



CIONET

DISCUSSION SUMMARY

ENTERPRISES
DON'T HAVE AN AI
PROBLEM, THEY
HAVE A DATA
PROBLEM



There's a comforting story enterprises tell themselves about AI: that success comes down to picking the right model, the right vendor, the right roadmap. But like most best laid plans it rarely survives contact with reality. Two organisations running the same AI stack can produce very different results, and the gap seldom has much to do with the model itself.

Ask why, and the same word kept coming up across sectors with almost nothing else in common: ownership: Who's accountable for a dataset? Who decides what "good" looks like? Who is willing to be measured against what it produces?

A recent CIONET Community gathering of senior technology leaders, which featured expert speakers from luxury goods, energy, government and professional services, made that pattern hard to miss: even with different regulatory pressures, and different starting points, they had almost identical diagnoses of where their AI ambitions were actually getting stuck.

The model was never the constraint

More than one leader pushed back on the idea that AI maturity is mainly about picking the right tool. The pattern they described was consistent: initiatives stalling not because the technology underperformed, but because the data feeding it was inconsistent, ownerless, or simply not trusted enough to act on. One framing recurred: programmes fail when they're run as IT projects with an end date, rather than as business capabilities somebody owns indefinitely.

Several attendees commented something close to the same thing, independently naming poor data foundations, not the AI itself, as the real constraint. One comment captured the room's mood: the advantage doesn't come from which model you choose. It comes from what's built underneath it, and the culture that decides how it gets used.

Ownership is the real governance battle

Governance produced the sharpest anecdote of the evening. One leader recalled being told "We need to clean up all the data" as a precondition for doing anything useful with it, and being able to do the maths on the spot: at that rate, they'd still be at it in thirty years. The

alternative floated, rougher but more workable, was an 80/20 split: govern properly the minority of data that actually drives decisions, and deliberately leave the long tail alone.

Governance mechanics, though, turned out to be the easier problem. Ownership was harder to pin down, and one leader put it most sharply - in their organisation data ownership had been pushed onto technology almost by default, not because it belonged there, but because nobody in the business had claimed it. The result: a team maintaining systems was quietly patching business definitions it had no authority to set, and didn't especially want. "I don't want to own data at all," as the leader put it. Framed more constructively, at least one attendee made the same point: data needs to sit with the people accountable for the decisions it informs, not with whoever happens to hold the infrastructure.

Two other examples, from very different contexts, pointed at the same underlying problem from another angle. Matrix structures, technology specialists embedded inside business units, but still reporting centrally, came up as a partial fix, credited with narrowing the gap only when paired with sustained relationship management rather than a one-off reorganisation. Separately, a data programme that had deliberately pushed decision-making out to the business regressed once the small number of people driving it moved on; teams reverted to hoarding data instead of sharing it. Different situations, same lesson: ownership doesn't hold unless someone keeps actively defending it.

Escaping Perpetual Pilot Mode

This frustration surfaced more than once: organisations running numerous AI pilots with almost nothing reaching production. One leader described inheriting more than forty AI "experiments" on arrival, and not one of which had scaled. The explanation on offer was blunt: a culture that celebrates the small technical win and stops there, rather than doing the harder, duller integration work that turns a demo into something that actually runs.

One proposed fix inverts the usual effort split. Most organisations, in this telling, put roughly 80% of budget into proving individual use cases work in isolation, leaving a fifth for the necessary last mile: integrating with existing systems, managing the change internally, decommissioning whatever the new capability replaces. Flip that ratio, was the argument, and pilots start reaching production.

A related idea, one level down: build reusable components (assets, interfaces, compliance checks) that plug into more than one use case, instead of re-solving the same integration problem every time a new experiment shows promise.

One expert offered a more hopeful data point: over the past six months, commercial and revenue leaders have started turning up to these conversations alongside technology, rather than being briefed after the decisions are made. Asked about cost, in his account, they increasingly answer in terms of new revenue and new products rather than savings, a shift he reads as a genuine change in how these initiatives get funded.

Proving value without chasing savings

Cost came up constantly, and the room didn't agree on how to talk about it. One leader was matter-of-fact about savings being the point: legacy-system modernisation using AI was on track to finish three years early. Two others, both in commercial businesses under pressure to grow, argued closer to the opposite: that leading with cost savings is usually the wrong story to tell, because "we're cutting the team" is a much harder conversation to have than "we found new revenue." One of them was candid about running two versions of the same argument at once: cost discipline internally, where cutting is simply expected, and growth and value in board conversations, where it actually counts. The difference wasn't really philosophical. A regulated body spending public money answers for cost in a way a retailer mid-turnaround usually doesn't.

Where the value argument worked, it worked at real scale. One case, described in some detail: a customer segmentation project that cut processing time from 72 hours to 120 seconds, opened a 28-day promotional window that hadn't existed before, and generated eight-figure incremental revenue within six months. In the public-sector case, the equivalent wasn't a revenue number but a specific traceable outcome; the point made there was that "we need solid data foundations for AI" used to be enough to justify a budget on its own, and as of roughly six months ago, it stopped being enough.

Underneath both versions of the argument sat a harder, shared problem: the running costs of agentic systems are difficult to forecast, because usage driven by autonomous agents doesn't scale the way historic human-driven API traffic did.

More than one leader described building continuous monitoring across cost, token usage and compliance for exactly this reason, alongside a tougher line in vendor negotiation, one recounted a well-known consultancy whose AI bill "went absolutely through the roof" after being rolled out with no constraints at all. Nobody pretended the vendor relationship was collegial: there was open acknowledgement that some partners, systems integrators especially, will capture value that should belong to the client if nobody stops them.

The human-agent equation

A recurring anxiety concerned what AI does to people, not to processes. One idea, offered without much hedging: business leaders don't need deep technical expertise, just enough fluency, "30% literacy" in the phrase used, to ask sharp questions and challenge what a system produces.

A more provocative framing came from elsewhere in the discussion: judge the combination of a person and their AI agents together, rather than either alone, since some organisations are already tracking machine usage the way they once tracked hours worked. Set against that was a cautionary case: an organisation that replaced a large share of its customer-facing workforce with autonomous agents, saw clear short-term cost savings, and then lost more than a third of its customers, because empathy and complicated problem-solving proved harder to automate convincingly than expected.

The takeaway wasn't "don't automate." It was: know exactly which parts of the job you're handing over, and don't guess on the empathy-heavy ones.

Where the human hand still matters

One leader, put this most starkly: audiences notice, and mock, when AI-generated content carries visible "fingerprints" (the phrase used was people getting "laughed at" online over AI images with obvious tells). The response there was to confine automation to internal, operational work, removing a coat hanger from a product photo, say, and keep anything customer-facing or creative under close human control. Less because AI couldn't do it, more

because a brand's whole value proposition can rest on the appearance, and human craftsmanship.

At least one attendee made almost the identical point: whatever AI ends up doing, protect the people who bring curiosity and creativity to the work. It's a narrower point than it first sounds, and a useful one: knowing where deliberately not to automate is becoming as strategic a decision as knowing where to.

What comes next

Nothing in this discussion suggested that the more capable organisations simply had better AI. What they had were older, less glamorous decisions already made: who owns the data, how rigorously "good" gets defined, how ruthlessly priorities get cut. Those decisions were compounding, in both directions.

None of what leaders described as increasingly important is really a technology problem. The discipline to refuse governance projects that never finish and pilots that never reach production sits with people, not systems. So does the commercial toughness needed to negotiate with AI vendors in a market where usage costs are volatile and vendors know it before the client does. And so, still, does the judgement to decide where automation earns its place and where it just quietly costs a business the trust it depends on.

For leaders who feel behind, the more useful diagnosis is probably not "we need better AI." It's ownership, culture, and the discipline of deciding what's actually worth governing properly. That's a slower, less exciting answer than a model upgrade. It also happens to be the one every sector in the room kept arriving at, without much prompting.





About CIONET

CIONET is the leading community of more than 10,000 digital leaders in 20+ countries across Europe, Asia, and the Americas. Through this global presence CIONET orchestrates peer-to-peer interactions focused on the most important business and technology issues of the day. CIONET members join over a thousand international and regional live and virtual events annually, ranging from roundtables, programs for peer-to-peer exchange of expertise, community networking events, to large international gatherings. Its members testify that CIONET is an impartial and value adding platform that helps them use the wisdom of the (IT) crowd, to acquire expertise, advance their professional development, analyse and solve IT issues, and accelerate beneficial outcomes within their organisation.