



How Spectra Logic & Versity Built an Exascale Archive for the Texas Advanced Computing Center

Thank you for downloading this Spectra Logic case study. Carahsoft is the government solutions provider for Spectra Logic cybersecurity solutions available via GSA 2GIT, NASA SEWP V, NASPO ValuePoint, The Quilt, and other contract vehicles.

To learn how to take the next step toward acquiring Spectra Logic's solutions, please check out the following resources and information:



For additional resources:
carah.io/SpectraResources



For upcoming events:
carah.io/SpectraEvents



For additional Commvault solutions:
carah.io/SpectraSolutions



For additional cyber solutions:
carah.io/Cybersecurity



To set up a meeting:
SpectraLogic@carahsoft.com
888-662-2724



To purchase, check out the contract vehicles available for procurement:
carah.io/SpectraContracts

For more information, contact Carahsoft or our reseller partners:
SpectraLogic@carahsoft.com | 888-662-2724



CASE STUDY

How Spectra Logic & Versity Built an Exascale Archive for the Texas Advanced Computing Center



Image: TACC, Edward Kemper

The Texas Advanced Computing Center (TACC) needed a scalable storage solution capable of managing an exabyte of data generated by Horizon, its new AI supercomputer. Here's why TACC selected Spectra Logic,® Versity, and Dell Technologies for its exascale archive.

THE CHALLENGE

As TACC prepared to launch the Horizon supercomputer, it faced a major storage challenge: **How to manage the massive, AI-driven datasets the system would generate?** Existing infrastructure couldn't scale to exabyte levels. TACC needed a long-term archive that ensured speed, reliability, and flexibility.

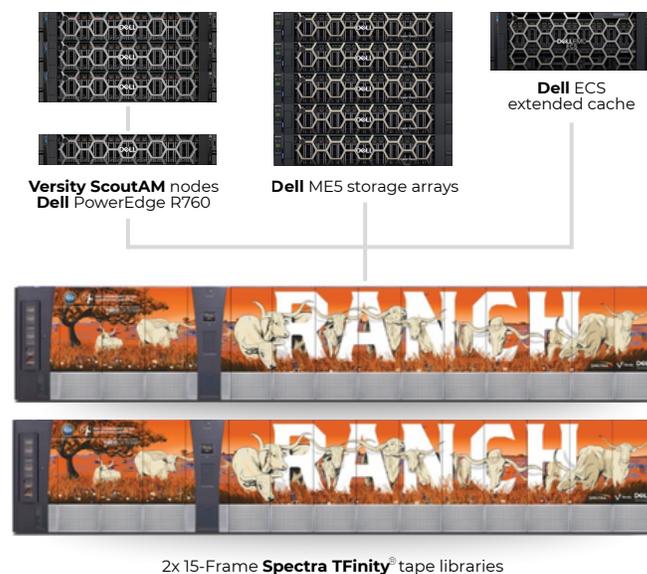
THE SOLUTION

TACC deployed Ranch, a next-generation archive system designed to scale with Horizon's demands. **It combines two Spectra Logic TFinity® tape libraries with Versity ScoutAM software**, supported by Dell servers and high-speed storage tiers. This integrated solution delivers superior scalability, intelligent data management, and disaster recovery framework.

THE RESULT

TACC now has a cost-effective exascale system that can support Horizon's workloads. Ranch safeguards data against downtime or loss and ensures long-term storage reliability. With this robust foundation in place, Horizon can drive scientific breakthroughs while keeping every discovery secure and accessible.

THE SYSTEM



THE ENVIRONMENT SNAPSHOT

- 2x 15-frame Spectra TFinity® tape libraries
- Versity ScoutAM nodes
- 20x LTO-9 drives*
- Dell PowerEdge R760 servers
- 5x Dell ME5 storage arrays
- Dell ECS extended cache

*Future plan: 16 LTO-10 and 4 LTO-9 drives



THE STORY

The Texas Advanced Computing Center (TACC) at The University of Texas at Austin has long led the nation in high-performance computing.

Since its founding in 2001, TACC has operated some of the most powerful academic supercomputers in the U.S. — from Ranger in 2007 to today's 40 petaFLOP/s Frontera.

Now, TACC has entered a new era of this legacy with Horizon, a next-gen, 400 petaFLOP/s high-performance supercomputer that is a cornerstone of the U.S. National Science Foundation Leadership-Class Computing Facility (NSF LCCF).

Horizon isn't just an incremental upgrade. It delivers a tenfold improvement in simulation performance and **a 100x increase in artificial intelligence capabilities** compared to Frontera.

Horizon is the largest supercomputer in the NSF portfolio dedicated to open, academic research — empowering scientists to explore new questions in fields ranging from climate modeling and genomics to quantum mechanics and cosmology.

But with such transformative capability came a significant challenge: How to handle the enormous, AI-driven volumes of data that Horizon will generate?

TACC needed a storage strategy that could reliably, affordably store an exabyte of data at a single site.

ABOUT THE TEXAS ADVANCED COMPUTING CENTER (TACC)



Image: Texas Advanced Computing Center (TACC)

The Texas Advanced Computing Center (TACC) at The University of Texas at Austin is the leading academic supercomputing center in the country. TACC delivers world-class and innovative systems, tools, software, and expertise to researchers who seek to make an impact in the world, and advance discovery across disciplines.

Since 2001, TACC has provided access to a wide range of computational resources — from high-performance computing and AI at scale, to storage and visualization — and offered staff expertise, thought leadership, support, training, and education and outreach.

TACC is also home to the U.S. National Science Foundation Leadership-Class Computing Facility (NSF LCCF), which is poised to revolutionize computational research and discovery.

TACC's approach is people-centered, its funding ensures sustainability, and its staff are dedicated to making the power of computing available to all.

MANAGING EXASCALE DATA QUANTITIES

When you're dealing with data at the petabyte or exabyte scale, one storage medium stands out: tape.

“We’ve always had pretty big tape archives at TACC,” says David Cooper, Senior Systems Administrator of HPC Large-Scale Systems at TACC. “Not only does tape provide more capacity for the money, but it’s also more stable and secure for long-term storage. When data is written to tape, we don’t have to worry about it quite as much as other storage mediums like disk or flash.”

To support the Frontera supercomputer, TACC uses a Quantum Scalar i6000 tape library with LTO-9 drives. While that system meets Frontera’s needs, it could not scale to support the massive influx of data Horizon would generate each year or provide the performance required for exascale AI workloads.

Recognizing this, TACC made a proactive shift to a fully modern archive built for Horizon’s scale. The team chose to implement two Spectra Logic TFinity tape libraries paired with Versity ScoutAM software, Dell PowerEdge R760 servers, and Dell ME5 storage arrays.

Aptly named Ranch, this system is installed in a prominent, high-traffic area of TACC to help educate and inspire a diverse range of visitors — from students to industry leaders.

As a nod to its Texas roots, the libraries are wrapped in **custom graphics** featuring longhorn cattle, blending technical excellence with local pride.



When data is written to tape, we don’t have to worry about it quite as much as other storage mediums.

David Cooper

Sr. Systems Administrator
HPC Large-Scale Systems
Texas Advanced Computing Center (TACC)



Image: TACC, Jorge Salazar

A MATCH MADE FOR HIGH-VOLUME DATA MANAGEMENT

When planning Horizon's new storage archive, TACC knew it needed more than just a larger tape library. It also needed an integrated solution that could scale to an exabyte — and beyond.

Together, **Spectra Logic and Varsity** provide the scalability, openness, and reliability that align with TACC's priorities.

The Spectra TFinity is already trusted by some of the world's leading high-performance computing sites to support multi-exabyte environments, such as CERN, NOAA, and more.

Varsity ScoutAM complements the time-tested storage strength of the TFinity tape library with an open, vendor-neutral software layer designed for exabyte-scale archives that integrates seamlessly with TACC's systems, providing unified S3 and POSIX access.

This proven combination of secure storage and intelligent data management gave TACC confidence in a solution that could meet Horizon's demands for decades to come.

"Scalability was the most crucial factor for us, followed by price — we usually try to get as much storage 'bang for the buck' as we can," explains Cooper.

ScoutAM also introduces a key performance advantage for TACC's researchers: an extended caching system that accelerates data access.

Horizon's users will be generating and analyzing far larger datasets than ever before, so minimizing wait times for data retrieval is critical.

To address this, the Ranch archive includes a 4 PB flash tier and a 16 PB object storage tier as a high-speed cache in front of the tape library. ScoutAM intelligently stages active data across these tiers, **allowing researchers to rapidly retrieve files.**

This system keeps Horizon's massive output both preserved and readily accessible — a key differentiator for AI-driven and simulation-heavy workloads.



Unlike other vendors that tie solutions to specific filesystems or tape formats, the Varsity platform isn't locked down. **It gives us more flexibility.**

Junseong Heo
Storage Manager
Advanced Computing Systems
Texas Advanced Computing Center (TACC)

SUPERIOR TECHNICAL SUPPORT

Horizon will play a key role in the U.S. National Artificial Intelligence Research Resource Pilot (NAIRR) — an initiative launched in 2024 to provide researchers across the country with access to **advanced AI infrastructure**.

With so much agency-backed investment riding on Horizon's success, even brief system downtimes are unacceptable. That made reliable, responsive technical support a top priority for TACC.

“For a system of such a substantial size, inevitably there will be technical problems,” says Cooper. **“We really don’t want to be down or have catastrophic data loss. A high level of support and software recovery systems are very important for us.”**

This is where the combined Spectra and Versity solution truly shines.

On the hardware side, TACC's technical team has already tapped into the Spectra Assisted Self Maintenance (ASM) Kit, an industry-first support solution that minimizes downtime by keeping a curated set of spare parts on-site for immediate use — with remote guidance available when needed.

On the software side, Versity ScoutAM reinforces resilience with a built-in disaster recovery architecture. Every data element and its metadata are stored on the mass storage media, guaranteeing complete recovery without proprietary tools and ensuring long-term accessibility.

ScoutAM supports asynchronous read-only remote replication, mirroring all data to one or more disaster recovery or secondary sites, so the entire system can be quickly and easily rebuilt if needed.

“Versity has a very elegant and simple disaster recovery solution,” says Cooper, noting that it reassured the team that even in the worst case, the archive could be restored with integrity.



We needed a new system that could keep up with the demands of our users and compute clusters.

The Spectra TFinity and Versity ScoutAM solution **allows us to continue the high level of service our users expect.**

David Cooper

Sr. Systems Administrator
HPC Large-Scale Systems
Texas Advanced Computing Center (TACC)

The installation process further highlighted the partnership in action. Versity engineers worked hands-on with software configuration and rapidly responded to feedback.

In parallel, Spectra delivered a seamless deployment for TACC.

“We were impressed with how the Spectra team methodically assembled the library frame-by-frame and aligned the robotics to validate the system,” says Junseong Heo, Storage Manager of Advanced Computing Systems at TACC. **“That’s when we truly realized that the wonderful power of this system is in its simplicity.”**

THE NEXT CHAPTER OF AI COMPUTING

When Horizon launches in 2026, its deployment will mark a bold new chapter in TACC’s commitment to advancing open, academic research at unprecedented computational scales. With the Spectra TFinity libraries providing durable capacity and Versity ScoutAM software ensuring intelligent data management, **TACC’s storage foundation is scalable, resilient, and future-ready.**

Horizon now has an advanced archive that keeps pace with rapid AI-driven workloads, protects data integrity, and preserves discoveries for decades to come.

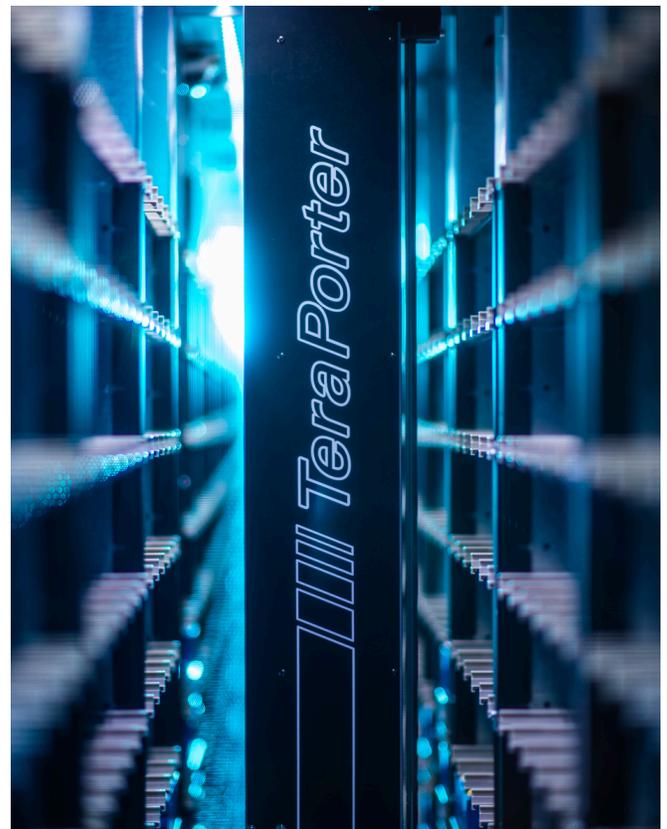
As Horizon powers scientific breakthroughs across disciplines, the storage system built with Spectra and Versity ensures that no discovery is lost — only unlocked.



We’d certainly recommend Spectra to other universities — in fact, it seems like most higher-ed research institutions already use Spectra solutions.

David Cooper

Sr. Systems Administrator
HPC Large-Scale Systems
Texas Advanced Computing Center (TACC)



The Spectra TeraPorter is the roving vertical arm inside the Spectra TFinity library that enables quick, seamless access to archived data.

Image: TACC, Edward Kemper