



AHEAD Helps University of Pennsylvania Build an AI and HPC Environment for Expanded Research Capacity

The Research Mandate

With more than 28,700 students, 12 schools, and a \$1.37 billion research budget, the University of Pennsylvania brings serious scale to serious research. That dedication to research was no more apparent than when Penn kicked off its Penn Advanced Research Computing Center (PARCC) initiative: A shared AI and high-performance computing environment, backed by high-speed storage and networking infrastructure.

The new PARCC facilities would offer stronger support for AI work, next-generation technology deployments, and smarter resource allocation. Plus, broader access to advanced GPU and CPU resources would power data-driven discovery for Penn's growing research community.

TOP TAKEAWAYS

- ✓ Built a 30-rack NVIDIA DGX SuperPOD deployment for advanced engineering capabilities.
- ✓ Brought compute, storage, networking, virtualization, and ongoing managed services together under one program.
- ✓ Doubled compute capacity for AI training and inference to support researchers across Penn's 12 schools.
- ✓ Moved Penn from RFP to a functional system in less than a year, effectively cutting the project timeline in half.



As demand for advanced computing grew, the **University of Pennsylvania** partnered with AHEAD to build PARCC, a shared AI and HPC environment with scale, speed, and staying power.

CLIENT



INDUSTRY



Higher Education

SCALE



28,700+

Students

12

Schools

\$1.37B

Research Budget

AHEAD FOCUS AREAS



AI-Enabled Infrastructure

Platform & Workload Modernization

Managed Services

The Questions (and Project) Got Bigger

Demand for advanced computing was climbing across campus. Penn needed to balance building PARCC alongside job growth, faculty priorities, and emerging AI workloads.

TALENT AND GROWTH

Faculty retention, along with attracting top talent and students, were major priorities. PARCC would need to support a wide variety of research projects, while providing enough headroom across Penn's schools and academic community.

INCREASING DESIGN COMPLEXITY

What started as a new computing center was rapidly expanding in scope. PARCC's success required a shared storage design, partner ecosystem coordination, and future expansion paths.

THE HARD TECHNICAL QUESTIONS

As the build continued, more questions kept popping up. How would Penn fund, design, engineer, and implement the shared storage infrastructure? How would it provision the necessary GPU and CPU capacity?

GOVERNANCE AND OPERATIONS

Even after PARCC was up and running, Penn needed a governance model to operationalize, maintain, and expand the environment over time and across other universities.

A good researcher will tell you: no plan succeeds without a program to guide it and trusted collaborators. PARCC needed a program to help it run and scale; Penn needed a partner that could pay equal attention to architecture, operations, and long-term support.



The Final Proposal

This project drew from multiple areas of AHEAD's expertise:

- **AI-Enabled Infrastructure:** Designing and integrating the GPU, storage, and networking infrastructure for expanded AI and HPC research.
- **Platform & Workload Modernization:** Turning that infrastructure into a shared, scalable platform and operating model the university could grow with.

AHEAD understood exactly how to help Penn launch, modernize, and automate its digital services and accelerate its AI and HPC initiatives. Together, they sat down to refine PARCC's requirements, improve its current design, and accelerate the deployment timeline.



The work unfolded across three phases:



Advise

The computing center, which started as a smaller eight-system plan, grew to a full 30-rack SuperPOD. AHEAD helped adjust the original architecture requirements and design, developing a phased strategy for centralized GPU and CPU capacity.

Why the Groundwork Mattered: Early design work clarified how the environment would operate and grow, while always keeping the larger campus vision in view.



Build

From there, the work became a full-platform effort. AHEAD coordinated with NVIDIA and Penn stakeholders, using an AHEAD Foundry deployment strategy to solve complex logistical challenges.

The final platform build included:

- NVIDIA DGX SuperPOD infrastructure
- 248 B200 GPUs and 47.0 TB of HBM3e VRAM
- Dell PowerEdge servers with AMD EPYC Genoa CPUs
- VAST storage, which combined multiple generations of flash and CPU hardware
- Ceph clusters for workload tiering and disaster recovery
- NVIDIA networking
- Cisco support for recovery infrastructure
- Red Hat OpenShift virtualization

So many moving parts that all had to line up across storage, networking, and virtualization layers. Fortunately, according to Ken Chaney, PARCC's Associate Director of AI and Service Technology, AHEAD worked "to make sure that we have good, solid relationships with all our hardware vendors, such as NVIDIA, Dell, and Vast."

How It Came to Life: By coordinating the ecosystem, AHEAD solved deployment issues in advance, so the environment could be assembled and stood up quickly.



Run

Post-deployment, AHEAD has remained closely involved with Penn. AHEAD Managed Services support the GPU and HPC cluster, along with storage and networking. Ongoing operational support and an on-site engineering presence help keep the environment stable as PARCC continues to evolve.

How It Kept Delivering: PARCC can keep growing as a living program with new workloads, new users, and the next wave of research demand.


A Bigger Lab Bench

The PARCC environment saw gains almost immediately in capacity and delivery speed:

Getting to the Good Stuff Sooner. Penn moved from RFP to a functional system in under 12 months. Chaney shared: “AHEAD helped us reduce the timeline for this project from what would have been an 18-month project after engagement down to a 9-month project.”

A Wider Field of Inquiry. The SuperPod will double the compute capacity available on campus for AI training and inference, giving researchers across Penn far more room to pursue their latest ambitions.

Ready for the Next Hypothesis. Penn is positioned to expand access to its advanced compute resources, better support interdisciplinary work, and give faculty and researchers infrastructure that keeps pace with the cutting-edge.



What's Next

PARCC's roadmap already reaches beyond the initial deployment. The program plans to offer general HPC services alongside a condo model for faculty who want to own hardware while taking advantage of co-located resources. As a strategic partner, AHEAD is helping Penn pursue new goals for its research community like equitable access, sustainable operations, workforce support, balanced infrastructure design, and governance for the long term.

“One of the exciting things about this project is that the one SuperPod that AHEAD is helping us put together in collaboration with NVIDIA will double the compute capacity available on campus for AI training and AI inference. The amount of research that we'll be able to get done across Penn is going to go up dramatically!

— Ken Chaney, Associate Director of AI and Service Technology, PARCC



Penn Advanced Research
Computing Center
UNIVERSITY of PENNSYLVANIA

AHEAD

Combining cloud-native capabilities in software and data engineering with an unparalleled track record of modernizing infrastructure, we're uniquely positioned to help accelerate the promise of digital transformation.

Visit us at ahead.com.

National Hubs

CHICAGO

444 W. Lake Street
Suite 3000
Chicago, IL 60606

NEW YORK

500 5th Avenue
Floor 17
New York, NY 10010

ATLANTA

1117 Perimeter Center
W406
Atlanta, GA 30338

SAN FRANCISCO

2000 Crow Canyon Place
Suite 250
San Ramon, CA 94583