EDT & PARTNERS

Leading AI with Intention:
A Guide for University Leaders

April 2026

Index

Executive Summary

Artificial intelligence is no longer something universities are deciding whether to adopt. It is already here: in the learning management system, the advising platform, the research pipeline, the plagiarism detector. It is in the productivity software that everyone who touches a computer uses for daily tasks, from mundane to complex. The question facing university leaders is not whether AI will transform their institution, but whether they will shape how that transformation unfolds.

The scale of what is already happening is significant. Over 90% of students and staff now use AI tools daily, and institutional adoption has increased sharply in a single year. Yet much of this growth is occurring through individual use, local experimentation, and point solutions rather than deliberate institutional design. AI is becoming ambient infrastructure as it accumulates across systems and practices that now structure institutional life.

This is not solely a technology or compliance question. It is, above all, a leadership question about what kind of institution you will become in light of such rapid technological change. Rather than operating with the greatest caution or fastest speed, the universities successfully navigating this moment are those that have been most intentional.

This paper makes a simple argument: that AI should be considered institutional infrastructure, and managing it well requires visibility, shared capability, and people genuinely at the center. Doing so will create conditions for the whole institution to benefit, not just the departments that have already embraced AI, but every student, educator, and staff member the institution serves.

Three things that distinguish institutions leading on AI:

- **They know where AI is operating,** and use that visibility to build on what's working, not just manage what's risky.
  - → Proactively monitor shadow AI and provide both application and network-level guardrails and institutional governance.
- **They build shared capability,** so AI use compounds across the institution rather than staying siloed in departments.
  - → Think multi-stakeholder project-based AI use cases, libraries of prompts and quick prototype validation.
- **They keep people at the center,** ensuring that human judgment, not algorithmic output, drives the decisions that matter most.
  - → Different forms of human-in-the-loop design are embedded in both the technology and the decision-making process. Innovation is not measured on time saved or output but quality, outcomes and differentiation.

SECTION 1: From Experimentation to Infrastructure

AI is Already Here

Three years after generative AI entered mainstream awareness, the question for universities is no longer whether AI will be adopted. It already has been adopted widely, rapidly, and largely organically. Coursera's 2026 global study found over 95% of students and educators use AI in educational contexts.[1] On the staff side, 90% of higher education professionals now use AI, and institutional adoption jumped 17 percentage points in a single year.[2]

AI represents an infrastructure shift as significant, in its way, as the arrival of the internet or cloud computing. AI is now woven into how students learn, how faculty teach, and how institutions operate. It is embedded in the learning management system (LMS), the email system, the advising platform, and the research pipeline. Like electricity or Wi-Fi, it is ambient infrastructure, present whether or not institutions have explicitly chosen to deploy it.

The opportunity at this moment is considerable. Institutions that understand where AI is operating, and that create the conditions for it to be used well and responsibly, are positioned to deliver meaningfully better outcomes for students, faculty, and staff. The institutions best positioned to seize that opportunity are those that get ahead of it, not by slowing adoption but by making it coherent.

The Visibility Challenges

The primary challenge is not adoption. It is that so much adoption is currently invisible to the institution. Over 56% of higher education workers use shadow AI, or tools acquired outside institutional knowledge or review.[3] Meanwhile, vendors are embedding AI capabilities into products universities already own, often without disclosure. And the rapid proliferation of AI across organizational functions — called agent sprawl — means that even well-intentioned institutions can find themselves with duplicated capabilities, inconsistent data handling, and no single point of accountability.

The consequences are not theoretical. In higher education, where institutions hold vast quantities of student data protected by Family Educational Rights and Privacy Act (FERPA) and equivalent regulations, fragmented AI adoption creates specific risks: regulatory liability, vendor relationships that have outpaced contract terms, and inconsistent experiences for students depending on which department or classroom they happen to be in.

None of this is cause for alarm, but it is cause for intention. The institutions navigating this well are not the ones that have slowed AI adoption. They are the ones that have made it legible, coordinated, and purposeful.

1 AI and Higher Education. (February 2026). Coursera.
https://blog.coursera.org/ai-in-higher-education-report-2026/
2 Artificial Intelligence in Higher Education: From Widespread Adoption to Strategic Integration. (March 2026). Ellucian.
https://www.ellucian.com/newsroom/ellucians-3rd-annual-higher-education-ai-survey-signals-shift-individual-ai-use
3 Data Shows AI 'Disconnect' in Higher Ed Workforce. (January 2026). *Inside Higher Ed*.
https://www.insidehighered.com/news/workplace/staff-issues/2026/01/13/data-shows-ai-disconnect-higher-ed-workforce

SECTION 2: What an AI Management System (AIMS) Makes Possible

AI as a Managed Institutional Capability

The response to shadow AI and agent sprawl is not more policy. Policies define acceptable behavior but they do not manage infrastructure. Universities already understand this distinction in other domains. Financial systems are governed through accounting frameworks, audit controls, and continuous monitoring. Research compliance is managed through institutional review boards, protocols, and oversight processes. Cybersecurity is not governed through a single policy document but through ongoing monitoring, incident response, and risk management systems. AI is no longer a collection of tools that can be governed individually. It is a distributed institutional capability embedded across teaching, research, and operations. Managing it requires the same shift.

What an Effective AIMS Provides

An effective AI management system (AIMS) does more than manage risk. It creates the conditions for AI use to compound across the institution. What one department learns, another can build on, so that every member of the community can engage with AI confidently, responsibly and ambitiously. At its core, an effective AIMS provides three essential capabilities:

1. **Visibility - so institutions can build on what's working**

   Institutions need to know where AI is operating and how it is being used, not as a surveillance exercise, but as the foundation for institutional learning. Without visibility, every team reinvents the wheel. With it, successful approaches can be identified, shared, and scaled. This includes AI tools produced centrally, AI features embedded in point solutions, department-level adoptions, and personal tools used in institutional workflows.

2. **Shared Capability - so AI use compounds rather than fragments**

   The most significant opportunity for institutional AI is not any single tool. It is the accumulated knowledge of how AI is being used effectively across the institution. An AIMS creates the infrastructure for that knowledge to circulate: shared libraries of testing prompts and workflows, structured and tested ways to evaluate and compare model outputs, and common frameworks for prototyping new use cases quickly and safely. Rather than each college or department building from scratch, the institution develops a growing base of shared capability that every team can draw on and contribute to.

3. **Continuous Oversight - so people, not AI, drive what matters**

   AI systems cannot be approved once and forgotten. The systems evolve, outputs change, and the impact on students, staff, and institutional decisions needs to be tracked over time. Effective oversight means monitoring AI outputs and behavior on an ongoing basis, conducting periodic assessments of risk and impact, reviewing vendor relationships as models are updated, and measuring AI use against the outcomes that matter.

SECTION 3: AI Stewardship

Leading AI with Intention

AI governance gets misclassified as a technology problem, and that framing leads institutions in the wrong direction. What's actually at stake cuts across academic policy, research integrity, data governance, vendor management, and institutional reputation in ways that connect people and practices beyond the IT department. Leading AI with intention means recognizing this from the outset: that the decisions shaping how AI is adopted, monitored, and evolved are fundamentally institutional choices, not technical ones.

In practice, those choices surface as questions like: What kind of learning environment do we want to create? How do we protect student data while enabling innovation? Who is best positioned to make decisions when AI produces unexpected or unintended outcomes? These, and many more, are questions for presidents, provosts, and academic leaders, not just technology teams.

The Emerging CAIO Role in Higher Education

Many universities are responding by establishing a Chief AI Officer (CAIO) as a coordination point, not centralized control, that spans academic and administrative domains. Institutions including George Mason University, UCLA, and the University of Utah have created CAIO positions to align AI strategy with institutional priorities, manage risk, and bridge faculty, staff, administration, and technology teams.

Crucially, the CAIO is not a rebranded CIO/CTO. Where the CIO's mandate centers on infrastructure and systems, the CAIO's scope is fundamentally about people, practice, and institutional culture. In many institutions, the position reports directly to the Provost or President, a reporting structure that reflects its cross-cutting nature. In some cases, it is explicitly designed to be transitional: establish governance structures that eventually become distributed across institutional leadership rather than concentrated in a single office.

Centralized vs. Federated Governance

Higher education governance is inherently federated, and effective AI governance works with that grain rather than against it. Departments, colleges, and administrative units exercise significant autonomy, particularly in teaching and research. Any governance model that ignores this will be bypassed quietly or rejected loudly.

The model that works is one where the center establishes shared standards, common risk frameworks, and vendor policies, while local units retain authority over how those standards are applied within their disciplines.

Virginia Tech's cross-institutional AI Working Committee with domain-focused subgroups spanning research, teaching, and administration, reflects this logic,[4] as does UCLA's cross-functional advisory structure, which was explicitly designed to bring decentralized campus entities under cohesive guidance without overriding their autonomy.[5]

4 Responsible and Ethical Artificial Intelligence: Framework for Virginia Tech. (July 2025). Virginia Tech AI Working Group Final Report. https://ai.vt.edu/content/dam/ai_vt_edu/Responsible-and-Ethical-AI-Framework-for-Virginia-Tech-(v1.0-Sep-2025).pdf

5 Artificial Intelligence Governance. UCLA. https://dts.ucla.edu/initiatives/ai/governance

Clarifying Ownership

One practical gap many institutions encounter is not policy. It is clarity about who owns what. As AI becomes embedded in advising systems, research pipelines, and administrative workflows, questions naturally arise: who is responsible for an AI tool that produces unexpected results? Who monitors vendor behavior as models are updated? Who decides when systems should be retired?

These are not edge cases, but the normal operational realities of deploying AI at scale. Institutions that navigate them well typically have a few things in common: clear executive sponsorship, defined roles across units, particularly IT, academic leadership, and legal counsel — and a shared understanding of how decisions get escalated when something goes wrong. The goal is not bureaucratic control, but instead making sure that human judgment remains genuinely at the center.

SECTION 4: Designing the AI-Native University

Building the AI-Native University

Becoming an AI-native university involves many dimensions of change for an established organization. However, when you build an institution from the beginning, you have a genuine opportunity to apply a holistic AI view of what higher education should be.

EDT&Partners is currently working with a new institution building from the ground up, with AI as its operating infrastructure from day one. It has no legacy systems to integrate, no incumbent vendors to negotiate around, and no established workflows to protect. What it has is a clear mandate: scale from one thousand to ten thousand students within a year without compromising the quality of the student experience.

The constraints of radical growth, limited time, and finite resources have shaped every AI decision the institution made. And the decisions it has made offer a useful frame for any university thinking about where to start.

1. Start where AI creates the most immediate institutional value.

The institution did not begin with a complete transformation vision. It began where many institutions have already deployed AI: the chatbot. Specifically, a recruitment chatbot as the first point of contact between a prospective student and the university. This is a familiar starting point, but one the institution chose deliberately rather than by default. Recruitment is where operational efficiency translates most directly into institutional revenue, and visibility into that process was the foundation everything else would build on. Critically, the AI does not operate alone. Human advisors remain part of the sequence, intervening at the moments where a personal connection matters most.

The institution did not stop at recruitment. The next high-friction moment in the student journey was pre-qualification: the collection and review of documents like test scores and transcripts that must happen before a human advisor meaningfully enters the conversation. This process is time-consuming, largely administrative, and poorly suited to human attention at scale. AI agents now manage it by gathering, organizing, and screening materials so that advisors engage only with students who are genuinely ready for that conversation.

The result for both recruitment and admissions is a blended model where AI handles volume and repetitive tasks, and humans focus where their judgment and relationship matter most.

2. Scale the things that humans cannot.

The second priority was assessment. At one thousand students, human graders are manageable, but at ten thousand, they become a bottleneck that threatens the institution's ability to grow. Because accreditors and academic integrity both require human final judgment, the institution's response was not to hand grading and assessment to AI, but to use AI to dramatically accelerate the process. AI generates rubrics, structures feedback, and surfaces criteria that help teachers make faster, more consistent decisions. The teacher remains the decision-maker and AI removes the administrative weight that was slowing the decision down.

This is institutional ownership in practice: AI accelerates the process, but the institution never cedes the decision.

3. Build learning content, experience, and support as a single system.

Course creation follows the same logic. It begins with defining the learning objectives and outcomes for each course, the pedagogical starting point that shapes the entire design. AI then generates the technical scripts, and learning objects are assembled and published to the LMS. A brief expert recording, based on two minutes of footage from a subject matter expert, becomes a full avatar-delivered lecture through AI video generation tools. An AI tutor is embedded in every course, with access only to that course's knowledge base so it will not speculate beyond its domain. The AI tutor is designed to support and monitor learning, not provide answers.

4. Connect what was always siloed.

What holds all of this together is an orchestration layer, not a database overhaul. Existing systems, including the LMS, the CRM, the student information system, all remain in place, while AI pulls relevant information from each, surfaces connections that were previously invisible, and acts on them in real time.

This orchestration layer means that a student who is struggling in a course triggers a sequence — AI tutor escalation, human advisor outreach, a recommended resource — that draws on data from multiple systems that previously never spoke to each other. Patterns of engagement across student services are correlated with AI tutoring use and course performance, and become visible at the institutional level. Interventions that were previously made on instinct or with incomplete information can now be made on evidence.

This is what measurement infrastructure actually looks like when it is built into operations rather than added afterward. The data was always there but what was missing was the orchestration layer to make it speak.

What makes this work is not technology but the careful decisions underneath it to determine what data matters, what connections are meaningful, and what actions are appropriate. When those decisions are made deliberately and governed with intention, the result is an institution that learns as it grows.

The cost argument

There is a financial dimension to this that university leaders should consider. Institutions that own their AI infrastructure rather than paying per-seat SaaS licensing face a fundamentally different cost curve as they scale: one where costs grow far more slowly than enrollment.

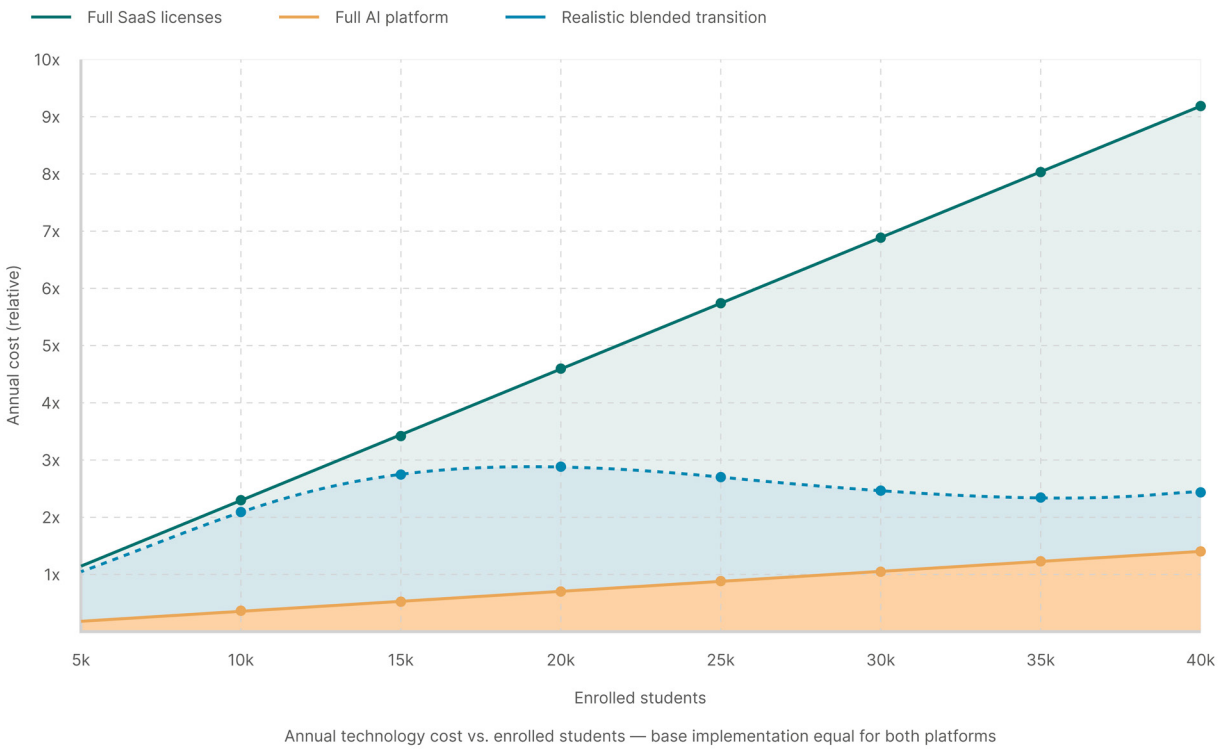

Annual technology cost vs. enrolled students — base implementation equal for both platforms

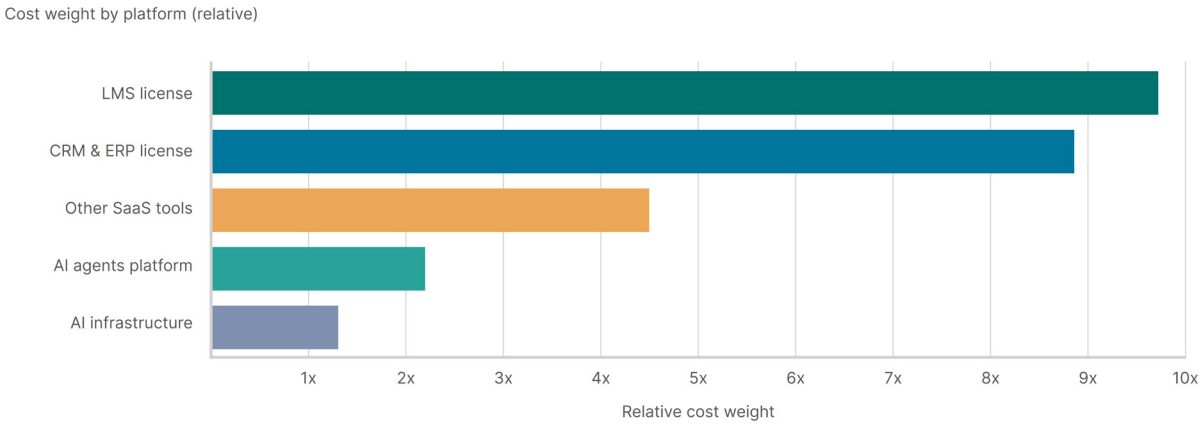

The most realistic path for most institutions is not a full replacement of existing systems but a blended transition: maintaining core SaaS infrastructure while building AI capability on top of it. EDT&Partners' experience suggests that the cost advantages of owned AI infrastructure over traditional SaaS licensing, at scale, can be substantial, and that the gap widens as enrollment grows. The implication for institutions planning growth or looking for cost savings is worth taking seriously.

From one institution's experience to yours

This institution's experience is not necessarily a template. It is a new institution, building without legacy constraints, with growth goals specific to the market it seeks to serve. The decisions it made about where to start, what to automate, where to insist on human judgment were shaped by its particular circumstances.

But the logic underneath those decisions is transferable. Visibility before strategy. Institutional ownership built into the design. Measurement infrastructure as an operational reality, not a reporting afterthought. Every institution that is serious about AI governance is navigating toward the same questions, regardless of where it is starting from.

The questions that follow are an invitation to locate your institution honestly — not where you aspire to be, but where you actually are.

SECTION 5: Leading with Intention: A Reflection for University Leaders

AI technologies are rapidly changing. The institutions that will navigate it well are the ones that know what they are navigating for.

Visibility - Do you know where AI is operating?

1. If a student asked you exactly how AI is being used in their learning experience — in the platforms they use, the feedback they receive and the systems that support them — could you give them an honest and complete answer?

2. When a vendor updates a product your institution uses and adds new AI capabilities, how do you find out and how quickly?

3. Where in your institution is AI being used most ambitiously? Do you know, and are those experiments visible enough so that others can learn from them?

Shared Capability - Are you building institutional knowledge or reinventing the wheel?

1. When a faculty member or administrator finds an effective way to use AI in their work, what happens to that knowledge? Does it stay with them, or does it circulate?

2. Does your institution have a shared way for staff, faculty, and students to access tested AI tools and approaches, or is everyone starting from scratch?

3. How are you measuring whether AI use is actually improving outcomes for students, and who is responsible for that question?

Continuous Oversight - Are people, not AI, driving the decisions that matter?

1. For the AI systems most consequential to your students, who is responsible for their outputs, and how would you know if something was going wrong?

2. When did you last review a significant AI vendor relationship, not just the contract, but whether the tool is still performing as intended and whether the vendor has changed what it does?

3. Are there decisions in your institution that have quietly shifted from human judgment to algorithmic output? If so, was that a deliberate choice and are the trade-offs and benefits known?

AI Stewardship - Are you leading with intention?

1. Does your institution have a shared sense of what AI is for, what it should enable, and what it should never replace?

2. Are the people most affected by AI — students, faculty, staff — part of the conversation about how AI is used, or are those conversations happening without them?

3. A year from now, what would it mean for your institution to have led on AI well?

About EDT&Partners

EDT&Partners is a global education and technology consulting firm transforming the education and knowledge ecosystem through purposeful technology, strategic expertise, and hands-on execution. With a presence in over 16 countries, the company partners with universities, schools, governments, publishers, and organizations worldwide to help them transform responsibly and unlock human potential in the AI era.

For more insights from EDT&Partners follow us on LinkedIn or visit The EDiT Journal.

Learn more about our work in AI governance and higher education.

About the Authors

This report was led by Nina Huntemann, with contributions from the EDT&Partners team.

Nina Huntemann is a Senior Consultant at EDT&Partners. She previously served as Chief Academic Officer at Chegg and VP of Learning at edX. She works with higher education institutions and edtech companies on AI strategy, evidence-based teaching and learning, and student experience.

Pablo Langa is the Founder and Managing Partner of EDT&Partners. With over two decades of international experience, he has led hundreds of projects at the intersection of education, business, and technology—helping governments, institutions, and organizations of all sizes to define strategy, transform operations, and scale impact. Pablo brings senior leadership experience from some of the world's most influential education organizations, including Blackboard and Headspring.

Laurie Forcier is the VP for Strategy at EDT&Partners. She specialises in helping innovative education organizations articulate their purpose, translate complex ideas into accessible narratives, and build relationships that drive meaningful collaboration. She is co-author of Intelligence Unleashed: An Argument for AI in Education and has held leadership roles at Harvard University, Pearson, and ISTE.

Yuri Rangel is the VP of Growth at EDT&Partners, where he leads global expansion and go-to-market strategy for education organizations and EdTech companies. He specialises in internationalisation, AI adoption, and building scalable growth models that connect strategy to measurable impact.

Want to explore what this could mean for your institution? Get in touch.

EDT & PARTNERS