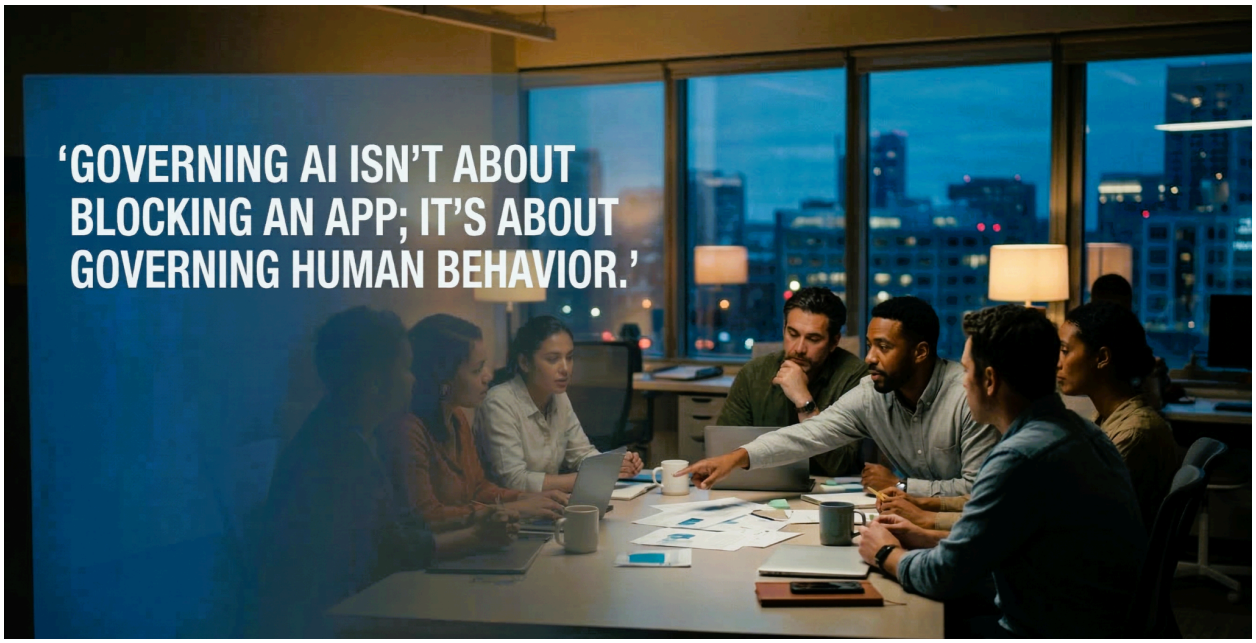

Your Firewall Can't Stop This One

Why AI governance requires a completely different playbook

Your Firewall Can’t Stop This One

Why AI governance requires a completely different playbook

This all started with a repost from my colleague, **Russ Marsh**, who shared a question that looked simple on the surface:

“If your IT staff are firing requests into free AI agents to solve issues about your architecture and operations, are they exposing: A) your IP, B) your technical architecture, C) your security gaps, or D) all of the above?”

The answer is **D. It’s always D.**
But the real issue isn’t *what* is being exposed — it’s why we keep being surprised by it, and why our response keeps failing.

© 2026 José A. Cedeño · Cofresí Consulting Services · All rights reserved · cofresi-consulting.net

The Brain Hasn’t Caught Up

Technology has outpaced our institutions, our governance models, and frankly, our instincts. We built compliance frameworks for a world where systems were discrete, access was controlled, and the threat surface fit neatly on a whiteboard.

That world is gone.

Humans are wired for immediacy. When an engineer hits a wall at 2pm, they’re not thinking about data-classification policy. They’re thinking: *I need an answer now.* So they open the fastest, smartest tool available — an AI agent with no NDA, no DPA, and no obligation to your organization.

This isn’t a people problem. It’s a design problem. And we’ve been applying the wrong solution for years.

The Playbook We Keep Running

When a new threat emerges, organizations reach for the same three levers:

- **Policy** — write a rule, publish it, add it to the handbook.
- **Training** — schedule the annual awareness session.
- **Enforcement** — monitor, audit, discipline.

These tools were built for a perimeter-based world. AI breaks every assumption underneath them.

Employees use AI from their phones, their home laptops, their instincts. You can’t train people out of a behavior that feels indistinguishable from productivity. And you can’t enforce a policy most employees don’t even realize they’re violating.

Here’s the line that captures the entire shift:

“*The bigger issue is this: our traditional governance model of policy, training, and enforcement was never built for tools this accessible and this embedded in daily problem-solving. All the compliance rules we’re implementing today are really just attempts to catch up with the evolution of technology. AI isn’t an app you can block at the firewall. It’s a behavior. And governing behavior requires a different playbook.*”

That’s the pivot point. That’s the truth most organizations haven’t absorbed yet.

What the New Playbook Looks Like

Most think pieces stop at the problem. The value is in what comes next.

1. Stop trying to prevent — start trying to channel.

© 2026 José A. Cedeño · Cofresí Consulting Services · All rights reserved · cofresi-consulting.net

Your staff *will* use AI. The question is where. Enterprise-licensed tools with DPAs aren't a luxury; they're the pragmatic middle ground between prohibition that doesn't work and exposure that does damage.

2. Make the secure path the easy path.

Behavior follows friction. If the approved tool requires three logins and a ticket, people will use the free one in a browser tab. If it's one click away and works well, they'll choose it.

3. Move the conversation to the boardroom.

This isn't an IT problem — it's a business-risk problem. The exposure happening right now (your architecture, your IP, your security gaps) is not theoretical. It's ongoing. Executives need to understand that the threat is already inside the building, and it walked in through the help desk.

4. Invest in private and on-prem AI for sensitive environments.

For critical infrastructure, regulated data, or proprietary systems, the long-term answer is AI that never leaves your environment. The technology exists. The ROI is clear. The barrier is organizational will — and that starts at the top.

The Confident Conclusion

Here's what executives need to hear: **this is solvable.**

Not by locking everything down. Not by pretending your workforce won't use the most powerful productivity tools available. But by accepting that AI governance is a new discipline — one that borrows almost nothing from the compliance frameworks of the past decade — and investing in building it properly.

Organizations that treat this as a technology problem will keep losing. The ones that treat it as a human-behavior problem, supported by the right technology, will get ahead of it.

The firewall was never going to stop this one. But the right strategy will.

---

*This post is part of an ongoing series on the intersection of human behavior, technology adoption, and organizational resilience. It goes along with this article: **The Biological Mismatch: A Personal Reflection on Technology, Power, and the Future** — where I explore why we're running 21st-century software on 50,000-year-old hardware, and what that means for how we govern, lead, and adapt in an AI-native world.*

© 2026 José A. Cedeño · Cofresí Consulting Services · All rights reserved · cofresi-consulting.net