ACTIVE Archive

# CASE STUDY: ORNL
## Oak Ridge National Laboratory

### CHALLENGE:

*Oak Ridge National Laboratory is a premier source of computational science research and home to the United States' most powerful supercomputer for open science, Titan. Every month the leadership-class machine generates almost 1 PB of archive data, which must be managed and stored for easy and efficient file recall beyond the timeframe of the normal purge policy of the centerwide file system. To account for continued data growth, the laboratory needed new technology that would offer data redundancy and integrity for its scientific community of more than 1,200 users worldwide.*

### SOLUTION:

*A file-based active archive featuring:*

- *HPSS software*
- *Six enterprise-class tape libraries with the capacity for 60,000 tapes*
- *RAIT 4+1 technology, which provides faster throughput and parity*
- *An 18 PB disk cache*

# ORNL Enhances Data Integrity and Accessibility with Active Archive Solutions

Oak Ridge National Laboratory (ORNL), managed by UT-Batelle, LLC, is currently the largest US Department of Energy (DOE) science and energy laboratory, conducting basic and applied research to solve energy and security problems. Since its establishment in 1943, ORNL has grown to employ more than 4,500 staff members, including scientists and engineers in more than 100 disciplines. ORNL's $1.5 billion budget enables the laboratory to support DOE's national missions of discovery, clean energy, and security through leadership in four major areas of science and technology: neutrons, computing, materials, and nuclear.

A majority of ORNL's research in these categories is published in open literature, and many of its facilities are open to researchers from universities, other national laboratories, and industry.

**ORNL's Titan Supercomputer**

ORNL is home to the United States' most powerful supercomputer for open science, Titan. Titan is capable of 27 petaflops and can handle quadrillions of calculations simultaneously for scientific simulations. Located at the Oak Ridge Leadership Computing Facility (OLCF), a DOE Office of Science User Facility at ORNL, the system provides the ability to simulate increasingly complex and realistic models and reduce time to solution.

OAK RIDGE National Laboratory

# CASE STUDY Active Archive Alliance

## RESULTS/BENEFITS:

- *TAPE - Tape has a lower cost per GB of storage than disk, is more reliable and has superior archivability (30+ year lifetime). Tape is also scalable for future demands.*
- *RAIT - (Redundant Array of Independent Tapes)*

*The overall mount time with RAIT is shorter than serial mounting and the only extra cost is the parity drive. RAIT provides data redundancy at a fraction of the cost of multiple copies.*

- *18 PB DISK CACHE - The large disk cache enables faster file recall and serves as a buffer by ingesting data and then writing it to tape. Data stays in the large disk cache for a longer period of time to allow quicker access to archived data.*

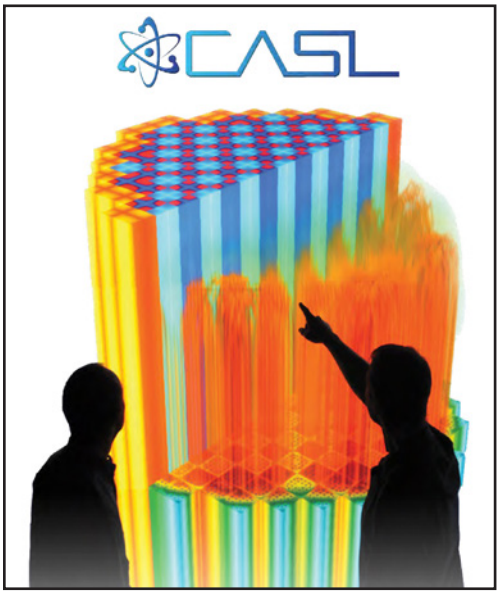


*An image of advanced modeling and simulation research for a light water nuclear reactor.*

**Scientists gain access to the supercomputer and other OLCF resources through three allocation programs:**

**1.** The Innovative and Novel Computational Impact on Theory and Experiment program, which supports larger, high-impact projects.

**2.** The Advanced Scientific Computing Research Leadership Computing Challenge, which supports projects related to the energy mission of DOE's Office of Science.

**3.** The OLCF's Director's Discretionary program, which aids new high-performance computing users in exploring nationally important topics.

At any given time, more than 180 projects are housed at the OLCF under the three allocation programs. More than 1,200 users have access to the supercomputer and its file storage systems, and simulation data must be stored so users can quickly and efficiently access datasets as needed. When the Center for Computational Sciences implemented the HPSS for its archival storage in 1998, the center was producing 300 GB of data per month and provided around 1 TB of disk storage space. Today the OLCF produces 3,000 times more data – approximately 1 PB of data per month. Data growth is predicted to continue due to advancing technologies in areas such as climate sciences and high-energy physics.

To keep up with increasing data demands, the OLCF needed an active archive solution that would ensure data integrity and enable fast file recall for its users, who access their data frequently. It also needed a solution that would allow for data redundancy without doubling the cost of its archive. This solution had to provide scalability to meet the center's future data demands as it increasingly generates data for some of the nation's largest scientific projects.

**Active Archive Storage Solutions: RAIT and a Large Disk Cache**

The center chose to implement an active archive solution that included Redundant Array of Independent Tapes (RAIT) technology as well as new enterprise-class tape drives and an 18 PB disk cache. RAIT technology stripes incoming data across multiple tapes and includes a parity striped across the tapes, allowing for data reconstruction in the event a tape becomes damaged or lost. RAIT was an appealing active archive solution because the overall mount time with RAIT is shorter than with serial mounting – and the only cost incurred is the extra parity tape.

> **"Our current-generation tapes run at 252 MB a second, but if you're running four of these with a RAIT 4+1, then you'll get a GB a second. That's four times faster than serial writing."**
>
> ***Quinn Mitchell, OLCF High-Performance Computing Storage System Administrator***

Since RAIT was implemented in January 2015, the center has placed more than 15 PB and 42,585,959 files – on RAIT tapes. The rate of migration for different tape technologies varies, but rough estimates indicate the center will have all but the most difficult tapes migrated to RAIT before the fall of 2017. The center regularly purchases new technologies when they are released and is now using the newest tape

# CASE STUDY Active Archive Alliance

> **“ These active archive upgrades were crucial to ensuring our users’ data is both accessible and fault-tolerant so they can continue performing high-priority research at our facilities. Our storage-intensive users have been very pleased with our new data storage capabilities. ”**
>
> ***Jack Wells, Director of Science for the National Center for Computational Sciences at ORNL***

technology. The center currently has more than 120 tape drives and the ability to house 60,000 tapes. While the archive has the potential capacity for 498 PB of data, it currently has 107 PB of tape storage. Of this storage amount, 59 PB is occupied at present.

The OLCF also implemented a large disk cache to support its active archive solution. The transition from a 2 PB disk cache to an 18 PB disk cache means users can access their data quickly and for longer time periods. The large disk cache offers faster file access because data can live on disk for approximately 18 months before it is removed from the disk cache. The cache acts as a buffer, so the center can ingest large amounts of data and then transfer them to tape over a period of time. Users can now quickly move multi-petabyte projects while benefiting from the long term storage advantages of tape.

Because it can take years for some researchers to analyze datasets, implementing new active archive storage solutions is important for providing researchers with an efficient and cost effective system for storing and retrieving their information for the duration of their respective projects.

**The Future of OLCF’s Active Archive**

Currently the OLCF’s data ingest rate is the highest it has ever been. In 2016 alone the center experienced a 20 percent growth in data. Storing at exascale volumes at the OLCF would only require the addition of another tape library, but achieving exascale data rates is more difficult. With an eye toward future exascale machines, the staff is constantly on the lookout for strategies such as RAIT that might also benefit center ingest and outgest rates in the future.

“It’s those data rates we still need to attain,” Mitchell said. “But we are looking at best-of-breed solutions all the time, whether those be for the disk cache or tape layer or for the application managing those hierarchical storage systems. We are always evaluating our current storage system to find the best active archive solutions to meet both our center’s needs and the needs of the next generation of computational scientists at the OLCF.”

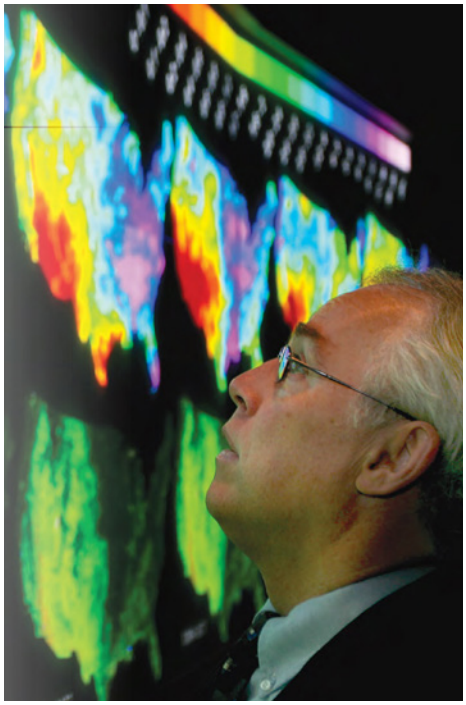

*Titan supercomputer simulation of climate changes in the atmosphere being viewed by James Hack, director of the National Center for Computational Sciences at ORNL.*


*Oak Ridge National Labs is the home of the Titan supercomputer. Titan is a Cray XK7 system that contains 18,688 nodes, each built from a 16-core AMD Opteron 6274 processor and an NVIDIA Tesla K20X GPU accelerator. Titan also has 710 terabytes of memory.*