

The logo for ACTIVE Archive features the word "ACTIVE" in white on a black background, followed by "Archive" in white on an orange background. The logo is set against a background of orange and black curved lines that create a sense of motion and depth.

ACTIVE Archive

CHALLENGE:

A global scientific community comprising about 5,000 users accessing massive and growing volumes of data required a storage infrastructure optimized for active archive usage.

SOLUTION:

Tape-based Active Archive using four automated tape libraries with pass-through capabilities – 128 tape drives, 35 of which are fast-access devices while 93 are high-capacity devices.

RESULTS/BENEFITS:

- *2009 internal study showed 99.9991 percent reliability rate*
- *Maintained low cost per GB with tape*
- *Ability to scale to support exponential growth, without corresponding budget increase*
- *Achieved data transfer rates of up to 1GB/second*
- *Worked to accelerate and streamline data access times for users*

Why Tape-Based Active Archives?

- *High data reliability*
- *Extremely cost effective and energy efficient*
- *Massive Scalability*
- *Simplified long-term storage management*

CASE STUDY

National Energy Research Scientific Computing Center

NERSC Exceeds Reliability Standards With Tape-Based Active Archive

Research Facility Accelerates Access to Data while Supporting Exponential Growth

Founded in 1974, the National Energy Research Scientific Computing Center (NERSC) is the primary scientific computing facility for the Office of Science in the U.S. Department of Energy. NERSC is located at Lawrence Berkeley National Laboratory's Oakland Scientific Facility in Oakland, California and is mandated with providing computational resources and expertise for scientific research to about 5,000 scientists at national laboratories and universities across the United States, as well as their international collaborators. A division of Lawrence Berkeley National Laboratory, NERSC supports more than 400 different projects at any given time, and maintains a growing archive of more than 140 million files.

NERSC also boasts the 8th fastest computer in the world according to the November 2011 listing of TOP500 supercomputing sites, with 1.3 petaflops of processing power. The facility is connected to a network called ESnet, which facilitates the transfer of large scientific data sets between NERSC and other supercomputing centers and experimental facilities around the world.

"NERSC provides some of the largest open computing and storage systems available to the global scientific community," said Jason Hick, Storage System

Group Lead at NERSC. "At any given moment, there are about 35 people logged into the archive system, including users from as far away as Europe or Asia, to researchers at various universities across the United States. We support a global distribution of users in every sense."

Delivering and Managing Increasing Data

Data growth typically hovers between 50 to 70 percent each year. The NERSC archive takes in approximately 50 TB of data each and every day – roughly the equivalent of its annual data intake a decade ago. With an indefinite retention period on the archive, NERSC's storage capacity and reliability requirements are significant.

NERSC stores data from some of the largest experimental devices in the world, including the Large Hadron Collider in Europe, the Daya Bay Neutrino detector in China, the Planck satellite, and the Joint Genome Institute in nearby Walnut Creek. In addition, the archive stores data from climate modeling, protein folding, and diverse set of energy and science simulations performed at NERSC. "The archive is a critical piece of the increasingly data-driven scientific picture," says NERSC Director, Kathy Yelick. "Scientists can share enormous data sets with colleagues and can come back months or even years later to reanalyze it from a different perspective or confirm scientific findings."

CASE STUDY Active Archive Alliance

Because NERSC supports such a broad range of scientific projects, ranging from some of the largest computational problems to data-intensive science applications, it required a flexible, efficient archival storage solution to meet the varying needs of its global users. Providing storage for scientists' findings was a key objective, as was the ability to access data quickly when needed. In order to keep costs under control, NERSC needed a storage solution that could keep pace with data proliferation. And finally, it needed to be reliable at scale to ensure user data was available when needed.

Accessing Data with Active Archive

Given the scientific community's heavy reliance on data, and its frequent need to refer to older findings to support current initiatives, NERSC maintains archived data indefinitely. Relying on disk drives alone would be cost prohibitive given the center's data growth and corresponding storage scalability needs. To address this, NERSC implemented an Active Archive, a combined storage solution consisting of the High Performance Storage System software, disk, and tape hardware. The Active Archive provides an effortless way for users to access all of their data. It also simplifies data storage and management for both researchers and storage administrators while keeping costs down.

"If data can't be archived indefinitely, it shouldn't be called an archive," said Hick. "Implementing an Active Archive was ideal for NERSC. We could support the high read rates that our users demanded while retaining data efficiently, reliably and cost effectively."

Users can access their data from any of NERSC's systems, using Lustre for local scratch and IBM's GPFS as its center-wide file system. Tape accounts for more than 75 percent of the storage environment, and the tape-based Active Archive boasts approximately a 30 percent daily read rate. This is in stark contrast to traditional tape archives that typically handle 95 percent or more write operations.

Exceeding Reliability Standards

Between July 2009 and December 2010,

“Tape technology is delivering double the capacity every two years, which is very closely aligned with our typical annual growth rate of between 50 and 70 percent,” said Hick. “At that rate, we can effectively keep a stable budget and still handle exponential growth with the technology of our Active Archive – and that’s exciting!”

NERSC began replacing its existing tape infrastructure with newer versions of tape in its Active Archive. Hick and his team migrated 40,489 tape cartridges during this time, which involved reading 22,065,763 meters of tape – the same distance as flying from San Francisco to Tokyo to Paris to Nova Scotia. The tapes ranged in age from two to 12 years.

During this massive migration process, Hick and his team observed the center's actual tape data reliability within its Active Archive. The findings flew in the face of conventional wisdom: 99.9991 percent of tapes were 100 percent readable, representing a 0.00009 percent error rate. Of the more than 40,000 tapes that were read, only 35 contained some data that couldn't be accessed. The unreadable data accounted for only 178 meters of the 22,065,763 total meters of tape.

"We constantly hear that tape is risky, but that hasn't been our experience," said Hick. "In fact, our recent migration further validated our belief that tape is reliable, and supported our long-standing practice of keeping a single copy of data."

Achieving this level of reliability on tape doesn't come without effort, but having a single copy of data on tape encourages regular intervention when issues arise.

Tape accounts for the vast majority of NERSC's storage capacity. NERSC's 26 PB of data is split between two distinct systems: one system stores 14PB of backup data,

while the other system acts as a scientific archive and holds 12 PB of data. NERSC's tape environment is optimized for its Active Archive usage and contains a mix of access and capacity-oriented enterprise tape drives.

A New Level of Scalability & Efficiency

In addition to the proven reliability of NERSC's Active Archive tape infrastructure, the center's investment in tape means that it can continue to support its exponential data growth, without an exponential budget.

Early adoption of new tape capacity has also offered immediate operational savings for NERSC. The new tape capacity reduces their tape media consumption and reduces the need to purchase more tape libraries to hold increasing amounts of data. Typically, NERSC cuts its media costs in half upon adopting new capacity tape technology.

NERSC's users also benefit from data transfer times of up to 1 GB/second from the archive. In a scientific environment that creates and pushes enormous volumes of research data, quick access to large files eliminates waiting time and results in greater productivity. With a daily read rate of 30 percent, it's clear that NERSC's community is defining its archival storage system as an Active Archive.

Looking Ahead

One of NERSC's future initiatives includes a remote disaster recovery strategy. Hick and his team want to offer users the option of placing a second copy of their data at a remote site. This will be ideal for users who want extra assurance that their data will be "disaster protected," according to Hick.

In the meantime, thanks to statistics that back up the reliability, efficiency, and cost effectiveness of its existing tape-based Active Archive infrastructure, NERSC will stay the course with tape and maintain a single copy of its data.