



200 IT Leaders **Six Months**

What's actually blocking AI progress

**FIELD INTELLIGENCE FROM VERSION 1'S
ENTERPRISE TECHNOLOGY LEADERSHIP FORUMS**

What's blocking AI progress

Over the past six months, Version 1 has hosted more than 200 senior IT and technology leaders across a series of closed-door forums and roundtable events, bringing together CTOs, CIOs, Heads of Technology, and senior programme leaders from commercial organisations across the UK and Ireland. These were senior operators with budgets, boards to answer to, and live AI programmes to manage, and the conversations reflected that.

What follows is our analysis of the most consistent themes that emerged across those events. Not what the analyst reports say, not what the vendor briefings claim, but what the people actually running these programmes told us, in their own words.

What every leader had in common

Before getting to the specific findings, it is worth establishing the common starting point, because it matters for how these findings should be read.

Every organisation represented had invested in AI in some form, most had live pilots or proof-of-concept programmes underway, and a significant proportion were operating under explicit executive-level mandates to scale. Despite all of that, very few leaders felt their organisations were scaling at the pace their boards expected, and the gap between stated ambition and actual delivery was a recurring source of tension throughout our forums. This is our analysis of why that gap exists and what the organisations closing it are doing differently.



What to expect from this report and who it's for

Each finding comes with our analysis of the pattern, what the organisations getting it right are doing differently, and what it means in practice. Each finding also closes with a concrete recommendation you can act on immediately.

This report is based on conversations with senior technology leaders from commercial organisations across the UK and Ireland, the majority operating in regulated sectors. The findings are specific to the private sector, where the investment mandates, governance pressures, and delivery challenges differ materially from public sector contexts.

FINDING 1:

Governance has moved. Has your AI programme?

A photograph of Clare McConkey, a woman with long brown hair, wearing a black blazer over a white top. She is holding a white folder and gesturing with her right hand. A name tag on her left lapel reads 'Clare McConkey', 'Director - Financial Services', 'Version 1', and 'VERSION 1'.

Eighteen months ago, the honest answer for any leader operating in a commercially sensitive or regulated environment was that the tooling for responsible enterprise AI deployment was still maturing, security frameworks were incomplete, model behaviour was harder to audit at scale

The responsible default was to run controlled pilots while waiting for the governance landscape to settle into something more workable.

That has changed materially, and the leaders we spoke to who are operating in these environments are no longer treating governance as a reason to slow down. They are treating it as a capability to build, and the distinction matters because it changes where the investment goes and what the programme team spends its time on.

What we see working consistently is a technology-agnostic centre of excellence: a governing function that defines value at the process level, sets and enforces standards across the organisation, and retains control of outcomes regardless of which vendor's technology sits underneath it. This does not require a large central team; several of the most effective examples we heard about were running lean functions with the ability to scale capacity in and out as specific initiatives demanded it, rather than trying to staff for peak demand across the whole portfolio.



“The organisations still treating governance as an obstacle tend to be the ones that have not yet built this central function. The ones that have built it find that governance becomes a competitive advantage, not a constraint.”

What has replaced governance as the primary blocker is something more structural: the absence of a clear, working decision framework for what to build internally, what to buy from the market, and what to bring in as an externally managed capability. The organisations still stalled at the planning stage are, in the majority of cases, the ones that have not yet resolved this question, and without resolving it, every new initiative triggers the same debate from scratch.

RECOMMENDATIONS

Build governance as a capability, not a control gate.

Establish a small, technology-agnostic centre of excellence with clear authority over AI standards and outcomes measurement. Define your build-buy-partner framework explicitly so that every new initiative has a decision path rather than triggering a debate. If you are in a regulated environment, design your QA framework now, in parallel with your next delivery cycle, before you need it.



FINDING 2:

The real reason you're still stuck piloting


Most teams built pilots to prove a concept rather than to prove business value, and these are meaningfully different objectives that produce different outcomes. A proof of concept is designed to demonstrate that something is technically feasible. A business value test is designed to show what happens to a specific, measurable process metric when AI is applied to it.

The organisations that have broken out of the pilot trap are consistently the ones that made this distinction early and built their

test design around demonstrating value rather than demonstrating possibility. One customer shared that they had catalogued more than 150 potential AI use cases across their operation. By the end of a structured prioritisation exercise, they had identified over 200, so the exercise had not narrowed their options at all, but had instead surfaced use cases they had previously missed. More importantly, it had forced a different question to the centre of the conversation: not “what can AI do here?” but “which specific outcome do we want to move, and what is the fastest path to measuring whether we have moved it?” That reframe is what unlocked progress.



What specific outcome do you want to move?



What is the fastest path to measuring that move?

This reframe is how you will unlock progress



“The teams making the consistent leap from pilot to production share one thing: they identify a well-defined process, run a time-bound test with a measurable outcome, act on the result, and move to the next one. They do not attempt to transform the entire operation in a single programme.”

The outcomes from this approach are concrete and repeatable. One organisation reduced the time spent on a high-volume review workflow by around three quarters. Another brought a process that had previously run on a two-week cycle down to two days, with a credible path to hours. Neither result came from a transformation programme; both came from disciplined, sequential delivery against clearly defined process metrics.

RECOMMENDATIONS

Reframe your next pilot as a business value test.

Pick one process with a measurable baseline metric. Define what success looks like in numbers before you start. Set a time boundary. When the test concludes, act on the result and move to the next one. If you cannot define the metric upfront, that is a signal the use case needs more scoping, not more technology.





FINDING 3:

Could your AI business case be wrong by a factor of 10?

Most organisations built their AI investment cases during a period when the technology was genuinely less capable, significantly more expensive to run, and considerably harder to integrate with existing enterprise infrastructure.

Since those business cases were written and approved, the models have advanced substantially, the infrastructure costs have fallen, and the range of processes now within reach has expanded beyond what most teams anticipated. The business cases, however, have largely stayed where they were filed after sign-off.

We heard one leader describe setting a value target that felt genuinely ambitious at the time it was put to the board 18 months earlier. When that organisation worked through what the technology could now realistically deliver against their specific processes, the achievable number was closer to ten times the original figure. The original estimate was neither reckless nor poorly constructed. The technology had simply moved that fast in 18 months.



“Teams may be measuring success against a benchmark that is no longer fit for purpose, declaring programmes successful when they are in fact capturing only a fraction of what is achievable.”

The organisations moving fastest treat the AI business case as a living document reviewed quarterly rather than a capital approval paper filed after sign-off. If your value assessment is more than six months old, stress-test the assumptions against what the tools can actually do now.

RECOMMENDATIONS

Reopen your AI business case.

If it is more than six months old, the assumptions are likely outdated. Run a targeted reassessment against your highest-priority processes using current model capabilities and current infrastructure costs.

The gap between what you approved and what is now achievable may be significant, and your board needs to know that.



FINDING 4:

The change ratio most programmes get backwards

The figure we heard quoted most often across our forums, from leaders at different organisations who had arrived at it entirely independently, is that roughly 70% of the work in a successful AI programme is change management and roughly 30% is the technology. We heard this ratio enough times, across enough different contexts, that we have stopped treating it as an anecdote and started treating it as a finding.

The technology works. That is no longer where the majority of AI programmes fail

or stall. The constraint is adoption: getting people to use the tools, trust the outputs, and genuinely change the way they work. Mandate alone rarely achieves this. Demonstrated value that spreads organically does.

The programmes gaining real traction are not the ones with the most thorough communications plans or the most polished training programmes. They are the ones that identified a critical mass of early adopters, gave them a genuinely useful tool, and allowed adoption to spread through peer influence rather than top-down instruction.





“Senior professionals who have built their careers on expert judgement do not readily cede that judgement to a model. The organisations navigating this well are not asking them to. They are positioning AI as something that handles the high-volume, lower-value elements of the role.”

If your current programme structure allocates the majority of senior attention to technology decisions and treats adoption and process redesign as secondary workstreams, the balance is off, and the delivery record will reflect that.

RECOMMENDATIONS

Audit where your senior programme time is going.

If more than 30% of it is spent on technology decisions rather than adoption, process redesign, and change management, rebalance the structure. Identify a small group of likely early adopters, give them something genuinely useful, and let adoption spread from there rather than pushing it from the top.



FINDING 5:

What scaling organisations do that others don't

The organisations consistently referenced in our forums as examples of genuine, sustained traction share a set of characteristics worth documenting directly.

They stopped asking “are we ready for AI?” at an organisational level and started asking “which process should we change next?” at an operational one. That shift in framing changes the governance model, the team structure, the success metrics, and the overall pace of delivery in ways that compound over time. They measure outcomes at the process level rather than the programme level. They do not declare

an AI programme successful in any aggregate sense. They declare a specific workflow improved by a specific, measurable amount and move to the next one, building a delivery record that sustains board confidence and continued investment far more effectively than programme-level reporting tends to do.

They have a clear, working position on what to build, what to buy, and what to bring in externally. This is not a policy document sitting in a governance framework; it is a practical decision framework their teams can apply to each new initiative without re-escalating to the centre every time.



**Stop asking
“Are we ready
for AI?”**



**Start asking
“Which process
should we
change next?”**

This shift changes everything



“The gap between these organisations and those still in the planning phase is growing. Each delivered use case builds the data, the organisational muscle memory, and the board confidence to move faster on the next one.”

They have also accepted that the underlying technology will keep changing and have structured their capability accordingly. The organisations most exposed to disruption are the ones that have built deep internal expertise around a single vendor's tooling. The ones best positioned for the next wave are the ones that have built around outcomes and processes, with the flexibility to bring in different underlying technology as it matures.

RECOMMENDATIONS

Stop measuring at programme level.

Define a single process metric for your next initiative, deliver against it, record it, and move on. Build a delivery register that tracks outcomes at the use case level.

That register is what sustains board confidence and investment through the difficult middle phase of a multi-year programme.



FINDING 6:

The AI delivery model that works. And the one that doesn't.

The delivery model we see producing consistent results brings three distinct capabilities together: a technology partner who keeps pace with what the models can actually do and applies that to real commercial processes; a delivery partner who understands change management and sector context; and a hyperscaler who provides the

infrastructure, security, and enterprise integration capability that commercial organisations require before putting production workloads on a platform.

None of these capabilities needs to be built entirely in-house, and the organisations attempting to develop all three internally are, without exception, the slowest moving in our experience.





“The model that fails consistently is the one that attempts to centralise all decisions and all capability in a single internal AI team. The people involved are rarely the problem.”

The model that fails consistently is the one that attempts to centralise all decisions and all capability in a single internal AI team. The people involved are rarely the problem. The pace of change in the underlying technology simply outstrips what any single team can maintain across the full breadth of a commercial AI portfolio while simultaneously delivering live programmes.

RECOMMENDATIONS

Map your capability gaps honestly.

Identify which of the three delivery components: technology currency, change and delivery expertise, and infrastructure at scale, your organisation can genuinely provide and which it cannot. Close the gaps through partnerships rather than attempting to build everything internally. The organisations moving fastest have been clear about what they need to own and what they do not.





LEARN MORE
VERSION1.COM

info@version1.com

If you would like to discuss where your organisation sits against any of these findings, Version 1 works with commercial organisations on exactly these questions.

Version 1 is a global leader in AI-enabled technology services, founded in Ireland in 1996. With over 4,500 employees globally and revenues exceeding €400 million, the company has a distinctive approach to business transformation. Its ability to co-create with customers delivers extraordinary outcomes by combining extensive industry knowledge and deep technology expertise. Version 1's vision is to be the catalyst for a better world where technology complements human innovation to solve the most complex challenges facing society.

