AI GOVERNANCE · TRAICEBACK INSIGHTS

The Compliance Officer's Expanding Perimeter

AI did not hand the compliance officer a new rulebook. It moved the line that let compliance treat technology as someone else's problem.

For nearly three years I set out to learn the technology. First the foundations of cybersecurity, then machine learning and the large language models that now unsettle everyone, then the security of AI systems themselves. I expected the payoff to be technical — and it was. But it also turned out to be more than that. What I carried out of it was not only how and why the models work, but where they belong in the compliance paradigm: a classic governance question, the one I had been asking my whole career in compliance, pointed at something new.

Ask most chief compliance officers whether network security sits within their mandate and the answer is an easy no. Firewalls, patch cadence, network segmentation — these are the province of IT and the managed service providers most advisers rely on. The compliance officer's job has never been to configure the plumbing. It has been to run a compliance program reasonably designed to prevent violations of the Advisers Act, and to supervise the people who do the technical work. That division of labor was sensible, defensible, and until recently, stable.

Artificial intelligence is breaking it — not by adding a new rule, but by dissolving the line the old division depended on: the line between an operational tool and regulated conduct.

Why the old line held

IT security could sit outside the compliance function because most infrastructure produces no conduct a regulator cares about. A firewall drafts nothing. A network switch gives no advice. When these systems fail, the firm has an operational problem — an outage, or at worst a breach handled under an incident-response plan. The compliance officer's interest was real but bounded, and handing the work to technical staff was not a retreat. It was the right allocation of expertise. They knew the machines. The compliance officer did not.

That boundary had already begun to move before AI arrived. The amendments to Regulation S-P, and the Commission's long attention to cybersecurity, pulled a piece of what looked like IT — safeguarding client information, incident response, breach notification — back into the compliance domain. Not because compliance officers began configuring controls, but because that work carries obligations that land on the program itself. The lesson was quiet and easy to miss: some of what everyone called IT had never been only IT. It was conduct that happened to run on a machine.

What AI changes

Start with what AI is not. It is not one thing, and it is not only the autonomous agents that draw the headlines. It is a machine-learning model scoring securities. It is a natural-language classifier routing client

email. It is a deep neural network buried in a vendor's product. It is a large language model drafting letters and summarizing research. The systems differ in how they work and in how clever they are. That difference is not the one that matters here.

What matters is that these systems, unlike a switch or a firewall, do not sit down in the wiring. They work inside the advisory process, next to the client. Give a model the job of drafting client-facing material and you have implicated the marketing rule. Give it research that feeds a recommendation and you have implicated the duty of care. Give it client data and you have implicated Regulation S-P and the duty of confidentiality. The tool now produces *conduct* — and conduct that falls within existing advisory obligations has always been the compliance officer's ground.

So the honest question is not whether a system is technical. Everything is technical. The question is where the system stands relative to the advice — and that is a line the compliance officer can draw. The drawing is a calculus, not a label, and three factors carry most of the weight.

The first is nearness to the advice. Some systems stay in the back office — the network, the storage, the locks on the door. Others sit up front, where the advice is made: research, recommendations, client communications, valuation, best execution. A firewall is far from the client. A model writing to the client is not. Nearness to the advice, not the sophistication of the technology, is what pulls a system across the line. This is why the breadth matters: a crude email classifier sitting close to the client can weigh more heavily in the calculus than a sophisticated model buried in the back office.

The second is review. A tool whose every output an adviser reads before it goes out is one thing. A tool that acts, or publishes, on its own is another. When a person stands at the gate, the system sits farther from the compliance officer's daily concern. Remove the gate and it moves closer. This is why the autonomous agents raise the stakes even when the underlying task has not changed — they do more of the work with no one reading it first.

The third is what happens when the system fails — and here the calculus turns from measurement into design. The instinct is to make the AI perfect, to drive the error rate to zero before trusting it. That instinct is wrong, and anyone who has worked in security already knows why. You do not secure a network by assuming it will never be breached. You assume it *will* break, and you build so that when it does, the damage is contained and the alarm sounds fast. AI deserves the same posture. A failure that stays inside the building — a slow server, a lost file — is one kind of harm. A failure that reaches a client, a filing, or a recommendation is another. The fiduciary obligation is not to deploy flawless AI, which is impossible. It is to deploy AI whose failures are bounded and caught: the blast radius held small by design, and the tripwire tuned sensitive enough, and quick enough, to sound before a contained problem becomes a client-facing one.

None of these three factors is a new idea. Nearness to the advice, the strength of the review, the reach and detectability of the harm — the compliance officer has weighed all three across an entire career. They are the tools used to decide what to supervise closely and what to leave alone. The calculus is not new. What is new is that the compliance officer was told, for years, that these systems were not theirs to weigh. AI is what forces the old tools back into use, on ground that was ruled off-limits.

Run the calculus and the answer is often plain. A firewall is far from the advice, sits behind many human hands, and fails inside the building. It belongs to IT. A model that drafts client material is close to the advice, may run with no one checking, and fails straight at the client. That one belongs to the compliance officer, whatever the org chart says.

Where the duty comes from

Notice what has not been said. No new rule. The scope did not expand because the Commission wrote something. The Commission withdrew its predictive-data-analytics proposal, and nothing changed, because nothing was waiting on it. The scope expanded because the tool moved into the work.

Follow that and the law arrives on its own. If AI output is conduct, and the conduct is already governed — by Advisers Act Section 206, by Rule 206(4)-7, by the books-and-records demands of Rule 204-2, by the fiduciary duty itself — then the rules for a firm's use of AI are not waiting to be written. They are the rules the firm already lives under, laid over a new instrument. That is firmer ground than any forecast of future rulemaking. It asks only that existing duties be read onto tools that plainly call for them.

The hard part

None of this asks the compliance officer to become an engineer. It does not require reading model weights, understanding backpropagation, or judging the number of epochs to run in a machine learning model. Building the controls stays with IT and the service providers, where it belongs. What cannot be handed off is the judgment of whether the firm's use of the tool squares with its duties. That judgment is the job. It outlives whatever machine the firm adopts next.

There is a skill gap, and it is real, but it is narrower than it looks. The compliance officer does not have to answer the technical questions. The compliance officer has to frame the right ones and recognize a weak answer — the same thing already done in every complex corner of the firm that is supervised but not personally run. In practice that means holding a few plain questions no one at most firms is holding today. Where is AI being used near the client, the research, the marketing, the data? For each use, what can go wrong in fiduciary terms, and is there a control that catches it — and catches it in time? Who reads the output before it reaches a client? Is the firm's reliance on the tool reasonable, and is the basis written down?

None of those is a firewall question. They are questions about supervision, about conduct, about whether reliance is reasonable — the compliance officer's native language, dressed in technical clothing. The compliance officer is not annexing IT's territory. The compliance officer is being asked to recognize territory that was arguably always theirs, and that AI has made too large to keep stepping around.

What holds and what does not

So can the old habit continue — hand the machine to IT and look away? For conventional infrastructure, largely yes. For AI, no. A poisoned document that pushes a model into flawed work is not an IT failure to pass down the hall. It is compromised material walking into the advisory process. The cause is technical. The problem is conduct. Treating it as merely an IT matter is the exact error AI will no longer permit.

The resolution is not a heroic compliance officer who has learned to read code. It is a plain division of labor, written down: the compliance officer holds the fiduciary lens, IT holds the controls, and the seam between them is on paper rather than assumed. What changes for the compliance officer is not the depth of technical knowledge required. It is the recognition that a growing share of what the firm calls technology now produces conduct the compliance officer already answers for.

Compliance officers spent years comfortable behind a line that kept the machines at arm's length. AI moved the line. It did not move it by regulatory fiat. It moved it by putting a machine inside the very work the compliance officer exists to supervise. The firms that adapt first will not be the ones with the most technical compliance officers. They will be the ones that stopped pretending the old line still held.

That, in the end, was what nearly three years of study left me with. Not just the technology, but the deeper understanding of where the new line is being drawn between the machines and the humans who must govern them. AI has only made that line harder to ignore.

Traiceback Solutions

Gordon Eng is Founder and Managing Director of Traiceback Solutions LLC, an AI governance advisory firm serving SEC-registered investment advisers and ERISA fiduciaries. geng@traiceback.com · traicebacksolutions.com