



How AI Impacts Datacenter Spending in Higher Education

ROI and ROV considerations in the age of AI growth



Hernán Londoño, Chief Technology and Innovation Strategist, Lenovo

In many of my recent conversations with technology leaders in higher ed, a recurring question is not whether they should invest in AI tools and hardware - it's about how those AI investments should be integrated into their organizations' overall data center growth plans.

To put a finer point on it, it's a no-brainer that universities will need to invest in AI. At a time when budgets are tight, staffing is constrained, and the number of competing priorities seems to grow by the day, the challenging part is creating a coordinated plan to keep modernizing your datacenter over time - but with AI investments now as top of mind.

Economic challenges unique to higher education

If I were to oversimplify the current higher education technology environment, I would describe it this way: institutions are being asked to modernize while carrying too much weight.

That weight comes in many forms. It comes from aging infrastructure. It comes from technical debt. It comes from understaffed teams that are already overextended maintaining the current state while also being expected to build the future state. It comes from growing cybersecurity expectations, increasing governance requirements, and rising pressure to support a wave of AI-related curiosity and demand from across the institution. It also comes from a basic and unavoidable reality: most institutions do not have the luxury of unlimited capital, unlimited operating dollars, or unlimited time[1].

New pressures around technology investments

Today, institutions simply cannot afford to be impulsive about technology investments. The days when tech spending could be justified primarily through aspiration, broad modernization language, or the promise of vague future benefit are gone. Finance leaders want clarity. Presidents want alignment. Boards want evidence. Provosts want relevance.

Decision making around datacenter and AI investments needs to move beyond excitement, anxiety, and the temptation to chase every emerging capability in the market. Institutions need to ground these decisions in two important but distinct concepts: return on investment, or ROI, and return on value, or ROV.

Considering both ROI and ROV

In my opinion, the most mature institutions are not choosing between ROI and ROV. They are using both. ROI provides discipline. ROV provides context. ROI helps answer whether the investment is economically responsible. ROV helps answer whether it is institutionally meaningful.

ROI is important. In a sector under financial pressure, institutions should absolutely ask what a technology investment returns in practical terms. However, in higher education, they are not the only questions.

ROV is about investments that may not produce a dramatic or immediate financial return and yet are still entirely justifiable, even necessary to the institution's larger goals.

Some examples of ROV considerations:

- A resilience investment may protect research continuity.
- A data governance investment may reduce institutional risk.
- A modernization effort may create the platform upon which future academic, research, and operational capabilities depend.
- An AI-related investment may strengthen the institution's ability to support teaching, research, student services, or administrative efficiency over time, even if the financial return in year one is modest or difficult to isolate cleanly.



Questions to ask about ROI

- Does it lower operational cost?
- Does it consolidate systems?
- Does it reduce downtime?
- Does it improve utilization?
- Does it free up scarce staff time?
- Does it reduce reliance on aging systems that are expensive to maintain and increasingly difficult to support?

Questions to ask about ROV

- Does the investment improve institutional capability?
- Does it reduce meaningful risk? Does it reduce downtime?
- Does it strengthen mission execution?
- Does it create optionality for the future?
- Does it improve institutional agility at a time when agility itself has become a strategic necessity?

Moving from investment urgency to solid planning

Datacenter and AI investments are now often discussed with a level of urgency that can create the wrong type of momentum. There is a tendency in some settings to believe that if AI is strategically important, then any spending related to AI must therefore also be strategic. I do not think that follows automatically. Wanting to be relevant in AI is not the same as being ready to invest wisely in it.

The same can be said of datacenter modernization. Refresh cycles are often treated as isolated technical events, but in many institutions, they are far more consequential than that. They determine flexibility. They determine resilience. They determine how efficiently a lean staff can operate. They determine whether future analytics, research, and AI ambitions will have solid footing or will instead sit atop an increasingly fragile foundation.

The real question is how institutions should decide.

Datacenter modernization requires strategy

I think one of the more underappreciated realities in higher education is that datacenter modernization is often discussed as though it were merely a technical upkeep exercise, when in fact it may be one of the most important strategic technology decisions an institution makes.

A modernized datacenter environment can:

- Improve uptime
- Simplify support
- Reduce operational friction
- Strengthen resilience
- Improve power and space efficiency
- Provide a much better platform for hybrid operations and future AI use cases.
- It can also help institutions whose teams are already thin by creating a more standardized and manageable environment



Consider a typical example: a mid-sized university has aging compute and storage platforms, backup challenges, inconsistent performance for data-intensive workloads, and increasing pressure from both researchers and administrators for more modern analytics and AI capability. The instinct in some settings might be to leap directly into a substantial AI infrastructure purchase.

I would recommend the following modernization process:

- First, stabilize and modernize the core. Improve compute, storage, networking, recoverability, and operational consistency.
- Reduce fragility. Improve provisioning and manageability.
- Create a platform that is more efficient to operate and more credible as a base for future growth.
- Lastly, expand into more targeted AI infrastructure where there is clear demand, governance, and a validated use case.

The point here is not caution for caution's sake. The point is that sequence affects value.

Decisions that should inform AI spending

AI is often discussed in language that suggests it is a discrete thing one buys. It is not. It is a stack of decisions. There are use-case decisions.

There are data decisions. There are governance decisions. There are infrastructure decisions. There are deployment-model decisions. There are security decisions. There are staffing and operational decisions. That is precisely why impulsive spending in this space can become problematic.

Institutions should begin with use cases that matter. Not every use case deserves investment, and not every use case requires the same kind of infrastructure. Some AI needs may be served through cloud-based services. Others may justify on-premises or hybrid approaches because of performance, privacy, sovereignty, or cost predictability concerns. Some use cases may belong in research environments. Others may belong in administrative workflows. Still others may require governance before they require technology.

Instead of simply asking how to become an AI-enabled institution, two of the best questions a technology leader can ask is:

1. What would meaningful success look like 12 months from now?
2. What is the smallest responsible set of investments needed to make that success possible?

A framework for technology decision-making: VALUE

Once you're ready to move forward with technology investments, it helps to apply a framework to help with your spending decisions. I've developed a quick mnemonic to help you remember the factors to consider before you make datacenter investment decisions.

V: Verify the institutional driver

Before evaluating any datacenter or AI investment, institutions should first ask a basic question: what exactly is driving this conversation?

- Is the driver operational inefficiency?
- Is it infrastructure fragility?
- Is it the need to support research growth?
- Is it a desire to enable AI use cases in teaching and learning?
- Is it data protection?
- Is it compliance?
- Is it pressure from leadership to "do something with AI"?

These are not the same drivers, and they should not be treated as if they are. The clearer the driver, the better the decision.

A: Assess readiness and dependencies

This is the step where many institutions discover that what they thought was a single investment is actually a chain of interdependent needs. AI readiness, for example, is not only about compute.

It is about data architecture, security, governance, identity, networking, storage, operating models, staffing, and support processes. Datacenter modernization is not just about replacing old equipment. It is about how power, cooling, lifecycle management, resilience, virtualization, automation, hybrid architecture, and team capacity all fit together.

In my experience, institutions get into trouble when they invest at the visible tip of the stack while neglecting the layers underneath that make the visible part sustainable.

L: Link to measurable outcomes

At some point we need to look at what we want to achieve as an institution. Every proposal should be tied to outcomes that can be measured, even if some of those outcomes are more operational than financial. Factors such as:

- Reduced downtime
- Faster provisioning
- Better energy efficiency
- Improved resource utilization
- Lower support burden
- Better performance for research workloads
- Stronger service continuity
- Greater adoption of supported AI use cases

are tangible outcomes that can be measured. If the institution cannot clearly define what success would look like, then the proposal is probably not yet mature enough to scale.

U: Understand the investment type

One of the reasons technology discussions become confusing is that not all investments are trying to accomplish the same thing. In my view, most datacenter and AI investments in higher education fit into one or more of four categories.

1. Efficiency investments - improve productivity, reduce waste, lower cost, simplify operations, or free up staff time.
2. Capability-building investments - create new institutional capacity, such as the ability to support advanced analytics, research computing, AI experimentation, or modern service delivery
3. Risk-reduction investments - improve resilience, security, recoverability, governance, lifecycle discipline, or compliance posture.
4. Strategic-advantage investments - position an institution to differentiate itself over time, whether through research strength, stronger academic offerings, better partnerships, or improved institutional competitiveness.

E: Execute in phases

Institutions should not attempt to solve the entire future in one procurement cycle. A measured approach is almost always the better one:

- Stabilize the core
- Modernize the foundation
- Pilot the right AI or data capabilities
- Validate value
- Then scale

Executing in phases is not a sign of hesitation. It is a sign of maturity.

Communicating the case for investment

Even the right investment can fail to gain support if it is communicated poorly.

Technology leaders must translate datacenter and AI proposals into the language of institutional consequence. That means the conversation changes depending on who is in the room.

- **With finance leaders** – the emphasis should be on lifecycle cost, supportability, utilization, avoided disruption, the cost of maintaining fragile systems, and the cost of deferral. In many cases, the strongest argument is not only what the institution gains by acting, but also what it risks by waiting.
- **With presidents** – the conversation should be centered on mission execution, institutional resilience, academic competitiveness, and the ability to support the institution’s future with credibility.
- **With provosts and academic leaders** – the language should shift toward research enablement, support for teaching innovation, academic productivity, and the institution’s ability to create better conditions for faculty and students.
- **With boards** – the case should be disciplined and strategic. What decision is being made? Why now? What risks are being reduced? What future capability is being enabled? What governance exists to ensure the institution is not acting impulsively?

In all these cases, I think there is one communication principle that tends to work especially well: do not present a technology investment as an isolated technology request. Present it as a measured institutional decision with a clear rationale, a defined sequence, and a transparent articulation of both the benefits of action and the consequences of inaction.

Top 10 technology considerations for higher ed leaders

<u>Collaborative Cybersecurity</u>	Building a cybersecurity culture of shared responsibility, end-user awareness and training, and improved access to security services and supports
<u>The Human Edge of AI</u>	Empowering students, faculty, and staff to engage with artificial intelligence tools critically, creatively, and safely
<u>Knowledge Management for Safer AI</u>	Mitigating the risks of artificial intelligence by integrating knowledge management into data governance, privacy, and ethics programs
<u>Data Analytics for Operational and Financial Insights</u>	Leveraging data analytics to provide insights into spending patterns, enrollment trends, and areas for cost savings and operational efficiencies
<u>Measured Approaches to New Technologies</u>	Making better technology investment decisions (or choosing not to invest) through clear cost, ROI, and legacy systems assessments
<u>Building a Data-Centric Culture Across the Institution</u>	Expanding and improving data access and unlocking the full potential of data as a strategic asset
<u>Technology Literacy for the Future Workforce</u>	Supporting discipline-specific technology training and education to enhance student success with in-demand technology skills
<u>From Reactive to Proactive</u>	Using data for scenario modeling, forecasting, and prediction to strengthen institutional agility and planning
<u>AI-Enabled Efficiencies and Growth</u>	Using artificial intelligence, robotic process automation, and other analytics capabilities to reduce operational costs, streamline processes, and improve strategic and business decision-making
<u>Decision-Maker Data Skills and Literacy</u>	Enhancing the value of institutional data by training and equipping decision-makers to use and interpret it properly

Technology investments need educated decision-making

Higher education does not need more technology ambition that lacks discipline. It needs better decision-making.

That is especially true now, as institutions try to modernize core infrastructure, respond to AI pressure, protect themselves against rising operational and security risk, and do all of this while carrying significant financial and staffing constraints. In that environment, the goal should not be to spend as little as possible, nor should it be to move as quickly as possible. The goal should be to decide as well as possible.

That means distinguishing between:

- Efficiency investments
- Capability-building investments
- Risk-reduction investments
- Strategic-advantage investments.

It means using ROI and ROV together. It means sequencing carefully. It means resisting impulsive AI spending. And it means recognizing that some of the most important technology decisions institutions make are not really about technology at all. They are about institutional fitness, institutional resilience, and institutional direction.

In the end, the institutions that will navigate this moment most effectively will not necessarily be the ones that buy the most. They will be the ones that think the clearest. That, in my opinion, is where the real value is.



About the Author

Dr. Hernán Londoño serves as Lenovo's Chief Technology and Innovation Strategist. With over 25 years of experience as a Chief Technology Officer (CTO) and Chief Information Security Officer (CISO) in higher education and as a Chief Technology and Innovation Strategist in private industry, he is known for his expertise in innovative academic solutions, cybersecurity, and strategic technology advancement.

Dr. Londoño has been a trusted advisor for several prominent cybersecurity organizations, including Vectra, Cyphort, and SANS Securing the Human. He is actively involved on various advisory boards, such as the National Applied AI Consortium (NAAIC) Business Industry Leadership Team, the FBI-InfraGard South Florida, and the Florida International University's ATOM Think Tank.

Dr. Londoño holds a B.S. and M.S. in Computer Science and Information Technology from Barry University, a Ph.D. in Computer Science from Nova Southeastern University, and a Post-Doctorate Certificate from Harvard Business School, focusing on AI in business competitiveness. Currently Dr. Londoño's academic research centers on the human aspect of the impact of AI in cybersecurity risk.

Sources:

1. <https://www.insidehighered.com/opinion/views/2024/02/05/most-colleges-finances-are-biggest-challenge-opinion>



Explore Lenovo Solutions
For Higher Education [here.](#)