

Step by Step: Higher Education's Path to AI Readiness



Introduction

Colleges and universities are exploring how AI can enhance research, student services and campus operations. A 2025 American Council on Education survey of 517 higher education leaders found that nine in 10 respondents had adopted the technology or expected to do so within two years.¹

But turning promising pilots into reliable, campuswide capabilities depends on something institutions have spent decades accumulating: data. Much of that important information — related to students, research and operations — remains distributed across administrative systems, departmental infrastructure and cloud applications, making it difficult to access and share.

“Institutions have so many different sources of data,” says Brian Cohen, vice president of the Center for Digital Education. “Being able to map, catalog, access and trust that data, while minimizing duplication, is an enormous effort.”

AI projects often need to pull information from more than one campus system. Some also require greater compute capacity than existing systems provide. These challenges can prevent a successful pilot from expanding across the institution. But overcoming them doesn’t necessarily require colleges and universities to replace everything they already have.

Campus leaders must weigh their AI ambitions against aging infrastructure, tight budgets and technology investments that may still have useful life. The path to AI readiness starts with understanding what a priority use case requires, identifying the data or infrastructure that limits it, and modernizing where it matters most.

By taking that incremental approach, institutions can build on existing investments while creating the foundation needed to deliver reliable value from AI.



AI-Related Campus Challenges

Data silos. Growing AI adoption puts new pressure on campus data systems because each use case draws on a different mix of information. Student support may require records from several administrative applications, whereas a research model could depend on large collections of video recordings or scientific measurements. Institutions often store these materials in different systems and formats.

“The university environment tends to be siloed by design,” says Scott Lewis, director of solutions consultants at Hitachi Vantara, a leading data infrastructure company. “Different colleges within the same university, research institutions and IT groups often operate separately. As a result, data is spread across campus. To me, that is the No. 1 challenge to innovating with AI in higher education.”

Unforeseen requirements. An AI pilot project rarely reflects what campuswide use will require because the team relies on a selected dataset and temporary



computing resources. Once the institution expands the project, continuous operation places greater demands on data availability, the protection of sensitive records and recovery planning.

Finances. Costs can change significantly when a pilot becomes an ongoing service. In many institutions, AI implementation has outpaced long-term budgeting. The 2025 EDUCAUSE AI Landscape Study, based on responses from 788 higher education professionals, found that 46% were implementing AI-focused initiatives, but only 19% were budgeting for anticipated costs of long-term AI use.²

Locally managed research infrastructure can compound the cost problem. Some research groups already operate their own servers, storage and specialized computing resources. Local investments can support immediate research needs, but they often lead several departments to purchase similar capabilities independently. “It’s too expensive to do AI five times for five departments,” Lewis says.

Your institution may need to modernize tools before AI can run reliably on campus.

Investing in different solutions across campus raises costs and makes it harder to extend a successful project beyond the group that created it.

Modernization. AI tools depend on campus systems to store data and deliver it promptly. If those systems are older or outmoded, an institution may need to update them before AI can run reliably across campus. As a result, higher education leaders are now planning for AI and modernization together. “AI and modernization are the same challenge on two different timelines,” Lewis says. “AI is making the case for modernization relevant to presidents and provosts, but modernization has to come first.”



Conducting a Readiness Assessment

Campus leaders can use a readiness assessment to examine the systems and data an AI initiative would rely on. In addition to inventorying an institution's full AI environment, this review will show how much additional work can be handled by existing equipment and when vendor support ends. The review should cover four factors:

1. **System lifecycle and capacity** – A review of existing hardware and software identifies the systems a proposed AI use case would need and shows how much additional work they can handle. A system with enough capacity and continued vendor support may remain in service. If a system is nearing the end of vendor support, leaders can plan an update or replacement.
2. **Data location** – An assessment shows where the datasets for an AI project are stored, including those managed by different departments and research groups. If a large research dataset is housed in a campus data center, transferring it to a cloud service for analysis takes time. The institution could instead opt to run the AI software in the campus data center, provided it has enough compute capacity.

3. **Risk and compliance** – A review identifies systems that hold protected student records or clinical data and verifies who can access them. Student records may be subject to FERPA, while some clinical information may fall under HIPAA. Existing permissions may need to change before sensitive data is used in an AI initiative.
4. **Service criticality** – Leaders should examine what would happen if a system stopped working. A learning management system outage could interrupt teaching, while a registration failure could prevent students from enrolling. Institutions should map these possibilities and build a strategy for preventing system failures and mitigating them if and when they occur.

An Incremental Approach

An assessment usually uncovers several areas that require attention before an AI solution can be deployed at scale. Addressing all of them at once may not be feasible or realistic. Instead, modernizing one layer at a time is a good strategy, and it should be driven by an institution's own mission goals.

For example, a university looking to expand AI-driven research may find that existing storage cannot deliver data fast enough to the GPUs it has purchased.

Improving storage could let researchers make fuller use of that investment. Meanwhile, if an institution wants to streamline and improve registration with AI, the systems supporting enrollment should take priority. An AI assessment helps leaders connect each proposed investment to a service they need to protect or an academic goal they want to pursue.

Assessments also show where replacement can wait. A system that still receives vendor support and can be secured is often worth retaining. Leaders can then check whether it has enough capacity for the planned workload. “Replacement is costly, and it’s not just the cost of new equipment,” Lewis says. “Replacing systems can also include recertification costs, new compliance challenges and training for staff.”

A deliberate, phased approach can also help align proposed infrastructure spending with campus budget cycles.

Each phase of a plan for adopting AI should focus on a result that presidents and financial leaders recognize. “Define valuable milestones for each phase and put them in metrics a president or CFO understands,” Lewis says. “Rather than say, ‘I need an extra petabyte of storage,’ say, ‘I need to reduce registration problems during the first week of classes.’”

This approach also helps leaders determine the best way to fund each phase. Buying equipment upfront is one approach, while a consumption model can align spending with the capacity used at each stage. If demand is likely to grow in stages, leaders can compare the cost of buying capacity for the full plan now with the cost of paying as use grows.

Each phase of your AI adoption plan should focus on a result that presidents can recognize.

Preparing Existing Data for AI

Once leaders decide which AI use case to prioritize, campus IT can use the assessment findings to arrange access to the datasets that project needs. The practical work has four parts:

- 1. Map relevant data** – Technology staff can turn the assessment findings into a record of where each dataset resides, who owns it and how an AI team can request access. That record may point to a central administrative application or a research group’s own servers. It gives project teams a way to find information held outside their department.
- 2. Classify data and govern access** – When an AI project crosses departmental boundaries, teams need to find the appropriate data owner and determine whether they have permission to use the information. Tags that identify a dataset’s contents and sensitivity can help staff efficiently find information and apply the right access rules. “Metadata becomes critical,” Lewis says. “If you can’t classify and tag your data, it’s not going to be useful for AI training, inferencing or fine-tuning.” Data owners can then decide which project teams may use the information under the institution’s rules.
- 3. Protect data and plan for recovery** – Technology and security staff should check that access restrictions still apply when an AI team uses student or clinical information. They should also confirm that the institution can restore the data the team needs if a system fails or records are lost.
- 4. Modernize the limiting layer** – If existing storage or computing equipment cannot handle a planned AI workload, leaders can upgrade the limiting component. Other supported equipment can remain in service while the team tests whether the change supports its intended work.



Choosing Where Workloads Run

The right place to run an AI workload can change between initial training and everyday use. A hybrid cloud approach lets a university run some workloads in its own data center and rent cloud capacity for others. If different workloads call for different providers, the university may also use more than one cloud platform, known as multicloud.

“Cloud doesn’t compete with on-premises infrastructure,” Lewis says. “Cloud is a tool, as is the on-premises data center. The question is where it makes sense to run a workload at a given time.”

For instance, AI models tend to require powerful GPUs during training and fine-tuning, but far less compute power once the model is running. Renting that capacity can make more financial sense than buying equipment the institution will no longer need when training ends.

Data location and sensitivity also affect where the university should run AI software. If a workload

Renting capacity in the cloud can make more sense than buying equipment.

depends on a large research dataset held on campus, IT should account for any data transfer charges when comparing cloud and on-campus options.

For a project using student or clinical records, data owners should review who would control access in the cloud and whether the service meets the university’s privacy and security rules. The university can keep those records on campus or consider a cloud service that meets its privacy and protection requirements.

Campus IT should reassess a placement decision when a project moves from development to regular use, since the amount of data transferred and the ongoing cost may change.

Establishing Access, Governance and Ownership

As leaders plan phased modernization, they also need agreement on who can approve access to the campus data an AI project requires. “Most of the barriers to adoption are not technical,” Lewis says. “Shared data requires shared governance. Decisions may require consensus across multiple departments or colleges.”

A governance process identifies who has authority to approve use of each dataset. It also gives departments a way to resolve requests that cross their usual boundaries.

Many institutions assign an executive sponsor to oversee specific AI implementations. This person can convene stakeholders and help align academic priorities with the project’s operational demands. For example, departments accustomed to choosing their own technology may worry about losing authority when their institution adopts shared infrastructure. An executive sponsor can bring departmental leaders into the discussion and make sure the proposed change serves an established campus objective. Depending on the university’s size and culture, that sponsor might be someone who works in academic affairs, operations or IT. Some institutions appoint more than one sponsor when a proposal affects several parts of the campus.

Faculty and researchers need a clear route for proposing AI pilot projects, too. “If you want to accelerate adoption, make governance the enabler,” Cohen says. In practice, a faculty team describes the purpose of a proposed pilot and identifies the records it needs. The governance process then shows the team who has authority to approve access and which rules apply. Connecting the pilot to a specific academic or service goal also gives campus leaders a basis for deciding whether to support it.

Campus leaders can revisit those rules as AI use cases develop. Policies written by different

Governance can serve as a way to promote AI adoption among faculty and staff.

offices sometimes conflict, or an older rule may no longer fit the work a team proposes. In a periodic review, leaders can resolve conflicts and agree on who will approve an AI project that uses data from more than one department. If a governance group receives similar proposals from several departments, its members can also consider whether one shared investment would serve them better than separate purchases.

Conclusion

An institution can begin with a priority AI use case and identify any necessary modernization for the project. In this regard, some institutions will get more from the systems they already have, some will need to optimize and modernize at the same time, and some will need new infrastructure before anything else.

“I don’t think there’s one option that will meet everyone’s needs,” Cohen says. “Institutions can choose different approaches for different parts of their strategy.”

After each phase, campus leaders can check whether their investment met its stated goal. The results can guide their next decision, and leaders can keep existing systems in service if they still meet the university’s needs.

“Start by making the data you already have visible, governed, protected and recoverable,” Lewis says. “Then modernize one layer at a time, showing incremental value along the way. You’re closer to being ready for AI than you may think.”

1. <https://www.acenet.edu/Documents/Pulse-Point-Survey-Fall2025.pdf>
2. <https://er.educause.edu/articles/2025/3/ai-procurement-in-higher-education-benefits-and-risks-of-emerging-tools>

This piece was written and produced by the Center for Digital Education Content Studio, with information and input from Hitachi Vantara.



Produced by the Center for Digital Education

The Center for Digital Education is a national research and advisory institute specializing in K-12 and higher education technology trends, policy and funding. The Center provides education and industry leaders with decision support and actionable insight to help effectively incorporate new technologies in the 21st century.

www.centerdigitaled.com

Hitachi Vantara

Sponsored by Hitachi Vantara

Hitachi Vantara is transforming the way data fuels innovation. A wholly owned subsidiary of Hitachi, Ltd., we're the data foundation the world's leading innovators rely on. Through data storage, resilient infrastructure systems, intelligent data management and AI-ready hybrid cloud storage — we're the data backbone for the mission-critical.

We bring agility and stability. Flexibility and security. Working closely with customers, we understand their aspirations and architect their data infrastructure that is built for sustainable business growth. As part of the Hitachi Group, we draw upon a global network of digital (IT) and operational (OT) expertise to empower industries in advancing the next innovation.

Government, State, & Local, Education | Hitachi Vantara