

The Eighth Conversation: Why Your Enterprise Needs a Chief Reliability Officer for AI

By Rahul Jindal

I have been in twenty AI steering committee meetings in the last year. Every one of them has the same seven people in the room and the same eighth question that no one in the room owns.

The seven are familiar. The CLO worries about risk and compliance. The CFO presses on ROI and unit economics. The CHRO holds the workforce conversation. The CTO defends the architecture. The COO worries about operational disruption. The CMO protects the customer experience. The CEO synthesizes. Each one is doing the job they were hired for.

The eighth question, the one nobody owns, is this: *is the AI we shipped six months ago still doing what we said it would?*

“Capability gets a seat at the table. Reliability does not. The Decay Tax is paid in the gap between them.”

Why the seven cannot answer it

Each of the seven is structurally disqualified.

The CTO will tell you the platform is up, the model serves traffic, the pipelines are green. None of that tells you the system's answers are still right. Uptime is not accuracy.

The COO will tell you the process runs. None of that tells you the process produces the same outcomes it did at launch. Throughput is not quality.

The CFO will tell you the invoice clears every month. None of that tells you the AI is delivering the ROI the original business case assumed.

The CLO is watching for catastrophic failure: the Air Canada moment, the discrimination lawsuit, the regulator's knock. Slow erosion does not register at that altitude.

The CHRO, the CMO, the CEO are even further from the instrumentation. The CMO learns from the customer escalation, not the dashboard. The CHRO learns from the team that quietly stopped using the AI tool. The CEO learns from the board.

By the time any of the seven hears about decay, the decay has already happened. The Decay Tax is the cumulative invoice for the time between when an AI system started slipping and when someone at this table noticed.

What the eighth conversation actually contains

The eighth conversation is not one question. It is a category of questions that today fall into the cracks between functions.

- **Decay measurement.** For each AI system in production, what is its half-life: the period after which 50% of its accuracy or utility erodes? Who owns that number?
- **Silent regression.** When a vendor swaps the underlying model, who catches the behavior change before users do?
- **Eval debt.** The gap between what was tested at launch and what users actually ask now. Who refreshes the eval set on a cadence?
- **Agent reliability.** In multi-step agentic systems, where does the compounding error live, and how is it bounded?
- **Incident response.** When an AI system produces a meaningfully wrong answer at scale, what is the runbook, who is paged, what is the SLA?
- **Lifecycle governance.** When does a model get retired? Who signs the deprecation memo? What does the migration plan look like?

Each of these is, today, somebody's tenth priority. None of them is anybody's first.

The role: Chief Reliability Officer for AI

The right structural answer is a named role with the eighth conversation as its first job. Different orgs will give it different titles (Chief Reliability Officer, Head of AI Operations,

Director of AI Assurance) but the function is the same: be the person whose calendar is built around the question the other seven are too busy to ask.

The closest existing analogue is the SRE (Site Reliability Engineer) function that Google formalized for software in the 2000s. SRE was not infrastructure (that was the platform team). It was not product (that was engineering). It was the named owner of *does the system actually work for users over time*, with hard SLOs, error budgets, and the political authority to block a launch that would burn the budget. SRE turned out to be the missing role. AI is in the same position now, twenty years earlier in its arc.

A CRO for AI inherits the same logic. The mandate is not capability; capability is owned by the data and ML teams. The mandate is *does this AI still do what we said it would, six months in, and how would we know if it stopped?*

“SRE is the precedent. AI Reliability is the same idea, twenty years earlier in its arc.”

Where the role sits

Three viable structural homes, in order of likely durability:

1. Inside the COO's office. This is where reliability instincts already live for everything else the business runs. The COO already owns business continuity, operational risk, and process integrity. Adding AI reliability is a natural extension. Pro: deep operational culture. Con: the COO's muscle is for human-process continuity, not for the specific shape of AI decay (model swaps, eval drift, agent loops). The role needs serious technical depth, which the COO's shop typically lacks.

2. As a peer C-suite role reporting to the CEO. The strongest political position. The CRO for AI sits at the same table as the CTO, with a mandate to question and, when necessary, block. Pro: real authority, real visibility, real budget. Con: most enterprises will resist creating a new C-level role; the politics of one more reorg cycle is expensive.

3. Inside the CTO's org with hard independence. The fastest to stand up. The CRO reports to the CTO but with an explicit mandate to publish findings independently, like

the model SREs and red teams that already exist inside frontier labs. Pro: fits existing org structure. Con: the conflict of interest is real. A CRO for AI who is graded by the CTO will under-report decay in CTO-favored systems.

The right answer for most enterprises is **Option 1 with technical depth bolted in** : a senior leader inside the COO's office with a small team of ML engineers, a budget for tooling, and a charter that specifically names the eighth conversation as their first job. Option 2 is correct for AI-native firms (Anthropic, OpenAI, and the few enterprises betting the whole company on AI). Option 3 is a stopgap, not a destination.

The 90-day mandate

If you appoint someone tomorrow, here is what their first ninety days should produce.

Day 1-30: the inventory. Every AI system in production. Owner, vendor, business process, last evaluation date, last model swap date. This is usually the first time anyone has assembled the list. The count alone is a useful surprise.

Day 31-60: the half-life baseline. For the top five systems by business criticality, design and run an eval that estimates current accuracy versus launch-day accuracy. Most enterprises will find that two of the five have degraded materially and nobody knew.

Day 61-90: the cadence and the budget. Define an evaluation cadence per system. Define an error budget per system. Stand up the runbook for a meaningful AI incident. Present to the board.

That is a ninety-day project that pays for the role for the year.

Why this is now urgent

The case for moving on this is timing. Enterprise AI deployment has crossed the volume threshold where decay starts producing visible failures at the press-cycle scale. Air Canada was twenty-four months ago. Cursor was a year. Replit was last summer. Each one was code-correct and system-broken at the moment it happened. Each one was, in retrospect, a Decay Tax event, paid by the company that had the failure, watched by every other CEO who is now wondering if their systems are quietly heading for the same headline.

The market will name this role within twelve months. The question for an enterprise leader is whether you want to be the one defining the role at your company, or the one reading about how Goldman Sachs or JPMorgan defined it first and wondering why your CRO for AI has the wrong charter.

What to do this week

You do not need a new C-suite role to host the eighth conversation today. You need a host. In your next steering committee, put one item on the agenda: *for each AI system in production, who is the named owner of its half-life?* If the answer is "nobody" for any system on the list, you have just identified your first hire, or your first charter. The Decay Tax is being paid either way. Naming the host is the only move that changes the math.

This essay extends two frameworks: the Decay Tax (what reliability erosion costs) and the Seven Conversations (the C-suite conversations that shape enterprise AI). It sits in the Multiframework Breakthroughs notebook as the bridge essay between the two. The corporate-finance cousin is Revenue that holds through a bad year: the same structural-resilience question applied to the top line, not the AI stack.

Shared privately. Please do not redistribute.