Success Story

# Unifying AI and HPC at the University of Wyoming with VAST Data

## Executive Summary

The University of Wyoming continually seeks to expand its research capabilities, an effort that recently culminated in the institution attaining R1 research status. Supporting this ambitious academic mission is the Advanced Research Computing Center (ARCC), a specialized division providing enterprise-class IT architecture for research to the entire campus. Tasked with modernizing its data environment to keep pace with dynamic high-performance computing (HPC) and artificial intelligence (AI) demands, ARCC deployed the VAST AI Operating System. By consolidating diverse data silos into a single VAST DataStore, the university achieved significantly higher write speeds, substantial data reduction, and a unified architecture that empowered its IT professionals to expand critical research services regionally.

## The Organization: World-Class Research Infrastructure

The University of Wyoming serves approximately 12,000 students and operates as a leading technological and scientific hub within the state. The Advanced Research Computing Center acts as the dedicated IT support arm for the university's Research and Economic Development Division, managing technological infrastructure for researchers.

Unlike many academic institutions of comparable size, the University of Wyoming maintains a highly specialized IT staff operating enterprise-grade systems to support research on campus. This strategic, continuous investment directly helped fuel the institution's recent elevation to R1 classification.

The ARCC team manages immense computing capabilities. Their primary cluster, Medicine Bow, boasts roughly 350 compute nodes encompassing 15,000 cores and 150 NVIDIA GPUs. Additionally, their newer Gannett Peak cluster features 60 NVIDIA GH200 Grace Hopper™ GPUs. Supporting this dense computing infrastructure requires a highly robust underlying architecture.

> *We are the Research and Economic Development Division's IT support arm. We do the whole gamut of enterprise-class architecture to support research on campus.*
>
> **Mike Killean**
> Director, Research Computing Infrastructure, ARCC

Unlike many academic institutions of comparable size, the University of Wyoming maintains a highly specialized IT staff operating enterprise-grade systems to support research on campus. This strategic, continuous investment directly helped fuel the institution's recent elevation to R1 classification.

The ARCC team manages immense computing capabilities. Their primary cluster, Medicine Bow, boasts roughly 350 compute nodes encompassing 15,000 cores and 150 NVIDIA GPUs. Additionally, their newer Gannett Peak cluster features 60 NVIDIA GH200 Grace Hopper GPUs. Supporting this dense computing infrastructure requires a highly robust underlying architecture.

## Overcoming Data Complexity to Accelerate Scientific Discovery

Between 2010 and 2024, ARCC relied on an assortment of distinct storage solutions to support its varying computing workloads. Managing these environments required multiple dedicated administrators. "We had all of these disparate storage systems, and we had to have two or three systems administrators to run all the storage that we had," Killean states.

As modern platforms became available, the team sought an all-NVMe, scale-out architecture capable of high throughput for computational clusters, yet cost-effective enough for static user data via Server Message Block (SMB). Crucially, they refused to rely on complex, legacy storage tiering methods. "We had this idea: first off, we're not doing any tiering, that is just a given. And second, we decided we wanted to go to all-NVMe," notes Killean.

## The Solution: A Singular Data Lake Powered by Vast

In evaluating options, ARCC partnered closely with its long-standing technology advisor, Cambridge Computer. Valued for a decade-plus relationship built on cutting through vendor marketing and prioritizing objective technical guidance, Cambridge Computer introduced the university to VAST Data. This trusted collaboration prompted ARCC to evaluate VAST and its Disaggregated Shared Everything (DASE) architecture. "The VAST technology spoke to us, the DASE architecture did."

ARCC deployed an 8x10 VAST cluster consisting of 32 CNodes and 10 DBoxes, with 3 PB of raw NVMe capacity. This consolidated infrastructure acts as the university's primary research-related storage platform, serving HPC operations over NFS, static SMB data, and archival workloads.

> *For our primary storage, we went from having multiple GPFS systems, using Ceph, using block storage for different purposes, all to VAST.*
>
> **Mike Killean,** Director, Research Computing Infrastructure, ARCC.

# The Results: Scale, Speed, and Simplicity

## Performance throughput

Prior to deploying VAST Data, ARCC considered 30 GB/s throughput a minimum requirement to support their research workloads. VAST fundamentally altered those performance metrics by significantly increasing peak write speeds. “When we did our benchmark and that was what our goal was in 2024, it was 30 gigabytes per second write. With VAST we were able to get 90 gigabytes per second write,” explains Killean.

## Multi-cluster unification

Rather than dedicating separate storage arrays to each compute cluster, the performance of VAST allowed ARCC to unify three distinct compute environments. “We went from having a dedicated storage system per cluster to connecting three clusters with different uses, all using the same centralized VAST DataStore and we’re not exceeding throughput, hitting any limits, or overrunning the system,” notes Killean.

## Economics with VAST Similarity data reduction

By leveraging VAST Similarity Data Reduction, ARCC effectively multiplied its raw capacity by 3 times, making all-NVMe storage highly cost-effective. “We will be able to store up to 9 petabytes due to the three-to-one data deduplication rate,” Killean details.

> *Because we are freeing up FTEs from having to do storage, we’re able to put more time into things that matter most for our users*
>
> **Mike Killean**
> Director, Research Computing Infrastructure, ARCC

## Reclaiming administrative resources

Native features like automatic snapshotting enabled self-service restorations, while streamlined NFSv4 setup simplified administration. By freeing staff from routine storage maintenance, ARCC expanded regional access to Wyoming community colleges as well as institutions like Colorado State University and the Rocky Mountain Advanced Computing Consortium (RMACC).

# A Unified Foundation for Future Discovery

The transition to VAST Data enabled the University of Wyoming to fulfill its vision of a unified, high-performance data lake. By eliminating traditional storage trade-offs, ARCC liberated valuable staff resources to expand support for advanced research across campus and regional institutions.

### About VAST Data

VAST Data is the AI Operating System company – powering the next generation of intelligent systems with a unified software infrastructure stack that was purpose-built to unlock the full potential of AI. The VAST AI OS consolidates foundational data and compute services and agentic execution into one scalable platform, enabling organizations to deploy and facilitate communication between AI agents, reason over real-time data, and automate complex workflows at global scale. Built on VAST’s breakthrough DASE architecture – the world’s first true parallel distributed system architecture that eliminates tradeoffs between performance, scale, simplicity, and resilience – VAST has transformed its modern infrastructure into a global fabric for reasoning AI. Learn more at vastdata.com and follow VAST Data on LinkedIn, YouTube and X.



©2026 VAST Data, Inc. All rights reserved. All trademarks belong to their respective owners.

10/2026 I 16210