Moving beyond the hype: How to scale AI successfully

thoughtworks

Design. Engineering. AI.

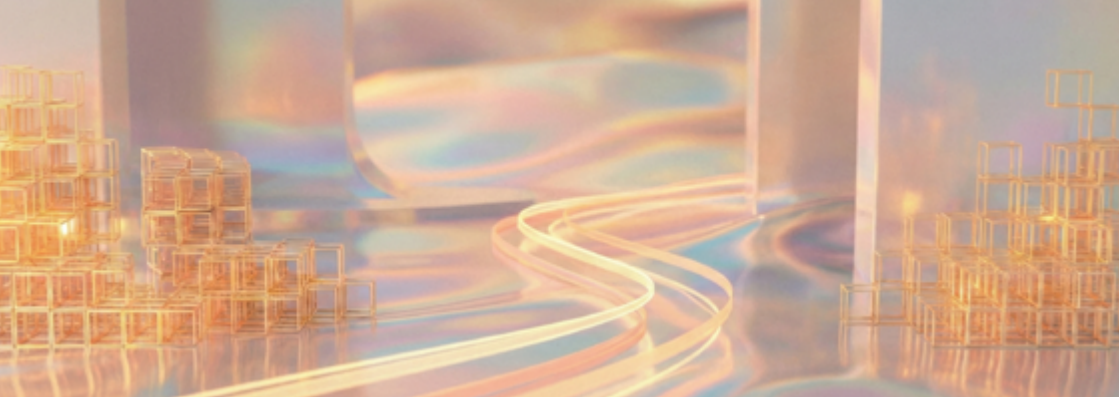

Organizations everywhere have embraced AI. Now, the challenge is scaling it into production to unlock value.

According to research from Gartner, the global AI software market is set to reach a total value of $297.9 billion by 2027. As of 2025, around 39% of worldwide organizations are now at the experimentation stage of Gartner's AI adoption curve.

Firms understand the potential of AI and how it could help them deliver transformational customer and business impact. They're actively experimenting with it and creating Proofs of Concept (POCs) to help them see exactly where AI can (and can't) deliver the value for their organization. But for most, that's where the real challenges begin.

Scaling AI POCs into production isn't easy. Even if your POCs have demonstrated high value and strategic impact, initiatives can grind to a halt quickly and fail to deliver their potential value.

There's a huge range of reasons why that happens. In this paper, we'll examine the most common barriers to scaling AI, and provide actionable advice to help you break them down and build AI readiness across your organization.

Introducing the FOREST framework

In most cases, the barriers to AI scaling don't cause issues because they're insurmountable — it's because their impact on the success of AI initiatives is significantly underestimated. Because there can be so many of them, teams often struggle to see the big picture.

To help organizations better understand the impact of each barrier and prepare for scaling AI, Thoughtworks created the FOREST framework for AI readiness. The framework breaks down the essential components for AI readiness into six areas:

- Foundational architecture
- Operating model
- Readiness of data
- Experience for humans + AI
- Strategic alignment
- Trustworthy AI

By building AI readiness in these six areas, organizations can remove the barriers to scaling AI and transition their POCs into production with relative ease.

Foundational architecture

Without the right foundational architecture, your AI plans simply can't progress. If essential elements of your AI architecture are missing — such as enabling platforms or updated processes — you'll quickly start encountering major barriers to AI scaling.

Here are three of the most common barriers encountered in this area, along with remediating actions and strategies that can help you avoid and overcome them:

Barrier #1: Inefficiency in model development, training and deployment

When core AI activities like model development, training and deployment are inefficient, error-prone and inconsistent, teams' patience quickly runs thin. This leads to lower engagement, slow progress, and ultimately, the abandonment of high-potential AI initiatives.

These processes can be challenging to establish, especially in teams with very clear and structured workflows. The missing piece for many organizations is a purpose-built platform to support them.

The right platform can be tightly integrated with existing systems to help you establish a robust foundation for AI training and development processes, and provide exceptional developer experiences by design.

Instead of having to slowly develop their own approach to model development, training and deployment, teams can engage with your platform to guide their journey. This makes the experience faster and more efficient for them and helps ensure consistency and adherence to AI governance standards.

Barrier #2: Slow iteration and problem-solving post-POC

AI models and use cases rarely deliver perfect outputs from their first iteration. Feedback, retraining and rework are all essential to get your POCs ready for scaling into production.

Often, processes around model retraining, evaluation, testing, and automated deployment are an afterthought for organizations experimenting with AI; they're only formed at the point where they're needed. Naturally, this stops many POCs in their tracks as teams establish how they're going to progress from the first iteration.

By taking the time to establish those processes upfront and build clear pipelines for feedback and model iteration, the optimization process can start straight after the first iteration — and continue long after your POC is in production.

Barrier #3: Slow decision-making and collaboration

AI initiatives demand continuous collaboration between multiple technical and domain teams, and autonomy for key decision-makers in those teams. It sounds paradoxical, which makes it a key sticking point for many organizations.

The balance between collaboration and autonomous decision-making must be made at the architectural level. The right architecture can ensure that all relevant stakeholders have the insight and autonomy to make informed decisions quickly, and enable continuous, transparent collaboration between stakeholders, domains and agents.

Operating model

AI development and AI-enabled operations are significant shifts in how teams work. If your operating model doesn't evolve alongside your technology, then your processes and AI aspirations quickly become misaligned, making it difficult to realize AI's full potential.

The most common barriers to AI success in this area are:

Barrier #1: Imbalances between specialization and cross-functional collaboration

Cross-functional collaboration is essential for the success of your AI initiatives. Technology and business teams must work closely together to ensure your chosen POCs are technically feasible and aligned with business goals.

Both teams' specializations are essential, but if their involvement in your projects is imbalanced, your outcomes will be skewed toward one team's area of expertise.

To overcome this, we recommend building bespoke cross-functional teams to support AI development and scaling. Both technical and business perspectives need a seat at the table, but they must also continuously collaborate. Building a team with the different perspectives and expertise required to deliver successful AI use cases ensures the balance is correct.

Barrier #2: Dependencies between distributed teams

When you have multiple teams working toward the same goal, each one needs a clearly defined role in the process. Without clarity on roles, dependencies slow innovation and iteration.

For the process to work efficiently, every team must understand the scope of their role, what's required of them, and when they

have to deliver it. Only by setting and upholding those roles and responsibilities can you ensure continuous cross-team delivery.

Barrier #3: Centralized management models

When orchestrating initiatives between multiple teams, it's natural to lean on centralized management models. But in practice, having centralized decision-makers responsible for coordinating every activity between teams can lead to project slowdowns.

Instead, look to evolve your operating model to empower each team to make as many of their own decisions as possible, within the boundaries of their role. It's a delicate balance, but decentralization is essential to enable continuous AI iteration, delivery and innovation.

Is your data ready for AI?

With AI, what you get out completely depends on what you put in. So, if your data isn't AI ready, neither is your organization.

The top barriers encountered in AI data readiness are:

Barrier #1: Data isn't discoverable, reusable or optimized for AI-driven applications

When building AI use cases for the first time, organizations need to take a careful look at their data strategy. If essential data assets aren't easily accessible and useable for AI use cases, that needs to be addressed before the AI development process begins.

Emerging approaches like Data as a Product can be hugely helpful here. Data as a Product enables domains to build data products for specific use cases, so they're AI ready by design.

Then, it makes those products highly visible and accessible to teams across the organization, so they can be easily reused for further AI use cases.

Barrier #2: Low data quality

As AI has matured into a mainstream technology, nobody has been able to escape the phrase "garbage in, garbage out." It's a cliché, but it holds true. If you train AI using low-quality data, it leads to unsatisfactory and even damaging outcomes from POC models and use cases.

To become AI ready, data must be supported by frameworks that ensure it's of a high enough quality to deliver top-notch outcomes. It must also be structured in the right way, as unstructured data is even harder to maintain at a consistently high quality.

Barrier #3: Difficulty translating AI outputs into the right actions

AI enables human decision-making. But, many decisions that organizations want to support with AI demand deep and contextualized data. Without that context, it's very hard to have confidence that the decisions you make are right for your business or customers.

In many cases, what's missing is a purpose-built semantic layer, such as a knowledge graph, to connect data and provide vital context. This layer can help bridge the gap between AI outputs and human actions, and helps organizations ensure that AI outputs are uniquely relevant to them.

Experiences for humans + AI

A common misconception about AI is that its purpose is to replace some of the roles of humans in key processes. In reality, it must augment and improve human experiences. And it's a two-way process — humans must work in ways that get the most out of AI, and AI must be engineered to support humans effectively.

The barriers most commonly encountered here are:

Barrier #1: Humans don't like using AI solutions

AI solutions can be technically valuable and sound, but if humans don't engage with them, their value will be limited, and motivation to scale them into production will be low.

To avoid this, AI models and solutions should be designed with a user-centric vision from the outset. You should set a clear definition of expectations around human-AI experiences and use it as your north star throughout development and scaling.

As human needs change, your AI solutions must change alongside them. User feedback and research must be integrated into the iteration process, so models remain aligned with what the humans using them want to get from them.

Barrier #2: Low user value

The same principles around experience expectations must also be applied to user value. Everyone involved in AI development and delivery should have a clear understanding of your model or solution's intended customer value.

If your initial iteration fails to deliver that value, teams must iterate toward it. Throughout the iteration process, user feedback and research should again be used to ensure the value you're trying to deliver is still relevant.

If user needs change and your solution doesn't change with them, you may end up with an AI solution that delivers very limited value. These solutions shouldn't be scaled into production, but need to be reworked until they're fully aligned with what users require.

Barrier #3: Difficulty engaging with the AI development process

If a specific team or set of stakeholders stops engaging with the AI development progress, everything can grind to a halt. One of the biggest reasons why that happens is that the AI development process is too difficult to engage with.

If you include diverse stakeholders and build cross-functional teams with different backgrounds and skillsets, the project needs to speak everyone's language. You may start with a utopian view of your cross-functional project, but things can quickly slip into the technical, limiting business stakeholders' ability to engage and provide input.

Continuous alignment between all relevant teams and individuals is essential for ensuring that doesn't happen.

Strategic alignment

To deliver high business value, AI use cases must be closely aligned with your strategic goals.

The most common barriers organizations encounter here are:

Barrier #1: No clear value articulation for POCs

If AI POCs deliver low value, they get dropped, and never reach the point of scaling. But low initial value doesn't mean that the

underlying concept behind that use case isn't capable of delivering value.

With every POC, it's important to conduct a detailed post-mortem before moving on. Assess whether the lack of value realization was down to technical issues, processes, or a poor alignment with business objectives.

If you spot exactly how the POC fell out of alignment with your business goals, you may be able to easily tweak and test it again. Suddenly, an idea you were ready to abandon could become one of your most valuable AI use cases.

Barrier #2: Inability to track and prove the value of AI POCs

You can't prove what you can't measure. So, one of the keys to AI readiness is choosing and tracking the right measures of success for your POCs. It's essential that the metrics you choose are relevant to your business. Start with an idea of what you want to achieve at a strategic level, then work back toward specific use case success metrics to build a true view of a POC's potential business value.

Barrier #3: Teams' inability to get the most from AI tools

Another big reason why a POC may not appear to deliver on its hypothesized value is that teams don't understand how to engage with it or integrate it into their workflows.

Upskilling and enablement are an essential part of AI readiness. It's not enough that your technology, processes and data are AI ready — your people must be prepared for it, too.

Leaders must determine exactly how new AI tools should be integrated into existing workflows, or how workflows should be changed to accommodate AI. Then, training and development

should be offered to ensure everyone fully understands how to get the most from their AI tools.

Trusted AI

If you can't trust the outputs of your AI solutions, you can't scale them. So, establishing robust governance frameworks is an essential element of AI readiness.

The barriers many organizations run into in this area are:

Barrier #1: Low or no visibility of what causes negative AI outcomes

If an AI POC starts generating or contributing to negative outcomes and you can't pinpoint why that's happened, you can't take the right steps to resolve it. So, in most cases, the POCs are abandoned or completely reworked, and never scaled into production.

To prevent this, you need to ensure you have the right evaluation frameworks and safety guardrails to assess model performance, mitigate risks and ensure ethical AI deployment. Just like with value metrics, visibility and traceability are everything.

Barrier #2: Imbalances between automated governance and human input

Humans and automated capabilities need to work together to monitor AI outputs and act quickly when needed. If responsibilities aren't clearly defined between the two, that process can break down. If humans aren't looped in at the right time, the entire "human in the loop" concept can break down.

Take the time to find the perfect balance between automated monitoring and alerting mechanisms, and human decision points. Then, make them very clear to everyone, so people know exactly where and when their input will be required.

Barrier #3: Governance and responsible AI are an afterthought

In many cases, governance and responsible AI aren't prioritized during the creation of AI POCs, because teams believe they're only relevant when AI outputs are used in a real business or customer context.

Firstly, this gives a false impression of how valuable the POC in question could be. If the right governance frameworks and rules haven't been applied, the POC won't reflect what it could do for people in a live production context. Secondly, this approach makes it extremely difficult to retroactively apply governance frameworks and rules to your new solution.

Governance and responsible AI should be built into all POCs from their inception. Then, when you find one you want to scale, you can do so with confidence, immediately.

Bringing everything together for scalable AI success

Clearly, there are many potential barriers that organizations can encounter when they want to scale AI POCs into production. But, by taking the right actions across the six dimensions of our FOREST framework, you can break down the barriers to AI scaling and realize its full value. Only by doing that can you become truly AI ready.

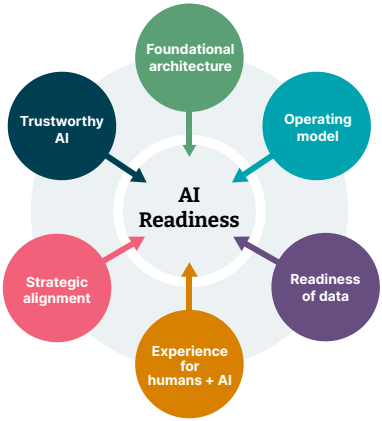

AI readiness FOREST framework

That means:

- Core **change and transformation**, including strategy, organization and governance evolution
- **Use case discovery**, inception and delivery
- The construction of robust and relevant **platform and architecture foundations**

The impact of our approach

Our proven frameworks maximize AI's impact on an organization and the people it serves. Every organization has different AI goals and ambitions, but we focus on five specific areas of impact that we know can have a transformational effect on any business.

1. Driving growth

- 79.5% growth in average basket size delivered for a fashion and design house

- 1 million users onboarded to a new AI-powered digital service in just 100 days for a leading loyalty program in Singapore

2. Delivering exceptional customer, employee and developer experiences

- 75% agent efficiency growth and a 20% improvement in customer satisfaction delivered by creating new AI assistants at a leading bank

3. Cutting costs

- $15.1m in cost reductions delivered within the first months for a major international airline through optimizing maintenance operations with data-driven decision-making.
- $80k annual savings delivered through process optimization and efficiencies at a global mobility leader.

4. Accelerating processes and delivery

- 20% time savings for obtaining customer and product insights at a global bank
- 75% increase in operational efficiency at a pharmaceutical company
- 2x to 5x speed gains from parallelizing tasks and optimizing code at a large pharmaceutical company

5. Managing and mitigating risk

- Thoughtworks has partnered with the United Nations to lead and promote responsible tech practices for emerging technologies such as AI and help organizations worldwide avoid negative impacts from their use.

Build AI readiness and start scaling high-value POCs into production today

No matter which barriers to AI scaling you're facing, Thoughtworks can help you overcome them. With decades of AI experience and expertise, we know exactly what it takes to become AI ready, and have proven frameworks and tools to help you achieve it.

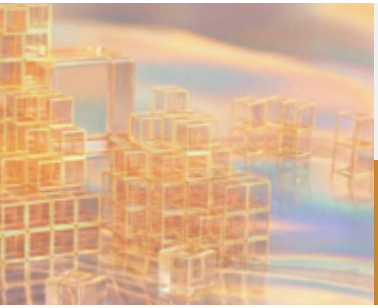

Find out more about our AI readiness and AI scaling services

Visit our website or talk to us to discover how we could help you start building high-value POCs and scale them into production.

Authors

Danilo Sato
Global VP of AI
at Thoughtworks

Danilo Sato is the Global VP of AI at Thoughtworks. He leads the global team that defines Thoughtworks' services, assets, accelerators, go-to-market and partnerships strategy to help organizations get AI-ready. He advises executive leaders on topics ranging from strategy and governance to building the products and platforms that bring the strategy to life. A former member of Thoughtworks' Office of the CTO, he developed unique approaches to Data + AI, such as Data Mesh and CD4ML. Author of "DevOps in Practice: Reliable and Automated Software Delivery" and an experienced international conference speaker, Danilo was named by DataIQ as one of UK's 100 most influential people in data in 2022, 2023 and 2024, and played a key role in ITV's award-winning Data Mesh implementation.

Tiankai Feng
Director for Data & AI Strategy
at Thoughtworks

He is the author of "Humanizing Data Strategy" & "Humanizing AI Strategy." With 12+ years experience in Data Analytics, Data Governance, Data Strategy and AI Transformations, he found a passion for the human side of technology: how to collaborate, communicate and be creative around Data + AI. Tiankai is a frequent public speaker at renown conferences like TEDx, Big Data London or data:unplugged, and is executive board member at DAMA Germany.

We are a global technology consultancy that delivers extraordinary impact by blending design, engineering and AI expertise.

For over 30 years, our culture of innovation and technological excellence has helped clients strengthen their enterprise systems, scale with agility and create seamless digital experiences.

We're dedicated to solving our clients' most critical challenges, combining AI and human ingenuity to turn their ambitious ideas into reality.

thoughtworks.com

/thoughtworks

Design. Engineering. AI.

Last updated August 2025