meter- (985-foot-) long, 0.001-meter- (0.03-inch-) wide vertical fracture (centerline of graphic). Concentration decreases from red near the detonation to blue at the surface as surface pressure variations cause the tracer gas to move up and down the fracture until it eventually reaches the surface.

to detect nearby underground nuclear explosions that do not fracture the surface. As a result, says Carrigan, an on-site inspection has a good chance of finding conclusive evidence for a clandestine nuclear explosion for several months afterward.

### Putting Treaty Evaders on Notice

"If detected, the radioisotope signals would be unequivocal," according to Bryant Hudson. "They would put treaty evaders on notice that they risk detection if they try to explode a nuclear device underground. We can't absolutely guarantee there won't be cheating, but we've made it more difficult."

Carrigan points out that because of political considerations, it may take some time to get a country to agree to an on-site inspection under the terms of the test ban treaty. The thinking of many experts has been that such inspections need to be conducted within a few days to capture evidence of a test. The Livermore team's work, however, shows that waiting weeks or even months to detect rare gases is not a problem and may well be advantageous, because the gases need time to arrive at the surface.

Team members caution that searching for tracer gases is only one of many detection tools. Other methods that might be used at a suspected test site include analyzing the printouts of seismographs for aftershocks from an explosion, looking for explosion-induced stress in plants and trees, drilling for explosion debris, examining the earth for fractures and craters, and searching for pipes and cables leading underground.

In discussing the work of the team, Carrigan attributes its accomplishments to a confluence of Lawrence Livermore strengths in computer simulation, geophysical theory, nuclear test containment, and radiochemistry. "Interdisciplinary collaboration made this work possible," he says.

*—Arnie Heller*

**Key Words:** Comprehensive Nuclear Test Ban Treaty, nuclear proliferation, nuclear treaty verification, NUFT (Non-Isothermal Unsaturated Flow and Transport).

*For further information contact Charles Carrigan (510) 422-3941 (carrigan1@llnl.gov).*





## Patents and Awards

*Each month in this space we report on the patents issued to and/or the awards received by Laboratory employees. Our goal is to showcase the distinguished scientific and technical achievements of our employees as well as to indicate the scale and scope of the work done at the Laboratory.*

### Patents

| Patent issued to | Patent title, number, and date of issue | Summary of disclosure |
|---|---|---|
| John W. Elmer<br>Alan T. Teruya<br>Dennis W. O'Brien | Modified Faraday Cup<br><br>U.S. Patent 5,554,926<br>September 10, 1996 | A rotatable modified Faraday cup incorporating tungsten slit blocks that are machined with an included angle, such as 10 degrees, and face away from the beam. The Faraday cup is used in a computer tomographic technique to measure the current density distribution in an electron beam swept across a narrow slit in the Faraday cup that is rotated in a stepped arrangement such that the beam waveform is recorded at regularly spaced angles by a digitizing storage oscilloscope. The recorded waveform provides the input for the computer tomographic technique. |
| Richard H. Sawicki | Apparatus for and Method of Correcting for Aberrations in a Light Beam<br><br>U.S. Patent 5,557,477<br>September 17, 1996 | A technique that uses two optical elements to correct for aberration in a light beam. The first optical element defines a flat, circular light-reflecting surface having opposite reinforced circumferential edges and a central post and is resiliently distortable, to a limited extent, into different concave and/or convex curvatures, which may be Gaussian-like, about the central axis. The second optical element acts on the first element to adjustably distort the light-reflecting surface into a particular, selected one of the different curvatures depending upon the aberrations to be corrected for and to fixedly maintain the curvature selected. |
| Richard W. Pekala | Organic Aerogels from the Sol-Gel Polymerization of Phenolic-Furfural Mixtures<br><br>U.S. Patent 5,556,892<br>September 17, 1996 | A phenolic-furfural aerogel produced by sol-gel polymerization of phenolic-furfural mixtures. The sol-gel polymerization of a phenolic-furfural mixture in dilute solution leads to a highly cross-linked network that can be supercritically dried to form a high-surface-area-foam. The new organic aerogel may be carbonized by pyrolyzing in an inert atmosphere at 1,050°C to produce a carbon aerogel. These porous materials have cell–pore sizes of less than 1,000 angstroms, density of 0.1 to 1.0 grams per cubic centimeter, and surface area of 350 to 1,000 square meters per gram. Dopants can be included in the aerogel. |

### Awards

**Bruce W. Shore**, a physicist in the Laboratory's Laser Programs Directorate, has been named winner of a **Humboldt Research Award** that will take him to the University of Kaiserslautern in Germany for a year of collaborative research with Professor Klaas Bergmann. The award recognizes the value of Shore's five-year collaboration with Bergmann on the behavior of atoms and molecules exposed to laser radiation. The award was made by the Alexander von Humboldt Foundation, which was established in 1860 to honor the German scientist, explorer, and advocate of international scientific collaboration who gives his name to Humboldt County. Today, the foundation provides opportunities for scholars from the U.S. to work with colleagues in Germany. Shore, a 25-year Livermore employee, specializes in optics and laser physics and is the author of the two-volume text, *The Theory of Coherent Atomic Excitation* (1990), and of numerous scholarly articles on theoretical atomic physics.

**David Seibel**, head of the Laboratory's Administrative Information Systems Department, has been honored for his long-term executive leadership with a 1996 **Information Management Quality Award** from the Department of Energy. Seibel, a Livermore employee for 27 years, was recognized for "providing nearly two decades of continuous vision for the future of business information systems at LLNL." DOE Chief Information Officer Woody Hall presented the award to Seibel at an October 30 ceremony in Denver.