navanta

IT SECURITY INSIGHTS | NOVEMBER 2026

AI Is Already in Your Institution. Is Your Security Policy Keeping Up?



Your institution may not have decided to adopt AI yet. Your employees and your vendors may have already made that decision for you.

Someone finds a tool that can summarize a document in seconds. A familiar software provider adds an AI feature in its next release. An employee copies part of an email into a chatbot because they need help drafting a response.

No project plan. No implementation meeting. No announcement that your institution has officially entered the AI era. It just starts being used.

That's why "we aren't really using AI yet" is becoming a dangerous assumption. You might be. You just might not know all the places it's happening.

Most people aren't trying to break the rules

The employee who finds an AI tool that saves an hour isn't trying to create a security problem. They're trying to get their job done. And there's the rub.

The work didn't disappear because the institution hasn't approved an AI tool. The report still needs to be summarized. The documents still need to be compared. The customer still needs an answer. If the only answer employees hear is no, some will find another way to get the work done.

That's how shadow AI can begin. It doesn't have to look particularly dramatic. Someone uploads a document to summarize it. Someone pastes text from an email into a personal AI account. Someone finds a free tool online that does in five minutes what normally takes an hour. The problem is that we don't always notice everything that travels with the information we are copying. A customer's name. An account number. An internal comment. Something buried farther down in an email chain.

Now it is not really an AI problem. It is a data problem. Good AI governance must account for how people work, not how we wish they worked.

 © 2026 Navanta™ All rights reserved.

The vendor you approved may not be the vendor you have today

Think about a software provider your institution approved two or three years ago. Same company. Same contract. Same login screen. But somewhere along the way, AI may have shown up.

It may be part of the product. It may be analyzing information in the background. The vendor may be using it internally in ways that never existed when you completed your original due diligence.

So, some of those questions need to be asked again. Does the vendor use AI in delivering its service? Does your data interact with it? Where does that data go? Is it used to train a model? Who can access it? What happens to it when your relationship with the vendor ends?

For new vendors, make those questions part of normal due diligence and contracting. For existing vendors, use annual reviews and product updates as a reason to ask again. Your institution doesn't need another vendor-management process just for AI. It needs its existing process to account for the technology vendors are using now.

Here's a better test of your AI policy

Stop someone in the hallway and ask: Which AI tools can you use for work, and what are you allowed to put into them? If the answer starts with, "I think...," you probably have some work to do.

Because the policy doesn't help much if nobody knows what it means when they are sitting at their desk trying to get something done. Employees need to know which tools are approved. They need to understand the difference between public, internal, and customer information. And they need to know where the hard stops are.

They also need somewhere to go when the answer isn't obvious. That doesn't require a longer AI policy. It requires a clearer one.

"It's just a pilot."

Pilots feel temporary. Small. Experimental. The data doesn't know that.

If a pilot can access information that needs to be protected, the institution still needs to know what the tool is doing with it. Someone still needs to own the use case. The right controls still need to be there.

There's another reason to get this right early. Useful pilots have a habit of sticking around. Someone finds a better way to review documents or reconcile information. It works. A colleague starts using it. Then another.

Before long, the pilot is not a pilot anymore. It's just how the work gets done. Trying to wrap governance around it afterward is a lot harder than putting guardrails in place at the beginning.

 © 2026 Navanta™ All rights reserved.

Someone will eventually ask why

AI decisions have a longer life than the meeting where they are made.

An auditor may ask why a tool was approved. An examiner may want to understand what information it can access. The board may ask what controls are in place. Someone who inherits the process two years from now may simply need to understand how it works.

The answer can't live in the memory of the people who were in the room. Document what was approved. Why it made sense. What information was involved. Who owns it. Where human review is required. What controls were put in place and whether they work.

This doesn't require putting AI into its own risk universe. The disciplines bankers have relied on for years still apply information security, privacy, vendor management, access controls, business continuity and risk assessment.

AI changes some of the questions. It doesn't erase everything you already know about managing risk.

Don't make security the reason good ideas never happen

There is a flip side to all of this. AI can take repetitive work off people's plates. Used well, it can give employees more time for the things technology isn't particularly good at: judgment, problem solving and relationships.

That's worth pursuing. The goal isn't zero risk. It's knowing which risks you are taking, deciding which ones are worth taking and putting controls around them that make sense. There is a big difference between accepting a risk you understand because the benefit is worth it and accepting a risk you didn't know you had.

So before worrying about the next AI tool your institution should adopt, answer four questions about the ones that may already be there:

1. What AI are we using?
2. What information can it touch?
3. Who is responsible for it?
4. Could we explain the decisions we have made if someone asked?

If you can't answer those four questions yet, hold off on the next AI tool. Start by finding those answers.

 © 2026 Navanta™ All rights reserved.