Building an Enterprise AI Governance Board from Zero

Case Study | January 2026

There was no playbook. No template I could pull from a framework document and hand to leadership. When I stood up the AI Governance Board at Queens University of Charlotte, I was working from first principles and 25 years of watching organizations get technology adoption wrong.

Most governance failures share the same root cause. Someone bought a tool, deployed it broadly, and asked the security team to figure out the risk after the fact. I was not going to let that happen with AI.

The Trigger

It started the way it always starts. Vendors showed up with demos. Leadership got excited. Departments started buying tools on their own.

Within a few months, I was finding AI-powered products embedded in workflows that had never gone through a security review. Marketing had one. Admissions had one. Individual faculty members were feeding institutional data into tools they found on their own. Nobody was asking what data these tools were ingesting, where it was being stored, or what the vendor's training policies looked like.

The behavior was not malicious. It was enthusiasm outrunning the process. But from a security perspective, shadow AI adoption is like shadow IT with a larger blast radius. A SaaS tool you did not approve might expose credentials. Your whole data environment could be exposed by an AI tool that you did not authorize.

I did not need to build a theoretical case for governance. I just had to show leadership the list of AI tools already running in our environment that nobody in IT had approved. That conversation lasted about ten minutes before I had the mandate to build a governance structure.

The Build

I started with stakeholders, not technology. People, process, technology. In that order. Always.

The board needed representation from IT, academic affairs, legal, HR, and student services. Not because I wanted a large committee. Because AI touches data flows across all of these functions, a governance structure that does not include the people who own the data is merely a security team talking to itself.

I kept the initial scope narrow on purpose. We were not trying to write a comprehensive AI policy for higher education. We were trying to answer three questions:

1. What AI tools are in our environment today?
2. Who approved them, and under what criteria?
3. What data are they accessing?

If you cannot answer those three questions, you are not ready for a policy discussion. You are still in discovery.

The Framework

I mapped our governance model to NIST AI RMF because it gave us a shared vocabulary without requiring everyone on the board to become a security practitioner. The framework breaks AI risk into four functions: Govern, Map, Measure, Manage. That structure made it easy to assign ownership.

Govern belonged to the board itself. We set the charter, the meeting cadence, and the escalation paths.

Map was my team's responsibility. We cataloged AI tools, data flows, and integration points across the environment. You can only manage the risk you have mapped.

Measure required defining what "acceptable" looked like for our institution. A research university and a 2,500-student teaching institution have different risk tolerances. We had to define ours, not borrow someone else's.

Manage was the ongoing work. Remediation tracking, exception handling, and periodic reassessment. This is the part most governance programs skip because it is not exciting. It is also the part that determines whether governance is real or performative.

What I Would Do Differently

I underestimated the communication challenge. Technical people and academic leaders think about risk differently. I spent too much time in the first few months translating between those two worlds when I should have established a shared risk vocabulary from day one.

I also should have started the data classification effort earlier. You cannot govern AI access to data you have not classified. We were doing governance and classification in parallel, which created friction that was avoidable.

And I would have built the tool inventory faster. Every week you spend without a complete picture of what is running in your environment is a week where a new tool gets adopted without review. Speed matters here more than perfection.

The Takeaway

Standing up an AI Governance Board is not a technology project. It is an organizational change management effort that happens to involve technology. The technical controls matter. The risk frameworks matter. But the thing that determines success or failure is whether you can get a room full of people who do not report to each other to agree on how decisions get made.

I have done this once now. The second time will be faster. Not because the technology is simpler, but because the organizational patterns repeat. Every enterprise has the same shadow AI problem. Every leadership team needs the same concrete inventory before they take AI risk seriously. Every governance board has to fight the same tendency to write policy before completing discovery.

If you are starting from zero, start with the inventory. Find every AI tool running in your environment that nobody approved. The rest follows.

Larnel Hight is a CISSP-certified security professional and founder of Render Defense, a cybersecurity consultancy specializing in AI security governance. I currently serve as a Senior Information Security Engineer at Queens University of Charlotte, and previously was an Enterprise Architect at Aflac