The Line of Accountability

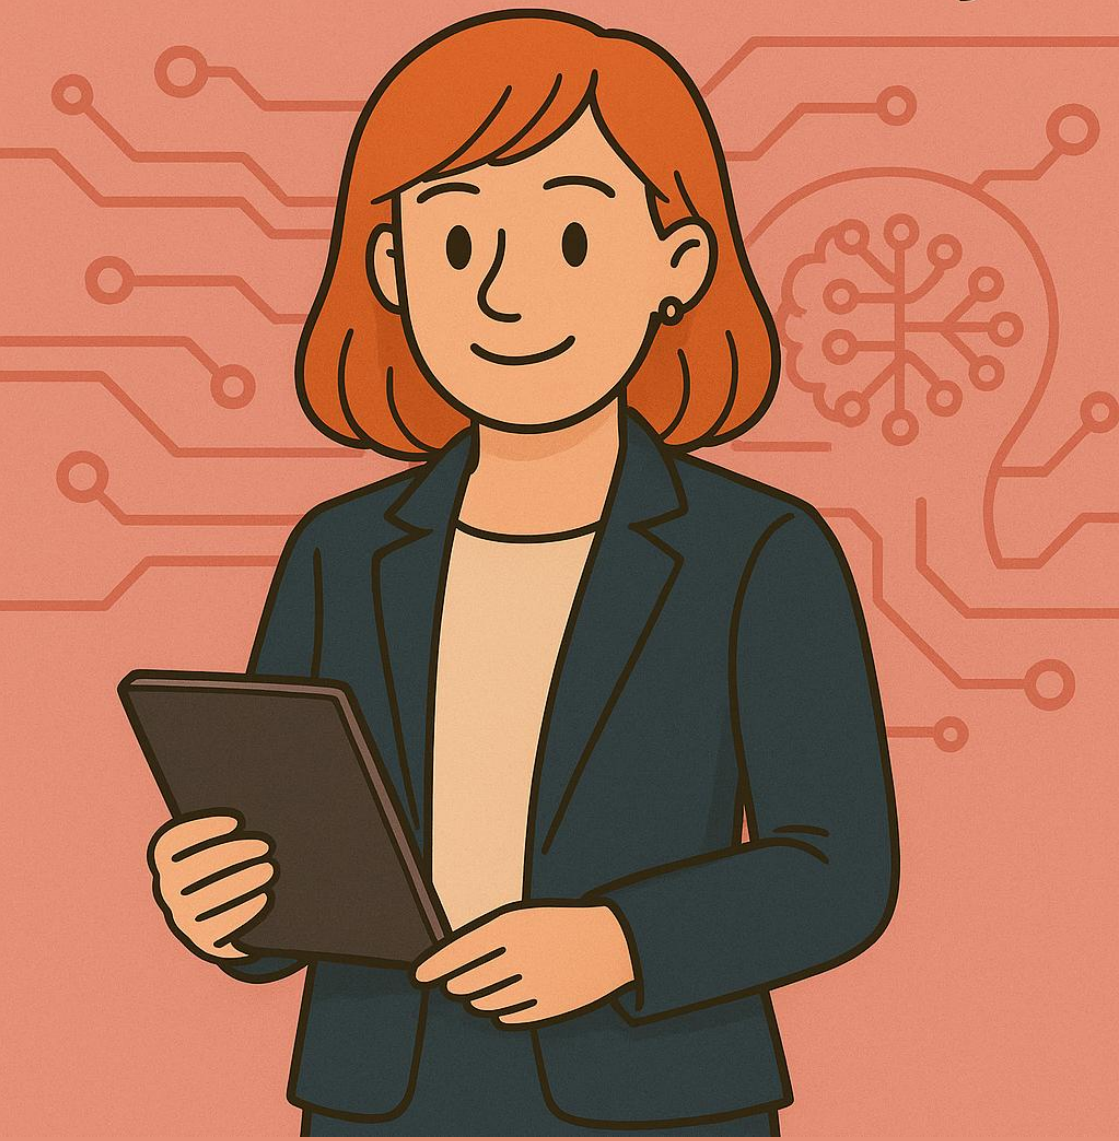

The Wake-Up Call

The email came in at 7:42 a.m. Subject line: URGENT: Client Credit Error – Escalated to CFO.

Ella scrolled through the thread. A high-value client had been downgraded by an automated credit risk model. The system had flagged anomalies in payment history, except there weren't any. A backend data ingestion error had misattributed another account's behaviour. The client's order had been halted. The sales team had lost face. The relationship was on the line.

The CFO's note at the bottom was blunt:

"Who approved this model? Who signed off on using it in production?"

Ella felt a quiet punch in the chest. She hadn't touched that model in months. It had been reviewed. Validated. Live. But the decision had evolved. Context had changed. Nobody caught it.

At 10:00 a.m., the executive leadership team assembled. Ella joined via video from her project room. She expected questions about the system. Instead, they asked about the process.

"Is this part of your AI governance framework?"

"Was it retrained?"

"Where's the human in the loop?"

"Who owns the outcome?"

Ella answered what she could. But the undertone was clear: if AI caused harm, someone would be held accountable.

Later that day, she met privately with Marcus from Risk. He didn't mince words.

"You built the scaffolding, Ella. That means the structure, the gaps, and the liabilities trace back to you. Even if you didn't write the code."

The truth was sobering. She hadn't broken anything. But she hadn't protected it either.

Key Learning: In AI, responsibility doesn't end at deployment. Real accountability begins when outcomes are questioned.

Model Drift

It took three days to find the root cause.

Ella sat in the analytics war room, surrounded by product managers, data scientists, and two deeply exhausted engineers. The credit risk model hadn't failed on day one. It had slowly degraded, like a compass that had been knocked off true north.

Model drift.

The term sounded almost poetic. In reality, it was brutally mundane. New data patterns had crept in, slowly changing the model's inputs, supplier delays, pricing anomalies, post-COVID variability. But because performance was acceptable and alerts hadn't triggered, no one had intervened. The model's accuracy had dropped just enough to cause harm, but not enough to set off alarms.

Ella clicked through the dashboards. A monthly retrain schedule had been proposed. Never implemented. The monitoring script hadn't been updated to reflect new decision thresholds.

She sighed. "This model didn't fail. We failed it."

Later, Ella briefed the AI Working Group.

"This wasn't just about the model. It's about stewardship. You can't just throw AI into production and hope it stays correct. Every model needs an owner. Every owner needs a protocol. And governance doesn't mean a PDF. It means practice."

The team sat quietly. For some, it was the first time they'd heard the term "model stewardship." For others, it was the first time it actually mattered.

In the hallway afterward, a junior product owner stopped her. "I thought these systems took care of themselves."

Ella paused.

"No. And if they do... be worried."

Key Learning: AI models degrade over time. So, governance must be continuous, not conditional.

Sign-Offs and Shadow Work

Ella was halfway through drafting a remediation plan when she got a Slack message from the Procurement Lead:

“Hey, is the supplier ranking tool still in scope for the AI Working Group? It’s live, btw. Scoring suppliers since Feb.”

Ella blinked. There was no record of that model in the AI registry. No approval form. No test log. No ethics checklist. It had never crossed her desk or anyone’s, for that matter.

She set up a meeting.

In the call, the Procurement team was upbeat. “It’s just a simple scoring matrix,” one manager explained. “It uses LLM output plus some heuristics. Cuts evaluation time in half.”

Ella pulled up a sample scorecard. A supplier had been flagged as “ethically high risk”. Based on what? The source was a scraped dataset, combined with ChatGPT summaries of public news mentions. No provenance. No legal review. No bias testing.

“How did this get deployed?” she asked flatly.

The answer was as old as time: “The vendor said it was ready. We didn’t want to wait.”

That week, Ella discovered three more “shadow” AI deployments. None had been reviewed. Two were still operating. One had already caused a supplier relationship to deteriorate.

She convened an emergency meeting with Legal and Compliance.

“We’re not dealing with AI usage anymore,” she said. “We’re dealing with AI exposure. And if we don’t map it, contain it, and govern it, someone else will. Like a regulator.”

That night, she stayed back and wrote three words on a whiteboard:

Discovery. Containment. Oversight.

The next phase wasn’t just about building AI right. It was about finding what had already been built wrong.

Key Learning: Unapproved AI use creates invisible risks. What you don’t govern can still harm you.

The Audit Room

Ella stood outside Conference Room 3A, laptop in hand. The audit team had arrived early.

This wasn't just any review. It was the company's first external audit of AI systems, triggered by new corporate governance requirements and a global shift toward algorithmic accountability.

The lead auditor, a sharp-eyed woman from a Big Four firm, greeted Ella with a nod.

"We'll need evidence of AI governance, design reviews, sign-off protocols, bias testing logs, model change history, and role assignments."

Ella opened her laptop and began presenting. She'd prepared a clean slide deck:

- AI Use Case Registry
- Model Lifecycle Policy
- RACI matrix for deployment
- Screenshots of the ethics checklist
- Retraining and drift monitoring documentation

The auditors flipped through. One stopped.

"This documentation is well structured," he said. "But how do you know people followed it?"

Ella hesitated.

"I... trust the process," she said.

The lead auditor looked up. "Governance isn't what's written. It's what's lived. Can you show us a traceable log of human interventions? Any sign-off records from Legal? Feedback loops from impacted teams?"

Ella's stomach dropped. Much of the structure was new. The culture hadn't caught up.

She promised to follow up with detailed logs and access records. But the message had landed.

On the way out, one auditor paused. “This is better than most. But the risk isn’t documentation gaps. It’s confidence without controls.”

That evening, Ella wrote a note to herself:

Build fewer frameworks. Embed more habits.

She’d have to go back to the teams. Not with slides, but with questions.

Key Learning: True AI governance is not documented. It’s demonstrated.

The Blowback

It started in the steering committee.

Ella had recommended a pause on the marketing attribution model after discovering it was over-attributing results to paid channels and under-valuing organic ones. Specifically, community-driven engagement that mattered to long-term brand equity.

She explained the model's flaw clearly. The data had shifted. Retraining was overdue. Risks were reputational, not just numerical.

But the Chief Marketing Officer was livid.

"You want us to stop running attribution analysis for two weeks? Do you realise we've just committed to board reporting on ROI uplift from our ad spend? Are you asking us to show up with nothing?"

Ella stayed calm. "I'm asking for a responsible delay to avoid drawing the wrong conclusions from a broken model."

The room was quiet. Then came the jab.

"Sometimes it feels like AI governance is just a reason to say no. We're trying to move fast and every time we do, you bring friction."

It stung.

After the meeting, a peer quietly pulled her aside. "You're doing the right thing, but it's a thankless position. People only love governance after a crisis."

Later that evening, Ella walked alone through the office. Governance fatigue wasn't just real, it was creeping in fast. Everyone wanted AI to work. No one wanted to slow down to make it work right.

She sat down and opened her laptop. Instead of updating the risk register, she wrote something different, a draft of talking points for a company-wide AI town hall.

It started with one sentence: "AI is not magic and friction is not failure."

If she was going to be seen as the blocker, then she'd make sure people understood why the brakes mattered.

Key Learning: Governance feels like friction, until you realise it's the only thing preventing a crash..

The Accountability Map

Ella walked into the strategy room with a single purpose: clarity.

For weeks, confusion had blurred the lines between AI ownership, approvals, and sign-offs. People assumed someone else was watching. No one really was. Now, Ella was here to fix that.

On the whiteboard, she drew four columns:

Use | Build | Approve | Own

She turned to the group, Legal, Risk, Compliance, Technology, and Business Unit leads. “We need to define who does what. Not who ‘consults’, but who’s accountable.”

One by one, the group populated the grid:

- Legal would approve data use and consent compliance.
- Tech would own deployment and versioning.
- Risk would validate thresholds and high-impact decision criteria.
- Business would sign off on AI decisions that affected customers, suppliers, or employees.

It took three hours and the conversation got heated.

“What if two models intersect?”

“Do we really want Legal involved every time?”

“What happens when a model changes post-release?”

Ella stayed firm. “This isn’t about slowing down. It’s about knowing who stands behind each outcome.”

They called it the AI Accountability Map.

Later, she met with the CFO to walk through it.

He scanned the roles and nodded slowly. “So if something goes wrong… we know who’s on the hook?”

Ella nodded. “And if it goes right, we know who deserves the credit.”

The map didn’t make AI less risky. It made AI less vague.

The next morning, the AI Working Group got an update: no model would go live without an accountability signature. One name. One owner.

And suddenly, everything became more real.

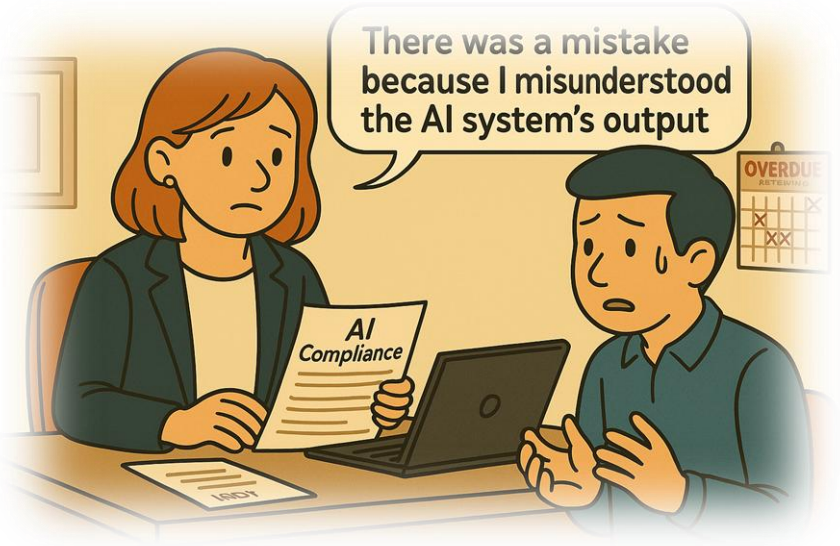

Key Learning: Accountability isn’t a concept. It’s a name on a decision with consequences.

The Risk That Spilled Over

The email subject was simple: URGENT – Employee Complaint Escalation (AI-related).

Ella opened the attachment. An internal complaint had been filed by a warehouse employee named Raj. He'd applied for a digital upskilling program run by HR, designed to help frontline workers transition into more technical roles.

Raj had been rejected automatically.

The rejection email stated: "Thank you for your application. Based on eligibility criteria, you were not shortlisted."

But Raj's manager had previously encouraged him to apply. He was qualified. Curious and frustrated, Raj raised a complaint and that's when things started unravelling.

The eligibility filtering had been handled by an AI model trained on past enrolment data. It had been tuned to prioritise "successful profiles", employees who had completed similar programs with high scores. But the training data was flawed. It overrepresented office workers. It penalised long tenure. It quietly excluded warehouse roles.

The AI hadn't just made a bad decision. It had institutionalised bias. Ella was furious.

She traced the model back to a vendor. HR had implemented it as part of a digital transformation sprint, bypassing the AI review process. "It was just a trial," they explained. "We didn't think it needed escalation."

But now it did. Raj had gone to the media. A journalist from a major outlet was asking questions about AI discrimination and fairness.

Legal scrambled. Comms drafted a response. Ella prepared a full model audit trail. In the middle of it all, she found herself face to face with Raj.

"I just wanted a chance," he said quietly. "Not a score."

Ella didn't have the words. But she knew this: they'd built something that dehumanised a decision that should have been humane.

Key Learning: Unchecked AI doesn't just get things wrong. It can hurt real people, silently and systematically.

Pressure and Precision

The boardroom was unusually full.

Ella stood at the front, flanked by the CIO and General Counsel. Across the table sat the Chair, the Audit Committee, and three non-executive directors who'd made AI oversight their new obsession.

The media coverage from Raj's complaint had shaken them. While the organisation had responded swiftly, pausing the model, apologising, offering Raj a tailored pathway. It had become clear this wasn't just a "tech" issue.

The Chair opened directly.

"We want to understand three things: one, what other risks like this exist; two, how we stop them; and three, who's responsible if they happen again."

Ella nodded. "Understood."

She clicked to the first slide, AI Exposure Map: Current vs. Controlled Use. Then came:

- A timeline of model reviews.
- A lifecycle policy with thresholds for escalation.
- An updated AI Charter draft, co-authored with Risk, HR, Legal, and Tech.
- A live example of accountability signatures for deployed models.

"Every AI system is now mapped to a named business owner, a technical steward, and a risk classification. No system runs without human oversight."

A director interjected: "That's fine in theory. But how will we know?"

Ella took a breath.

"Because the next time a decision hurts someone, we won't be asking who made the model. We'll already know who stood behind it."

The room was silent.

Then the Chair spoke.

“Ella, it’s rare for someone to advocate for guardrails in a way that still accelerates progress. Well done. But the pressure stays on. Keep proving this works.”

Ella nodded. She wasn’t after applause. What she wanted, needed, was sustained trust.

Key Learning: Accountability under pressure is the true test of leadership in AI.

The New Charter

The room wasn't fancy, just a borrowed training space with white walls, a screen, and too many power cords. But the people in it were the real deal: Legal, Risk, Compliance, HR, Technology, Finance, and Operations. Ella stood at the front with a simple objective.

"We're here to ratify a cross-functional AI Charter," she said. "This won't be a document that sits on SharePoint. It will govern every AI decision we make from now on."

The group had been meeting for weeks, hammering through debate after debate:

- Should low-risk AI tools be exempt from formal review?
- How do you handle generative models trained on sensitive internal data?
- Who signs off on public-facing AI outputs?
- What constitutes an AI incident worth escalating to the board?

Today was decision day.

They walked through the draft Charter section by section:

1. Definition of AI and Automation Boundaries
2. Governance and Oversight Structures
3. Risk Classifications and Treatment Plans
4. Ethical Principles: Fairness, Transparency, and Human Impact
5. Incident Escalation Frameworks
6. Board-Level Reporting and Responsibility Mapping

Some pages sparked friction. Others were locked in with quick consensus. But by noon, the final version was adopted.

Ella stared at the first printed copy. It wasn't perfect, but it was principled. It aligned accountability with influence. It gave the board visibility, teams confidence, and the public a reason to trust.

As the group disbanded, the CFO pulled Ella aside. “We need more than AI performance metrics from now on. Bring this Charter to every leadership session.”

She nodded. “Every AI model has a lifecycle. This Charter will be its conscience.”

Key Learning: Real AI governance lives not in frameworks, but in collective agreement on what’s non-negotiable.

What We Owe Each Other

The AI town hall was over.

Ella stood alone in the empty auditorium, staring at the final slide still projected on the screen: Trust is not the absence of failure, it's what you do next.

She'd spoken openly. About drift. About Raj. About governance friction, shadow systems, and board pressure. But also about something deeper: the duty of care owed to each other in a world increasingly shaped by algorithms.

Her voice had cracked slightly when she said it aloud:

"When AI fails, it fails in ways that people can't appeal, predict, or prepare for. That's why accountability isn't just a control, it's a contract."

A contract with colleagues, with customers, with society. One that says: we won't hide behind the system. One that doesn't wait for regulation to tell us what's right.

Outside, the sunset caught on the glass walls of the building. Ella gathered her things, walking slowly back to her floor. The AI Charter now lived in practice, not just policy. Each model had a signature. Each approval had a timestamp. And every team knew that AI wasn't a free pass, it was a responsibility.

In her office, she scribbled one last note before shutting her laptop:

We are not just building systems. We are deciding what kind of systems build us.

Ella knew the line of accountability wouldn't always be clear. But she also knew it had to exist, because leadership in AI wasn't about having all the answers. It was about showing up when the questions got hard.

And she would keep showing up.

Key Learning: In a world of intelligent systems, accountability is the most human act of leadership.

Credits & Copyright

Author

Dharshun Sridharan, Aresia

Illustrations / Image Creation
Generated with assistance from ChatGPT (OpenAI)

Contact
dsridharan@aresia.com