

Your board is measuring the wrong AI metric.

by Ben Samuels | General Manager EMEA, Invisible Technologies

Most enterprise AI programmes are succeeding at the wrong thing. Adoption is up. Licences are deployed. Users are active. In some organisations, employees are even being measured on log-ins and tool usage. And yet, in many cases, the underlying operation is unchanged.

Ben Samuels is General Manager EMEA at **Invisible Technologies**.

Invisible works with enterprises across financial services, pharmaceuticals, logistics and the public sector to move AI from pilot to production.

 **Invisible**
InvisibleTech.ai



I see this in almost every enterprise conversation I have. I ask leaders: how many AI pilots has your organisation run in the last two years? Then I ask: how many are delivering measurable outcomes such as real cost savings, revenue impact, risk reduction, faster throughput? The answer is often close to zero, in stark contrast to the current narrative.

“How many [of your AI pilots] are delivering measurable outcomes such as real cost savings, revenue impact, risk reduction, faster throughput?”

The adoption trap

Adoption measures whether people are using AI but provides no gauge on how the organisation is benefiting from it. The reason is that individual productivity gains don't translate into institutional results. They sit inside personal workflows, depend on individual behaviour and rarely survive team changes.

One COO I spoke with recently had built over 60,000 internal agents. It sounds impressive but in practice it created fragmentation. They had teams pulling in different directions, overlapping tools and problems being solved that were not anchored to company priorities. As she put it: “Everyone feels good. But we're pulling in too many directions.”

The inevitable outcome of decentralised AI is activity without alignment. The vast majority of enterprises are stuck at this first stage of adoption, getting people comfortable with the tools.

Stage two is where real value is created. To achieve ROI, you need structured data, systems designed for the edge cases and an operating model shift.

1. The data foundation problem

Most enterprise software is over twenty years old. It wasn't designed to talk to other systems, let alone feed an AI agent. Data sits across legacy ERPs, mainframes, inboxes, PDFs and spreadsheets and is unstructured, inconsistent and inaccessible. This is the reality of enterprise infrastructure and explains why AI is not yet reaching production at scale.

The good news is that you don't have to replace your technology stack. However, you do need to clean and organise your data so that agents can run reliably. Then you begin to unlock decades of institutional knowledge, customer behaviour and operational history to drive true advantage.

2. Enterprise-grade reliability

In our personal lives we tolerate AI being mostly right. In regulated, operational or customer-facing workflows, a system that is right 80% of the time is a liability. It has to be more reliable than what it replaces. Reaching that standard requires more than deploying the latest model. You need to understand the workflow deeply enough to know where human oversight is required.

Take pharmaceutical submissions. AI can draft, structure and cross-reference regulatory documents faster than any team. But a filing error has consequences that go well beyond the document. The human is not there to babysit the AI. They are called at specific points in the workflow, such as exception handling, where judgement and accountability matter most.

3. Who should own AI?

The organisations getting this right are not running AI as a technology programme. They are running it as a business transformation with technology as the enabler.

The leaders closest to revenue, operations and risk are the ones who understand where value is created, where the bottlenecks are and what failure actually costs. They need to own the strategy, set the priorities and be accountable for the outcomes. Technology teams are critical to making it work but they are not the right people to define what good looks like.

When the business owns AI and measures it against outcomes, things change.

When IT owns it, you get a roadmap.

Conclusion

The organisations pulling ahead are measuring impact: throughput, conversion, cost to serve, service levels, revenue and risk. The key governance question for any board is not ‘how many employees are using AI?’ but ‘where is AI changing the economics of our most important workflows?’

The companies that continue to optimise for usage will see incremental gains that never fully materialise. The companies that redesign how work gets done will create a compounding advantage that is difficult to catch.

The gap between AI enthusiasm and enterprise results is in the operating model.

That is where the real work begins.