

Op-ed: New York's AI problem isn't the RAISE Act. It's ownership.

[Gift Article](#)
[Share](#)


AI's promise is still not fully realized. (GETTY)

By Dan Leiva June 25, 2026 11:00 AM EDT

New York recently signed the RAISE Act into law, and it takes effect Jan. 1. This legislation requires the largest AI developers to publish safety protocols, report critical incidents within 72 hours, and submit to oversight from a new office inside the state Department of Financial Services.

It's a serious piece of legislation. And it does not address the AI governance problem most New York businesses actually have.

DAN LEIVA IN THE NEWS

The RAISE Act regulates frontier model developers: companies spending more than \$100 million in computing costs to train the most powerful AI systems on earth. That is a handful of firms. It is not the insurance company in Midtown processing claims through automated decision systems. It is not the financial services firm on Park Avenue using AI to approve or deny credit.

Those organizations are making thousands of AI-driven decisions every day. And most of them cannot answer a basic question: Who owns the outcome when the system is wrong?

This is not a regulation gap. It is a leadership gap.

The pattern I see repeating across industries is consistent. Organizations deploy AI, celebrate efficiency gains, and never redesign the operating model around what the technology now does.

The system works. The outcome fails. And nobody owns the space between.

Here is what that looks like in practice. A customer submits a request. One system classifies it. A second checks eligibility. A third applies policy rules. A fourth generates the response. A fifth routes it to the right channel.

If the answer is wrong, every team can truthfully say their step worked. Classification was accurate. Eligibility logic ran correctly. Policy applied as designed. The response generated without error. Routing delivered on time.

But the customer still got the wrong answer. And the gap between those steps belongs to no one.

This is the accountability pattern most governance frameworks miss. They are designed for single decisions. The frameworks break down when a decision is actually a sequence. And in New York's financial services, healthcare and insurance sectors, most AI-driven decisions are sequences.

Big thinking.
Within reach.

What works in a deck has
to work in the business.

[Learn more](#)

The RAISE Act will not fix this. It was not designed to. It regulates how models are built. It says nothing about how organizations operate them.

What New York business leaders need is not another compliance checklist. It is an operating model that answers three questions:

First, who owns the outcome of each major automated decision your company makes? Not a team. Not a committee. A named person with the authority to change the system when it produces the wrong result.

Second, what decision authority has your AI actually been given? Most organizations never explicitly define where AI can act on its own, where it should recommend and a human decides, and where humans must own the outcome entirely. Without those boundaries, automation expands by default.

Third, how do signals from customers and frontline employees reach the person with authority to act? If complaints travel through four layers of abstraction before reaching a decision-maker, you have a feedback sink.

New York is right to regulate AI safety at the model level. But the real risk for most businesses in this city is not a frontier model causing catastrophic harm. It is an operating model where AI makes consequential decisions and nobody is accountable for what those decisions produce.

Regulation tells you what is required. Ownership tells you what actually works.

New York has the first one. It still needs the second.

Dan Leiva is the founder of CXAmplify and the author of "Amplified: The Operator's Playbook for Scaling Human Potential in an AI World."