

Innovative new ways to archive data for supercomputers

March 22 2016



Los Alamos National Laboratory sits on top of a once-remote mesa in northern New Mexico with the Jemez mountains as a backdrop to research and innovation covering multi-disciplines from bioscience, sustainable energy sources, to plasma physics and new materials.

Seagate Technology and Los Alamos National Laboratory (Los Alamos) are researching a new storage tier to enable massive data archiving for supercomputing. The joint effort is aimed at determining innovative new ways to keep massive amounts of stored data available for rapid access, while also minimizing power consumption and improving the quality of data-driven research.

Under a Cooperative Research and Development Agreement, Seagate and Los Alamos are working together on power-managed disk and software solutions for deep data archiving, which represents one of the biggest challenges faced by organizations that must juggle increasingly massive amounts of data using very little additional energy.

"We see huge opportunities to bring power-aware cold storage to the market as organizations—including and beyond the Department of Energy—seek new ways to address data storage in resourceful ways," said Gary Grider, division leader of high-performance computing (HPC) at Los Alamos.

Typically, organizations using HPC systems must migrate their data from fast parallel-file systems to slower, more cost-efficient "cold" storage technologies. This approach, however, is problematic for organizations that require frequent access to research data, because that data tends to reside on static media, which is slow to access and could jeopardize the quality of research—especially when real-time analysis is critical.

The new dynamic storage research project will leverage the technology behind Seagate's recently released ClusterStor A200 system. This system uses automated, policy-driven hierarchical storage management to non-disruptively migrate data off expensive primary storage tiers while keeping it online for fast retrieval, which significantly reduces costs in time and operations. High-density, power-managed prototype disks are

intended be a key element of the research program.

This is the first of several projects Los Alamos and Seagate are planning to collaborate on to enhance, develop, and deploy cutting-edge solutions for some of the most challenging HPC storage problems. The organizations already have a long-standing relationship, with Seagate providing storage for the national laboratory's Trinity supercomputer. The supercomputer, which will soon conduct simulations crucial for national security, relies on the Cray Sonexion 2000 system Powered by Seagate for its scale-out storage capabilities.

"Los Alamos is pushing the frontier of high-performance computing, building systems to solve the most complex technical problems," said Ken Claffey, vice president and general manager, Seagate HPC systems business. "This new collaboration will allow us to explore advanced data archiving methods at the mission-critical government level, which could then influence big data in broader enterprise applications as well."

The archiving research program could result in future products that meet the deep-archiving [storage](#) needs of industries such as cloud providers, weather modeling, financial services, life sciences, and media and entertainment.

"Seagate's collaboration with Los Alamos exemplifies the technology innovations we are helping drive through unique federal partnerships—particularly in the national laboratory and intelligence communities," said Deb Oliver, president, Seagate government solutions. "Federal agencies are on the frontlines of today's [data](#) management challenges, and partnerships like this will help us excel in engineering new systems that not only meet their stringent demands, but can then be put to broader use, as well."

Provided by Los Alamos National Laboratory

Citation: Innovative new ways to archive data for supercomputers (2016, March 22) retrieved 4 May 2024 from <https://phys.org/news/2016-03-ways-archive-supercomputers.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.