Insight ON: AI

Still Piloting?

Here's How to Get AI into Production Faster

InsightAI

Most AI projects promise value but never deliver

Most technology leaders have been faced with this situation: a pilot is running, or maybe several are running in parallel, and the results are genuinely promising. The demos land well, the business case is solid, and there's real appetite to move forward. But somehow, months later, the programme is still not in production.

Governance sign-off takes longer than expected, the legacy system integration proves more complex than the initial scoping suggested, and cost projections increase as the scope becomes clearer. Important stakeholders become nervous about the unknown elements of new technology and become reluctant to move ahead. Gradually, the project's original momentum has drained away and there's not enough evidence that the benefits of full production would outweigh the effort.

This experience is extremely common. In fact, research from Insight's recent report **Navigating the Autonomous AI Trust Barrier: From Hesitation to High Velocity** found that 71% of organisations are still stuck in pilot or proof-of-concept stages of AI-driven transformation. With such a large number of firms investing in AI, it's clear that there is ambition and early results, but many organisations are stuck in limbo between a controlled experiment and an effective, operational asset that's adding value to the business – or evolving how it operates altogether.

We call this the **AI Implementation Gap**. It's the gap between companies who are still planning how to build or embed the right AI product, and those who have successfully deployed it.

For business leaders, harnessing AI to improve their organisation's efficiency and processes is one of today's defining operational challenges. And they are taking it seriously, including making structural changes to their tech infrastructure to make sure they're ready. In **The Digital Sovereignty Trilemma** we discovered that 85% of European organisations are already evaluating or deploying dedicated on-premises infrastructure specifically for AI – far more than have bridged the AI Implementation Gap and are using AI at scale in their businesses.

But even with investment, closing the AI Implementation Gap is not an easy task, or one with a one-size-fits-all approach. Businesses need to move ahead with their strongest use cases and find real value from AI. But just speeding up won't work if the use cases aren't right and the fundamentals aren't in place. We've found it's critical to think carefully about the structure of the journey and move ahead rapidly, but with a clear, flexible plan for generating long-term value.

A 10x improvement: how we closed our own AI Implementation Gap

At Insight AI, we needed a better way to turn completed project work into client-facing case studies – one of those important but not mission-critical tasks that many firms accept as slow, manual and inefficient.

Producing case studies seems simple in principle, but in practice, the process had no automation, no trigger to start the work, and no consistency in how it got done or the final results.

The project knowledge lived within our teams, who had to spend weeks pulling together information – while juggling new projects and priorities – before passing it to marketing for copywriting, editing, and design. By the time a case study was ready, the sales opportunity it was meant to support had often already moved on.

So, we built a solution for ourselves. AURA is an AI-powered knowledge platform that we use to make generating case studies automatic and easy. It means our sales teams have easy access to our comprehensive bank of case studies, and can quickly and easily use AI to craft narratives in local languages to show prospective clients our most relevant expertise.

When a project completes, the Delivery Lead receives an automatic prompt to fill in a structured questionnaire. That input feeds directly into AURA, which generates a polished draft that has proper governance baked in from the start.

Anonymisation-aware logic strips client and personal identifiers. Azure AD authentication, role-based access, and PII guardrails are built in by default rather than bolted on after. And because we built in operational visibility from the start, we can see how the platform is being used and where it needs to go next.

The result: case study generation is now 10× faster than the manual process.

We didn't wait for AURA to be perfect before we launched it. We built the right structure, applied the right governance, launched it, and kept improving. Because the foundations were right, the scope could expand. As it keeps evolving, the platform is now growing to support more advanced grouped case study creation, which means we can combine multiple relevant projects into thematic or whole-industry stories.

This is what closing the AI Implementation Gap actually looks like: it's not a single moment of transformation, but a structured journey with the discipline to keep moving forward by layering improvements onto a product with strong foundations.

The AI Implementation Gap is often a structural problem, not a technology problem

So why do so many pilots stall? In our experience, it usually comes down to four recurring blockers, and none of them are about the limitations of the technology itself.

Analysis paralysis

Most organisations have more potential AI use cases than they have capacity to pursue. Without a clear framework for deciding which ones are genuinely worth scaling, effort gets spread across too many experiments and nothing reaches production. When this happens, the number of pilots keeps expanding but the production deployments don't. Organisations need to know where the value is to prioritise what to work on – assessing this at the outset will make sure that the direction of travel is clear.

Cost creep

When designed poorly, AI projects can often look affordable in the pilot phase and considerably more expensive at scale, especially when they aren't properly stress-tested before investment. To prevent this, organisations need to work with experienced, fully transparent partners, and take due diligence seriously: **The Digital Sovereignty Trilemma** found that 56% of organisations don't even conduct Total Cost of Ownership assessments before deciding where to place significant workloads. Without radical transparency in the planning stages, it's easy to over-commit before the true cost picture becomes clear.

Governance and compliance anxiety

Particularly for organisations operating in EMEA, regulatory complexity creates genuine hesitation. Navigating GDPR, DORA and the EU AI Act is not simple and the consequences of getting it wrong matter. Anxiety about staying compliant in a rapidly changing regulatory landscape sits behind a lot of the conversations about projects that have potential but don't progress.

Over-investment before validation

At the other end of the spectrum, some AI-projects are built for scale before the concept has been properly stress-tested. When this happens, teams invest months in architecture and infrastructure for a use case that hasn't yet survived contact with real-world complexity. By the time it's clear the project needs to pivot, the sunk cost makes it feel impossible.

The structured path from pilot to production

Insight AI's answer to avoiding these blockers and closing the AI Implementation Gap is to find "Velocity Without Risk". This means charting a middle path between "just move fast and fix it all later" and "build out at scale and launch when it's all finished".

Using an agile approach to continually test and iterate is essential, but it's only part of the picture. Organisations should only feel comfortable moving ahead quickly when projects have clear, well-sequenced milestones and early proof that the outcome will provide something genuinely valuable – because building the wrong thing quickly isn't the answer.

This means working with a partner who can de-risk investments by making sure that budget is only allocated to high-yield, production-ready AI use cases. They should know which decisions need to be made early, which investments to front-load, and which questions will become blockers at the worst possible moment if they're left unanswered.

When we practice "Velocity Without Risk" for our clients, there are three things that mean they either close the AI Implementation Gap, or save their budget from projects that never go anywhere.

1. Fail-fast validation

Before any significant investment in build, we identify the kill points: the technical, commercial, or governance questions that would stop a use case from reaching production. The goal is to be wrong cheaply and quickly rather than expensively and late. We de-risk client investments by not developing projects that aren't feasible or won't create value, so budget only goes towards use cases that will be high-yield and production-ready.

2. Concept to Code in 5 Days

We use rapid prototyping to validate technical feasibility before committing to scale. This is a specific methodology that lets clients see something real, testable, and evidence-based within days, before deciding to go ahead. It changes the conversation from 'could this work?' to 'here's what it does.'

3. Milestone-based roadmaps

To avoid AI development becoming an opaque "black box", there needs to be radical transparency at every stage. This means all project elements have clear owners, clear stages, and clear exit criteria. Clients know exactly where they are and what the next step requires. This is what makes speed sustainable: the decisions that slow projects down are almost always the ones that weren't made clearly and transparently at the start.

Together, these help projects move at a rapid, sustainable pace, but they're also ways of making sure that when organisations do move fast, it's in the right direction.

Creating a production-grade MVP in 178 days: AMP's story

When AMP, a leader in sustainable energy, needed to build a platform to streamline end-to-end biomethane sustainability management, they chose to do it differently. Working with Insight AI, they made a deliberate decision to test whether Agentic AI embedded across the entire delivery lifecycle – development, QA, design, business analysis, and delivery management – could outperform classical delivery on a real commercial engagement, under real commercial pressure.

The approach delivered, with the results exceeding every expectation. 178 days after the project started, AMP had a fully functional, market-first platform in production. The team had solved a complex operational and compliance problem: the new digital platform streamlined their feedstock sustainability in biogas production, and ensured the traceability they needed for regulatory compliance.

The operating principle that made those numbers possible was straightforward: agents draft, humans decide. AI handled generation, synthesis, and execution; senior engineers retained responsibility for architecture, compliance, security, and the decisions requiring genuine judgment. The development team worked in one-week sprints – a velocity that was made possible because AI was absorbing the volume that would usually require significantly more people and time. Where classical delivery might take weeks to validate a design direction, functional prototypes could be turned around in one or two days.

"We didn't just use AI to write code – we built a system where AI operates as an autonomous team member with clear rules. It's not a chatbot – it's a guided team member."

Artiom Leontiev - Senior Software Developer, Insight AI

6.6x faster delivery	**60x** fewer bugs	**178** days to MVP
9x more code/ month	NPS **10** highest possible	**5.7x** dev speed increase

For the people working closest to the delivery, the pace and the approach were genuinely new territory, and the results made the case better than any upfront argument could have.

"I was a bit sceptical at the beginning, but after finding new ways of working every day that produced better and greater results, I became a believer."

Stefan Kolev -Business Analyst, Insight AI

Closing the AI Implementation Gap

The 71% of organisations stuck between pilot and production aren't there because they lack ambition or capability. Most have both. What the evidence consistently shows is a more structural problem: the lack of a clear path which can take organisations from a promising AI experiment to an operational system that unlocks significant value for the business – with the validation gates, the governance, the transparency, and the continuity of partnership actually required to achieve it.

That's what both AI projects in this piece have in common, despite being very different. We built AURA using our highly structured, fail-fast approach because, like our clients, we want to accelerate our internal processes and free up staff and budget for what really matters. AMP was a client engagement where the brief was ambitious, there was appetite to experiment and the team became believers because the structure held under real pressure.

For most businesses, closing the AI Implementation Gap doesn't require a huge leap of faith. It requires a partner who treats the path from experiment to production as something to be engineered rather than assumed, and who stays accountable at every stage – from strategy, to build, to the point where the system is running and the work of sustaining it begins. That's where ambition delivers.

If you're also navigating who builds your AI and what full-stack partnership looks like as EU AI Act obligations become enforceable on 2nd August 2026, we'll be addressing these issues in our next piece **Who's Running Your AI? What 'End-to-End' Actually Means**

nl.insight.com/ai