



**CHALLENGE:**

To implement a large data storage solution that could handle the Blue Waters increased data capacity while maintaining an extremely high service level for its end users. The goals were to have an environment tightly integrated with the file system on the Blue Waters machine, which is faster than a terabyte per second, and provide an archive system that would scale to 100 GB per second. Some key requirements were high enterprise-level performance, reliability, ready data accessibility and massively scalable capacity.

**SOLUTION:**

File-based active archive to enterprise tape:

- IBM HPSS software solution with DataDirect Networks disk cache
- Spectra Logic T-Finity enterprise tape library
- IBM TS1140 3592 JC enterprise data tape with 4.0TB native capacity

**RESULTS/BENEFITS:**

RESULTS of tape solution:

- 99.99% uptime
- 224 TS1140 drives at 240MB/s each currently at 52.7GB/s aggregate
- Expected to grow to 380PB of data storage over the course of five years.

**CASE STUDY: NCSA**

National Center for Supercomputing Applications

**NCSA Employs Active Archive for World's Largest Active File Repository**

**Active Archive for Blue Waters Meets Data-Intensive Needs of Scientists and Engineers**

Supported by the National Science Foundation (NSF), the state of Illinois and the University of Illinois, the Blue Waters project is designed to meet the compute-intensive, memory-intensive, and data-intensive needs of a wide range of scientists and engineers. The National Center for Supercomputing Applications (NCSA) provides 100% of the nearline storage for Blue Waters.


Scientists use the Blue Waters supercomputer for a diverse set of applications from predicting the behavior of hurricanes and tornadoes, analyzing complex biological systems, studying the evolution of the universe and simulating complex engineered systems like the power distribution system in airplanes and automobiles.

To handle its intensive data storage challenges, the NCSA implemented

an active archive consisting of Spectra Logic T-Finity tape libraries and IBM TS1140 Technology enterprise tape drives with a capacity of 4.0TB native per tape cartridge.

The Blue Waters system is one of the world's largest active file repositories stored on high capacity, highly reliable enterprise tape media. It held 8PBs of archived data after being in operation for a mere six months, and is expected to scale to a capacity of 380PB, which is the equivalent of 5,054 years of HD-TV video.

The active archive, which stores the supercomputer's nearline data, is connected to the Blue Waters system and uses IBM's HPSS HSM software. In the initial implementation, users of Blue Waters must move their own data from the file systems on Blue Waters to the nearline system. In the future, NCSA plans to enable seamless data movement between the disk cache and



# CASE STUDY Active Archive Alliance

the active archive tape storage based on enhanced Lustre software tree capabilities.

The active archive has 244 tape drives that provide the performance required today, and will add more tape drives to reach its targeted 100GB per second transfer rate between the disk cache and tape. NCSA has four Spectra Logic libraries, each with approximately 16,000 media slots, for a total of 64,000 media slots.

These tape libraries provide the Blue Waters project with the ability to cost effectively keep all nearline data accessible in an active repository, perform automated data integrity verification for the data store, and deliver high performance read/write rates.

## High Performance, Large Capacity

The active archive file system exceeds read/write rates of 58GB/s, and boasts 99.99% reliability based on uptime. The new system is also able to handle significantly more capacity as the amount of data escalates over the coming years. On NCSA's previous system, it took 19 years to store the first petabyte of data and another year to store the second petabyte of data. With Blue Waters, in the first six weeks NCSA already had the first petabyte of data stored.

“Using an active archive system along with enterprise tape and RAIT eliminates the need to duplicate tape data, which has led to dramatic savings. Having an active archive solution is essential for keeping all of our long-term data accessible and cost effective.”

*Michelle Butler, NCSA Senior Technical Manager*

With NCSA's active archive solution, users can now keep more data than previously allowed and have the confidence that their data is stored safely, securely and with easy access. The active archive allows users to make use of a larger data environment for the experiments and applications they run on the Blue Waters machine.

## Ease of Use

The active archive provides transparent file storage for users, and allows them to easily retrieve a file without needing to know where it is located. The system

provides a single interface for locating data, plus the ability to request information by file name. Users see data in an active file system, while the system itself moves the data between storage tiers based on policy and the size of the file. This data archive system allows the scientists and engineers to focus their efforts on core research with the assurance that their data is protected and easily accessible.

“In this high-performance computing environment, we are able to provide the best in class resource for our users. Making them successful makes us successful,” said Michelle Butler, Senior Technical Manager at NCSA.

## Room to Grow

There is plenty of room on the Blue Waters system. The active archive storage solution allows the NCSA to store the data without having to constantly watch its disk cache. Within the first six months, NCSA quickly ramped up its capacity usage, storing 8PB of data. The environment consists of 50 DELL 720XD servers that have two 40G Ethernet links, IB FDR for disk cache connectivity and then fiber channel cards for direct connect to the enterprise tape system. In addition, the Blue Waters nearline provides enough capacity to easily grow up to 380PB within the next five years.

## Dramatic Savings

The active archive enables NCSA to keep all of its nearline data accessible in an active repository and provides users online access to all data while also reducing the complexity and cost normally associated with very large datasets. It no longer has to duplicate tape to another system because of Redundant Array of Independent Tape (RAIT) software and therefore the total cost of ownership is much lower than it was before.



*A panoramic view of the Blue Waters complex including Sonexion units and networking hardware. Top image is merely the left half of the panorama shot, and the bottom image is the right half.*