

Vanderbilt ACCRE Modernizes Research Data Infrastructure with Hammerspace

Thank you for downloading this Hammerspace resource. Carahsoft is the sole federal government distributor for Hammerspace cybersecurity solutions available via NASA SEWP V, 2GIT, and other contract vehicles.

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

 For additional resources:
carah.io/hammerspaceresources




For upcoming events:
carah.io/hammerspaceevents

 For additional Hammerspace solutions:
carah.io/hammerspacesolutions



For additional cybersecurity solutions:
carah.io/cybersecuritysolutions

 To set up a meeting:
Hammerspace@carahsoft.com
703-871-8505



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



Customer

Vanderbilt University ACCRE

Use Case

Campus-wide HPC Resource
and Research Support

Contact Us

P: +1 (650) 750-8400

W: <https://hammerspace.com>

E: contact@hammerspace.com



Vanderbilt ACCRE Modernizes Research Data Infrastructure with Hammerspace

Overview

The Vanderbilt Advanced Computing Center for Research and Education (ACCRE) at Vanderbilt University is a campus-wide high-performance computing (HPC) resource and research support facility. Its core mission is to enable Vanderbilt faculty, researchers, and students to “explore and benefit from the new world of computing,” thereby tackling scientific, engineering, data, and cross-disciplinary questions of societal importance that would otherwise be difficult or impossible.

It supports researchers from many disciplines—genetics, physics, astronomy, engineering, social sciences, and beyond—enabling them to run large simulations, data analyses, machine-learning models, and other compute-intensive tasks.



As data volumes surged across hundreds of research projects, and the need for AI and GPU-accelerated workloads expanded, the Center sought a modern, flexible, high-performance data platform capable of supporting next-generation scientific computing. After a comprehensive evaluation, ACCRE selected Hammerspace to unify its storage infrastructure, reduce storage costs, and gain flexibility to deliver storage resources more to researchers in a more agile way.

The Opportunity to Deliver Better Services and Reduce Costs

ACCRE operates a large-scale HPC environment consisting of:

- 750 compute nodes and 80 GPU nodes with 320 NVIDIA GPUs
- Multiple storage systems, including:
 - ~5 PB of GPFS (legacy primary storage)
 - 5–7 PB of VDURA (current primary storage)
 - 27 PB of LStore-based archive storage developed in-house

The Center's challenges included:

- **Fragmented, Multi-Tier Storage Silos:** ACCRE historically relied on separate systems for performance, capacity, and archive storage, including GPFS, VDURA, and LStore. These silos made it difficult to efficiently move data across storage tiers or to provide researchers with predictable access to data.
- **High Storage Costs:** With expensive legacy primary file systems and storage appliances, storage costs were constraining ACCRE's ability to provide the best services to its researchers.
- **Need for Composability and Agility:** Every research project has different storage performance and capacity requirements. ACCRE needed a way to compose environments dynamically under one unified global namespace.
- **Avoiding Appliance Lock-In:** ACCRE wanted freedom from appliance-based storage systems and instead sought an open system that would allow it to use commodity server hardware across all tiers.

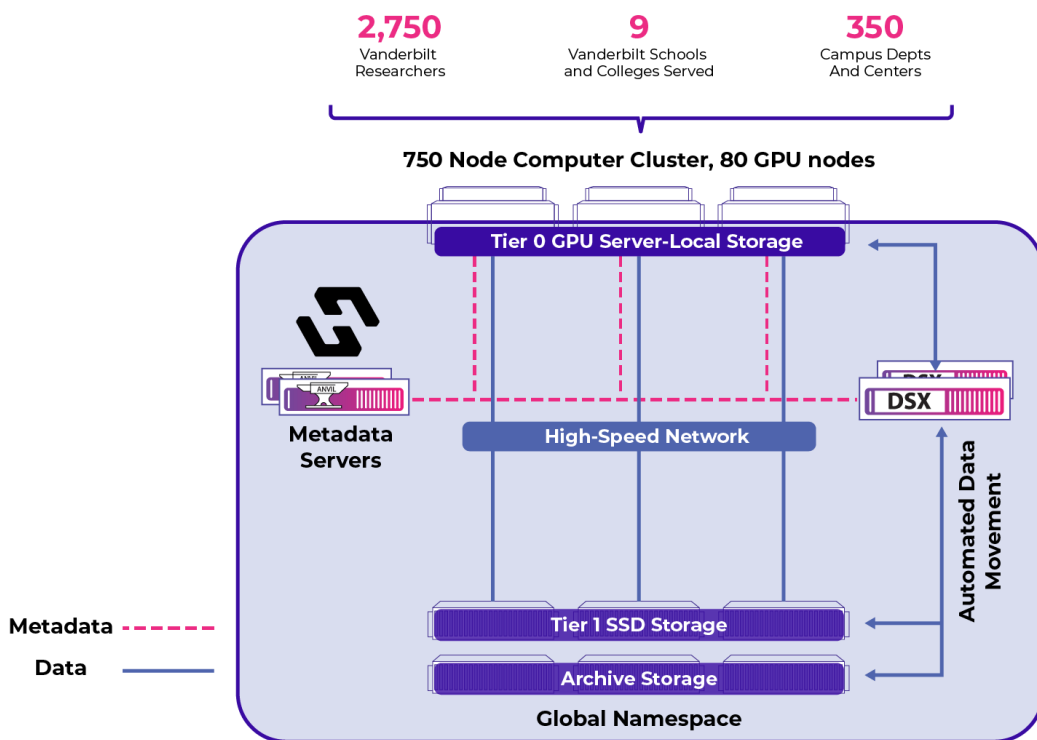
As data volumes surged across hundreds of research projects, and the need for AI and GPU-accelerated workloads expanded, the Center sought a modern, flexible, high-performance data platform capable of supporting next-generation scientific computing. After a comprehensive evaluation, ACCRE selected Hammerspace to unify its storage infrastructure, reduce storage costs, and gain flexibility to deliver storage resources more to researchers in a more agile way.



The Hammerspace Solution

ACCRES selected Hammerspace to deploy a 10-petabyte, next-generation research data environment built on a unified global namespace that spans performance, capacity, and archival storage. This architecture integrates:

- **Tier 0:** GPU/CPU server-local NVMe storage for high-performance AI and simulation workloads
- **Tier 1:** Newly purchased commodity storage servers (~500 TB initial deployment)
- **Tier 2/Archive:** Multi-petabyte capacity from LStore
- **2 Anvils + 4 DSX nodes** supporting metadata services and data orchestration across tiers



This deployment enables ACCRE to dynamically provision the right storage resources to each project, reconfigure environments quickly, and give researchers high-throughput access to data regardless of where it physically resides.

Hunter Hagewood, Executive Director of Research Computing Operations at ACCRE, noted that the Hammerspace Data Platform's "composable, open architecture" was a key factor in the decision, enabling the Center to unify GPU server-local, tiered, and archival storage into a single environment.

"Hammerspace...allows for immense flexibility, immense cost savings, and covers all cluster workload profiles," said Hunter.



Why ACCRE Chose Hammerspace

Unified Global Namespace Across All Tiers

Hammerspace allows ACCRE to present all storage tiers—performance, capacity, and archive—under one global namespace, eliminating silos and enabling seamless movement of data across compute and GPU nodes.

Open, Composable Architecture

The ability to use commodity hardware, dynamically compose storage resources, and avoid vendor lock-in aligned perfectly with the Center’s long-term strategy and existing LStore design principles.

Compelling Economic Advantages

By activating the previously underused CPU/GPU server local storage as a tier of high-performance shared storage, leveraging existing LStore archive capacity, and moving to commodity server hardware, ACCRE will reduce storage costs by 48% over the next few years.

Business Outcomes

✓ 48% Reduction in Storage Costs

Hammerspace enables ACCRE to pair commodity hardware with server-local NVMe and existing LStore resources, cutting their blended cost per TB nearly in half.

✓ Flexible, Composable Storage Delivery Model

ACCRE can now deliver agile storage services across research groups whose requirements shift week-to-week. Storage can be dynamically allocated, reconfigured, and provisioned without data migration.

✓ Unified Namespace for High-Performance Research Workloads

Researchers gain fast, consistent access to data across GPU nodes and storage tiers—improving collaboration and accelerating data-intensive workloads such as machine learning, large-scale simulation, and cross-disciplinary research.

✓ Freedom from Proprietary Appliances

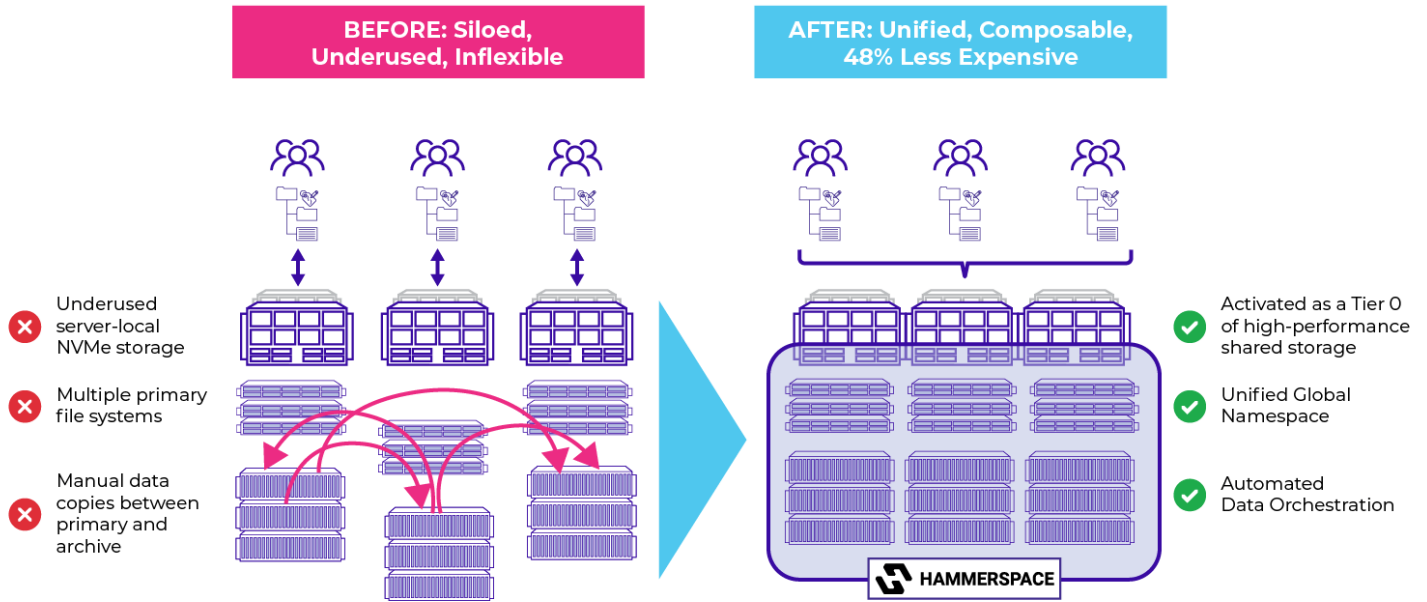
The Hammerspace Data Platform’s open design allows ACCRE to avoid cost-heavy storage appliances and embrace commodity hardware platforms across its next-generation architecture.

✓ Foundation for Future Storage Consolidation

Once Hammerspace proves value in this initial phase, ACCRE plans to expand the footprint and eventually migrate away from legacy GPFS and VDURA systems, fully unifying its storage under Hammerspace.



Conclusion



Vanderbilt ACCRE’s deployment of Hammerspace represents a strategic modernization of its research data infrastructure, one that enables greater flexibility, lower costs, and improved performance across AI, simulation, and data-driven science workloads.

By unifying GPU server-local Tier-0 storage, new commodity Tier-1 storage, and its multi-petabyte LStore archive under a single data platform, ACCRE has positioned itself to support the next generation of scientific discovery with unmatched agility and efficiency.

“Hammerspace’s Data Platform offers a composable, open architecture that lets us unify GPU server-local, tiered, and archival storage into a single data environment,” said Hunter Hagewood, Executive Director of Research Computing Operations at ACCRE. **“Combined with LStore, our largest and most advanced storage platform, we now have a long-term strategy for meeting strong capacity demand. This integration not only cuts our costs dramatically but also changes how we deliver compute and storage services to Vanderbilt researchers — making it easier to support the next generation of data-driven science.”**

