

Mind the Gap: Bridging from Sandbox to Scale

November 2025

💻🧠💡 CTO Craft x Ten10



Table of Contents

Introduction	3
Who We Surveyed	4
The AI "Pilot Purgatory" - Understanding the Gap	5
Architecting for Reality – Technical & Operational Hurdles	7
The Human Element – Talent & Teams	11
Value, Vision & The Future	13
Insights - Ten10 & The Scale Factory	15
Insights - CTO Craft	16
Methodology	17

Introduction

Welcome to Mind the Gap: Bridging from Sandbox to Scale

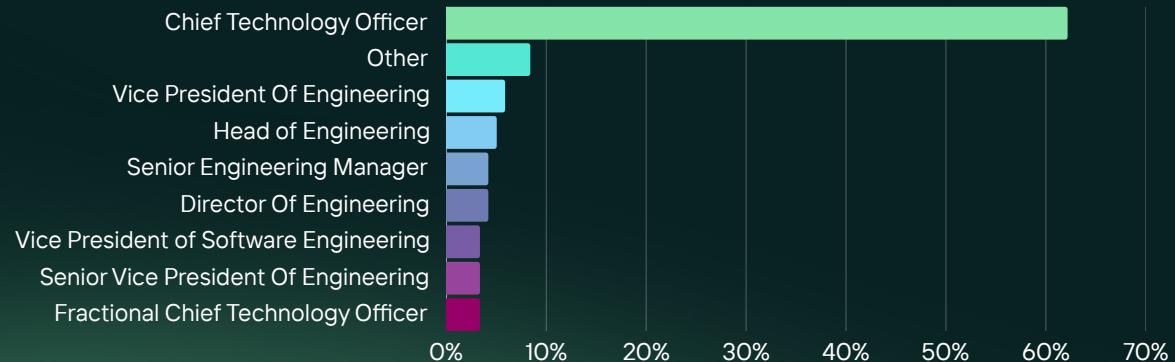
For senior technology leaders, the critical question has shifted from if AI can work to how to transition promising Proofs of Concept (PoCs) from the controlled "sandbox" into resilient, value-generating systems.

While investment in AI, particularly Generative AI (GenAI) and latterly Agentic AI, continues to rise, the ability to successfully operationalise these technologies remains the critical challenge. The journey from technical feasibility to scalable, production-ready impact is often derailed by complex, non-technical and organisational hurdles.

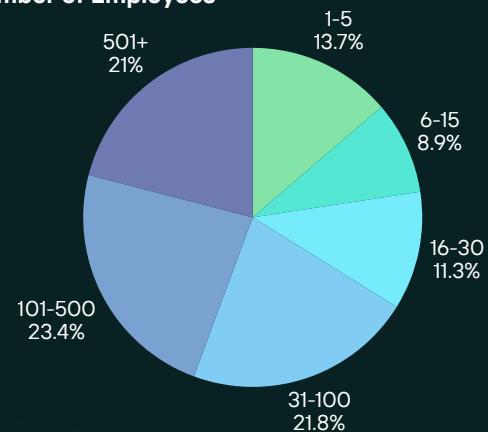
This report, based on an industry survey of technology leaders, provides a comprehensive view of these challenges and offers actionable intelligence to proactively anticipate and tackle the common, costly obstacles between a successful trial and scaled production.

Who we surveyed

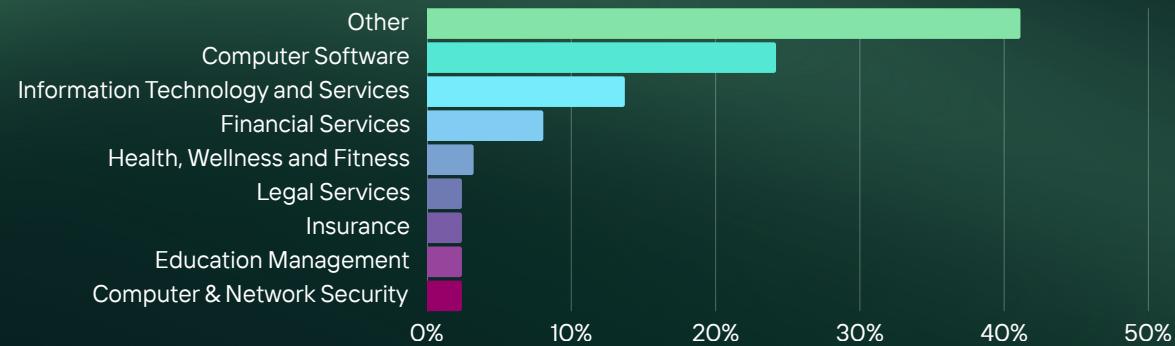
Job Title



Number of Employees



Industry Sector



Location

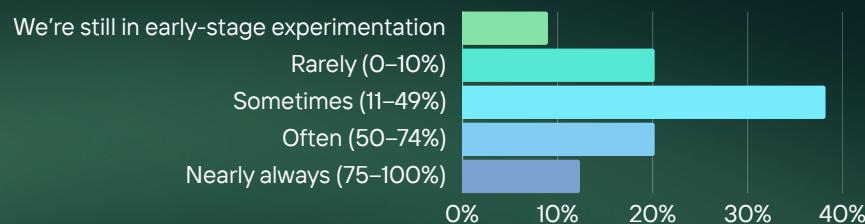


The AI "Pilot Purgatory" - Understanding the Gap

The AI Scale Chasm: A "Coin Flip" for Production Success

The data reveals a stark reality: most organisations are struggling to convert AI/ML experiments into tangible business impact. The successful transition of an AI PoC into a production system that delivers measurable business value is essentially "flipping a coin".

How often do your AI proof of concepts make it past experimentation and into production, where they deliver measurable business impact?



Key Findings on Success Rates:

- Sub-50% Success Rate:** 67% of organisations admit to having a sub-50% success rate for AI PoCs moving into production and delivering measurable business impact.
- Consistent Success is Rare:** Only 12% of leaders report achieving consistent success ("Nearly always"). This underscores the scarcity of repeatable, true AI scale across the industry.

The Non-Technical Bottleneck: The Organisational Wall

When AI projects fail to scale, the root cause is frequently non-technical, acting as the ultimate speed limit for innovation. The most critical factor derailing promising PoCs is a failure of organisational alignment.

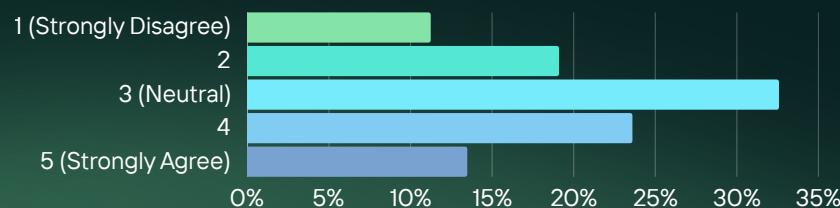
Top Non-Technical Barriers:

- Misaligned Stakeholder Expectations (The Pre-Launch Problem):** This is the single greatest non-technical factor cited in our survey. It centres on the failure to transition from a "science project" mindset to a "business product" mindset. Business leaders often see a successful PoC (e.g., 90% accuracy) as instantly deployable, ignoring the expensive, non-glamorous investment needed for MLOps, security, and integration. Data science teams compound this by sometimes over-promising short-term impact.
- Organisational Inertia and Prioritisation Shifts (The Scale-Up Problem):** Once a PoC is technically proven, internal bureaucracy often halts momentum. The primary friction points are:
 - Lack of Dedicated Product Ownership:** Many projects die because they are not treated as critical business capabilities requiring a dedicated product owner, maintenance budget, and integration roadmap.
 - Funding Friction:** Funding is available for exploration but often disappears for the less glamorous, expensive work of MLOps platform building and production hardening.

The Production-Readiness Gap: Why AI Projects are 'Flying Blind'

The survey responses reveal a concerning lack of foresight in the initial stages of AI development, creating significant technical debt that ultimately hinders scaling efforts. Many leaders are allowing Proofs of Concept (PoCs) to proceed without the necessary foundation for production:

To what extent do you agree with this statement: "Most of our AI prototypes are developed without 'production-readiness' in mind from day one, inadvertently creating significant technical debt that hinders future scaling."



- **Flying Blind:** A substantial **37%** of respondents effectively admit they are flying blind by agreeing or strongly agreeing that their AI prototypes are developed *without* production-readiness in mind from day one.
- **Lack of Confidence:** Conversely, only **30%** of organisations express confidence in their approach, actively disagreeing or strongly disagreeing with the statement, suggesting that fewer than one-third of tech leaders are consistently prioritising scale and stability from the outset.

This imbalance strongly suggests that organisations are struggling to enforce the discipline required for successful operationalisation. Developing models without a robust plan for MLOps, security, and integration guarantees a painful and costly re-engineering effort when it's time to transition from the sandbox to the live environment.



Flying Blind on Production Requirements

The survey reveals 37% of organisations develop AI prototypes without production-readiness in mind. While teams should avoid over-engineering at the experiment phase because there's a chance the project might be thrown away, be prepared for what you'll need to take it to production if successful: secure data access, lifecycle maintenance (development, deployment, monitoring, making updates), regulatory compliance, and integration with other systems. Too many organisations reach PoC success only to realise they haven't planned for any of this, which explains why two-thirds struggle to reach production.

Mike Mead
Chief Product & Technology Officer
The Scale Factory



Architecting for Reality – Technical & Operational Hurdles

The Technical Bottleneck: The Security, Data, and Integration Triad

When architecting AI for scale, the challenges shift from data science to foundational systems engineering.

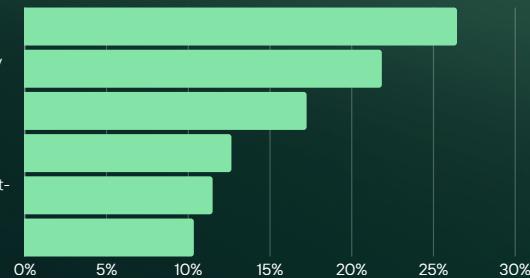
The most significant technical bottlenecks are not MLOps maturity itself, but the core non-functional requirements of software architecture: Security, Data, and System Integration. These three areas account for over 65% of the stated technical bottlenecks.

Top Technical Hurdles to Scaling:

1. Security, Privacy, and Compliance (26%): Once a model leaves the sandbox, critical architecture requirements around data protection and regulatory adherence become a significant, costly obstacle, often because they weren't addressed during the prototyping phase.

When architecting AI for scale on a cloud platform like AWS, which of the following represents your most significant technical bottleneck in moving from proof of concept to production? (Please select one)

- Ensuring robust security, privacy, and compliance throughout the entire AI lifecycle
- Advanced data management, governance, and quality for AI workloads
- Seamless integration of AI models into existing (more complex) systems
- None of the above
- Optimising cloud infrastructure and resources for cost-efficiency and performance at scale
- Establishing mature MLOps pipelines for continuous integration, delivery, and monitoring



2. Advanced Data Management, Governance, and Quality (21%): This validates the industry adage: "It's a data problem, not a model problem." Scaling AI demands reliable, high-quality data pipelines and rigorous governance to prevent model drift and ensure reliable performance.

3. Seamless Integration into Existing Systems (17%): An isolated model is a valueless model. Embedding the AI solution into the existing organisational ecosystem (e.g., microservices, legacy applications) is a major engineering hurdle. This reinforces the failure to design with "production-readiness" in mind.

Data Strategy: The Biggest Blocker

A lack of data strategy, or one that's simply not fit for purpose, is consistently the biggest blocker we see. Many businesses might have had a data warehouse project on the back burner. Now AI demands clean, accessible, well-governed data, forcing organisations to tackle problems they've ignored for years. Data may be poor quality, not in the right format, inaccessible, or simply not where it needs to be.

Solving your data problem delivers value far beyond AI: better business intelligence, compliance reporting, and improved operational decision making across your entire organisation.

Mike Mead
Chief Product & Technology Officer
The Scale Factory



The AI Pilot Purgatory

Architecting for Reality

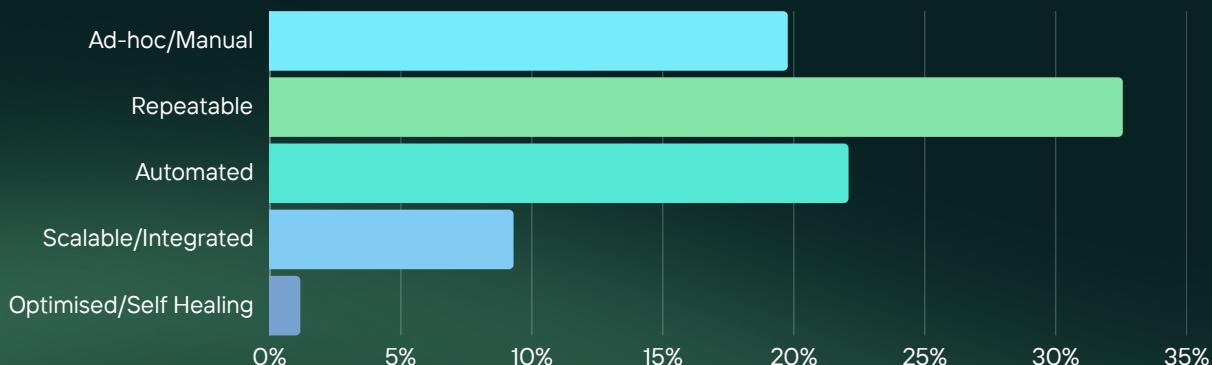
The Human Element

Value, Vision & The Future

The MLOps Maturity Hump

The core of the scaling problem is captured in MLOps maturity. **74%** of organisations are stuck in Stages 1, 2, or 3 (Manual, Repeatable, or only basic Automated MLOps). This means nearly three-quarters of the industry still relies on manual or partially automated processes for model deployment, monitoring, and retraining.

Where would you place your organisation on the MLOps maturity curve in terms of automating the deployment, monitoring, and retraining of AI models?



MLOps: Haven't We Solved This Already?

Three quarters of organisations remain stuck in early MLOps maturity stages with manual processes. When organisations describe their MLOps challenges, my first reaction is: haven't we solved this already with DevOps?

Over a decade ago, we faced identical challenges: manual deployments, ad-hoc monitoring, and gaps between development and production. We solved it through automation, collaboration, and standardised tooling. MLOps requires the same cultural shift that made DevOps successful. We shouldn't need to learn this lesson twice.

Mike Mead
Chief Product & Technology Officer
The Scale Factory



This lack of maturity directly limits speed, reliability, and security. Crucially, the organisational root cause of this low MLOps maturity is poor cross-functional collaboration. Over 59% of organisations rate their collaboration between Data Science, Engineering (DevOps/MLOps), and Business teams as Neutral or Ineffective. This low collaboration prevents the necessary process and feedback loops required to reach high MLOps maturity.

The Unpredictable Cost of Production AI

The most challenging aspect of managing cloud costs for production AI applications shows a remarkably even distribution across all cost centres.

When budgeting for production AI applications on AWS, which aspect of cloud costs is most challenging to accurately predict, manage, and optimise for long-term scalability? (Please select one)



22%

GPU/CPU compute costs
for model training and
inference



21%

Consumption of
specialised AWS AI/ML
services (e.g., SageMaker,
Rekognition, Comprehend)



18%

Unforeseen scaling needs
and peak demand
fluctuations



18%

None of the above



21%

The operational overhead
of MLOps tools, logging,
monitoring, and alerting



10%

Data storage, transfer, and
management (e.g., S3, EFS,
Redshift)

The challenge of accurately predicting and managing cloud costs for production AI applications is broadly distributed across several high-cost centres. While GPU/CPU costs were the most frequently cited challenge (22%), they only slightly exceed the difficulty in managing the consumption of specialised AWS services (20%). The lack of a single, dominant bottleneck highlights that technology leaders are struggling with financial predictability across hardware, specialised services, and unforeseen scaling needs, rather than failing on a single cost centre.

The Underlying Problem: Poor Cost Attribution

This mixed result suggests that the core difficulty lies in holistic cost attribution and optimisation.

As established in the previous section, MLOps maturity is low. This lack of automated, standardised pipeline management means costs spill across multiple, disconnected domains: compute, storage, and tooling. Without granular, centralised MLOps control, costs become functionally opaque.

The operational friction: The inability to accurately forecast the final, holistic cost of a successful deployment (from initial training through continuous monitoring) makes it nearly impossible for finance teams to establish a clear Return on Investment timeline for scaled AI systems. The solution requires mature operational discipline, not simply cheaper cloud bills.

[The AI Pilot Purgatory >](#)[Architecting for Reality >](#)[The Human Element >](#)[Value, Vision & The Future >](#)

Defining AI Value: Beyond Model Accuracy

In this survey, senior technology leaders consistently reject traditional technical metrics (e.g., model accuracy, latency) as the primary indicators of value for a deployed AI solution. The focus shifts to metrics that prove measurable impact on the balance sheet and operational efficiency.

The single most critical metric used to define and measure business value is not a single number, but rather the monetised value of process improvement.

Business Value Metric	Executive Focus
Process Automation Rate / FTE Hours Saved	The most frequently cited measure. Quantifies the direct reduction in operational costs by reducing manual effort in specific business processes (e.g., claims, support resolution).
Profit Margin Improvement	The ultimate benchmark. Measures the AI solution's contribution to either increasing the average value of a transaction or decreasing the cost associated with generating that revenue.
Customer Conversion or Churn Reduction	The critical measure for customer-facing AI. Assesses how the solution impacts user behavior, specifically the lift in conversion rates or a measurable reduction in customer attrition.



In partnership with

Ten10 THE SCALE FACTORY
Part of Ten10

The Human Element – Talent & Teams

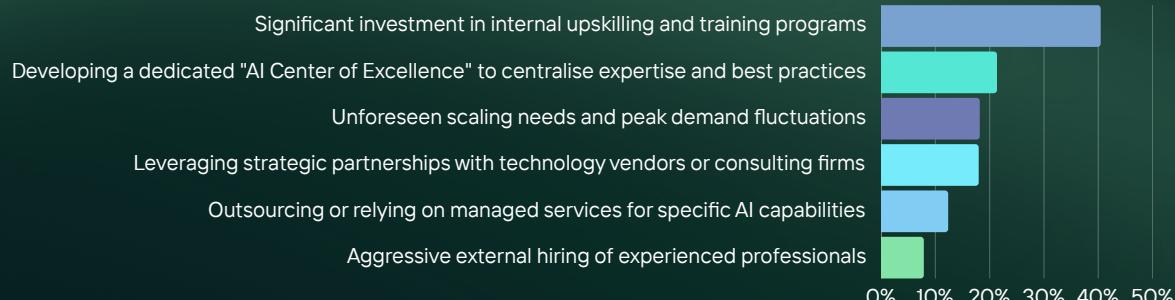
The Talent Gap Strategy: Investing Internally

When addressing the acute shortage of specialised AI/ML engineering and MLOps talent, organisations are clearly prioritising internal development over external acquisition or outsourcing. This reveals a strong desire to fix things themselves, despite the broader evidence suggesting they are struggling to successfully scale AI projects.

The primary strategies for closing the skill gap are:

- Internal Investment Dominates:** 40% of organisations are making significant investment in internal upskilling and training programs. This is the clear, dominant strategy, highlighting a preference for building long-term, internal capability.
- Centralising Expertise:** 21% are focused on developing a dedicated "AI Centre of Excellence" (CoE) to centralise expertise and best practices, another internally-focused approach.

When it comes to the acute shortage of specialised AI/ML engineering and MLOps talent required for production-scale initiatives, what is your organisation's primary strategy for closing these skill gaps?



The research highlights a paradox we see often: businesses are heavily investing in internal upskilling, yet without proven internal successes to model against, this approach struggles to gain traction. At the Ten10 Academy, we focus on breaking this cycle by equipping teams with the skills, frameworks, and confidence needed to drive maturity and achieve scalable success.

Jenny Briant
Academy Director
Ten10



The Paradox of Self-Help:

This strong preference for internal investment (40% upskilling, 21% CoE) creates a critical paradox for struggling businesses. If, as the data shows, two-thirds of organisations have a sub-50% PoC success rate and are still battling basic MLOps maturity (Stage 3 or lower), how effective can internal training be?

Without mature, successful internal pipelines to learn from, the effectiveness of self-led upskilling is severely limited, suggesting a crucial need for some strategic injection of external expertise to break the cycle of repeated failure and provide a scalable template to learn from.

The AI Pilot Purgatory

Architecting for Reality

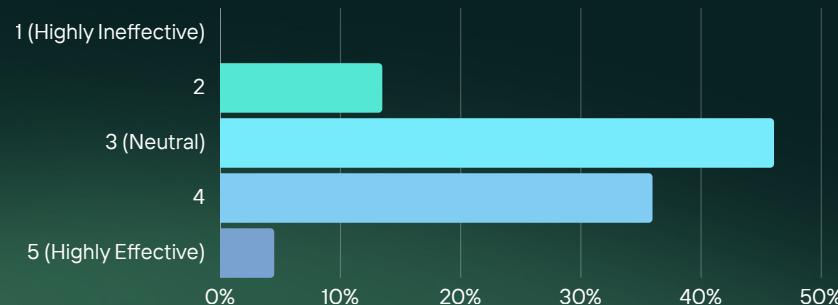
The Human Element

Value, Vision & The Future

The Organisational Root Cause of Low MLOps Maturity

The data establishes a clear cause-and-effect relationship between organisational alignment and technical scaling capability:

How effective is the collaboration between your data science, engineering (DevOps/MLOps), and business teams in bringing AI solutions from initial concept to successful production deployment and ongoing operation?



The Problem (MLOps Maturity)

74% of organisations are stuck in Stages 1, 2, or 3 (Manual, Repeatable, or only basic Automated MLOps).



The Root Cause (Collaboration)

Over 59% of organisations rate their cross-functional collaboration as Neutral or Ineffective.



The Strategic Link

Low Collaboration prevents the necessary cross-functional process and feedback loops required to reach High MLOps Maturity.

This alignment strongly suggests that the technical bottleneck is, in fact, an organisational failure to collaborate. The complexity of production AI demands seamless integration between Data Science (model building), Engineering (MLOps infrastructure), and Business (defining value and measuring ROI).

Without this synergy, organisations may continue to struggle, leading to the high failure rate of PoCs and the proliferation of technical debt documented throughout this report.

Organisational Misalignment

The survey identifies organisational misalignment as the single greatest non-technical barrier. Executive teams believe they're aligned until a PoC succeeds and suddenly everyone has different ideas about what 'production' means, what success looks like, and who's responsible for getting there. This obstacle derails more AI projects than any technical challenge.

Mike Mead
Chief Product & Technology Officer
The Scale Factory



The AI Pilot Purgatory Architecting for Reality The Human Element Value, Vision & The Future 

Value, Vision & The Future

From Proof-of-Concept to Production

The transition from successful AI pilot to scaled enterprise product requires a shift in priorities and investment focus. Experienced technology leaders consistently offered three core pieces of strategic advice for future-proofing AI investment based on the challenges they have already faced and learnt from.

If you could instantly solve one major impediment that currently prevents your organisation from scaling AI across the enterprise, what would it be?

1. The System Problem: MLOps Maturity and Tooling Standardisation

This emerged as the most frequently cited technical barrier, reflecting a fundamental lack of capability for continuous AI deployment. Leaders noted that their organisations lack a reliable, automated end-to-end pipeline A prerequisite for operationalising models efficiently.

- **Seamless Deployment (The "Last Mile"):** A critical gap is the inability to easily and quickly move a validated model from the experimentation environment directly into production.
- **Monitoring and Maintenance:** The absence of robust, automated systems to monitor key production metrics means models often become "stuck" and obsolete once deployed, lacking necessary updates or automated retraining triggers.
- **Standardisation:** There is a profound need for a uniform toolchain and process. Without it, every new AI project is forced to redundantly build its production environment from scratch, preventing resource efficiency and increasing technical debt.

2. The Foundation Problem: Data Governance, Quality, and Accessibility

A close second impediment focuses on the fundamental lack of clean, secure, and accessible data, which often underlies MLOps and scaling failures. If the foundation is weak, no matter how advanced the system, the house will fall.

- **Enterprise Data Access:** Organisations are slowed by bureaucratic and technical hurdles involved in unifying siloed data and obtaining necessary regulatory and security approvals to use that data for AI initiatives across different business units.
- **Data Quality and Lineage:** The poor quality, inconsistency, and lack of clear lineage for training data require constant, manual cleaning effort. This high-friction process undermines model reliability, particularly when models are moved into production environments.

3. The People Problem: Organisational Alignment and Talent Gap

While technical and data issues are prominent, a significant portion of responses pointed to the organisational context as the ultimate speed limit for AI scaling.

- **Talent Scarcity:** There is an acute shortage of MLOps/AI Engineering talent. Professionals capable of building and maintaining production-grade systems. This is distinct from the data scientists who focus primarily on model building in a sandbox environment.
- **Executive Buy-in:** Many organisations struggle to achieve consistent, enterprise-wide executive alignment on AI strategy and funding the central platform teams and tools required for scaled MLOps capabilities.

The AI Pilot Purgatory

Architecting for Reality

The Human Element

Value, Vision & The Future

What one piece of advice, gained from your own journey, would you offer a fellow CTO just starting their path to scale AI from experimental PoCs to robust, production-ready systems that deliver consistent business impact?

1. Prioritise MLOps Infrastructure Over Model Building

- **The Advice:** Invest in the operational pipeline first. Your primary goal is not a better model, but a reliable delivery system.
- **The Takeaway:** A 75% accurate model delivered reliably is infinitely more valuable than a 95% accurate model stuck in the sandbox. Success is measured by uptime and throughput, not just accuracy.

2. Focus on Business Value, Not Technical Accuracy

- **The Advice:** Shift the narrative immediately from technical output to financial outcome.
- **The Takeaway:** Ensure every AI project's success metric is a business KPI (e.g., X dollars saved per quarter, Y% improvement in conversion), not a technical one (e.g., 90% F1 score). An AI solution must prove its financial contribution from day one.

3. Solve the Data Problem Before the Model Problem

- **The Advice:** Recognise that scaling AI is fundamentally a data engineering challenge, not a modeling one.
- **The Takeaway:** Focus engineering efforts on creating clean, accessible, and compliant data APIs that serve the entire organisation. Model building is secondary; the quality of your data foundation determines your ultimate scale.

4. Hire for MLOps and Engineering, Not Just Data Science

- **The Advice:** Acknowledge the critical distinction between those who can build a model (Data Scientists) and those who can run it reliably at scale (MLOps/AI Engineers).
- **The Takeaway:** Build a balanced team. The MLOps Engineer is the "AI plumber" who ensures the system runs; the Data Scientist is the "AI chef" who creates the recipe. You cannot scale without the plumber.

A 75% accurate model delivered reliably is infinitely more valuable than a 95% accurate model stuck in the sandbox. Success is measured by uptime and throughput, not just accuracy.

The Path Forward

Successful organisations know AI scaling isn't just a technical problem, it's a people problem too. They build strong data foundations, understand production requirements early without over-engineering PoCs, apply proven DevOps principles through MLOps, and establish organisational alignment with clear ownership before implementation.

Mike Mead

Chief Product & Technology Officer
The Scale Factory



Insights

Ten10 is a leading technology consultancy, driving digital transformation through expertise in Cloud computing, Automation, and Quality Engineering. Together with The Scale Factory, our AWS cloud specialists, we deliver scalable, tailored solutions that empower businesses to innovate, adapt, and thrive.

Our findings present a candid snapshot of the current state of AI adoption across organisations. Despite increasing investment in AI, the leap from experimental results in controlled environments to real business impact remains uncertain, demonstrating that technical proficiency alone does not guarantee scale or value.

Some hurdles remain in the way, no matter what technology. The persistent misalignment between stakeholder expectations and business realities – where initial AI successes are assumed to be easily scalable – often leads to costly delays and frustration. This gap reflects a broader pattern: business units and technical teams are not always operating with shared objectives, particularly when moving from concept to operational product.

It is striking that over a third of respondents acknowledge developing AI prototypes without production-readiness in mind. This indicates a widespread lack of forward planning and highlights the prevalence of siloed experimentation, rather than a co-ordinated push towards robust solutions. Unsurprisingly, when it comes time to integrate with existing systems, concerns around security, data quality, and system compatibility surface as major technical hurdles. These findings reinforce the understanding that AI at scale is as much about robust engineering and data governance as it is about innovation in modelling.

Mike Mead
Chief Product & Technology Officer
The Scale Factory



Overall, we see an industry still finding its footing. Ambition is high, and experimentation is prolific, but the shift to reliable, scalable value remains a challenge. The clearest lesson is that successful AI adoption demands much more than technical success in a controlled environment; it requires a concerted focus on organisational readiness, collaboration, and operational discipline.

While the challenges outlined in this report may seem daunting, they are a testament to the ambition and forward-thinking nature of organisations striving to adopt AI. The fact that so many are investing in this journey reflects a collective understanding that AI has the potential to transform industries, drive innovation, and deliver significant value. These struggles are not signs of failure but rather growing pains on the path to achieving something truly impactful. We would like to thank everyone who participated in this survey for sharing their experience, insights, and advice with their peers.



Insights

CTO Craft provides a safe space for technology leaders to share ideas and challenges, network with other leaders and engage in online and in-person events to help accelerate their learning in the art of leading technology teams.

The ease with which we can now build AI prototypes is both a blessing and a curse. Thanks to the commoditisation of AI and powerful LLM tools, the barrier to entry for AI experimentation has practically evaporated, making it cheaper and faster than ever to get a Proof of Concept off the ground. But this ease has masked a far deeper, systemic problem: the journey from a successful, disposable trial to a resilient, value-generating system is proving to be a true ordeal for most organisations.

From what we can see from the responses to our survey, achieving measurable business impact from an AI PoC is essentially "flipping a coin": we found that 67% of organisations have a sub-50% success rate in making that leap. This isn't a lack of technical talent, though, it's a profound failure to bridge the organisational and engineering gap that appears once the initial excitement wears off.

The core issue lies in a failure of perspective: the inability to shift from a "science project" mentality to a "business product" one. When a data science team announces a model has hit 90% accuracy in the sandbox, business leaders can assume that it's instantly deployable. What they miss is the expensive, non-glamorous investment needed for MLOps, security, and integration. This misaligned stakeholder expectation is the single greatest non-technical barrier to success.

It's an organisational wall we keep hitting, where internal bureaucracy and prioritisation shifts halt momentum. Funding which is often readily available for the thrill of exploration disappears for the less exciting work of hardening the system for scale. A project dies not because the model failed, but because it was never treated as a critical business capability requiring dedicated product ownership.

This lack of foresight is why so many prototypes are doomed to fail at scale - 37% of respondents effectively admitted they are "flying blind" by developing their AI prototypes without production readiness in mind from day one. This creates inevitable and significant technical debt, explaining why 74% are stuck in manual or basic automated MLOps stages.

The technical friction is simply the symptom of a deeper illness: poor cross-functional collaboration. Over 59% of respondents rated the collaboration between Data Science, Engineering, and Business teams as Neutral or Ineffective. This breakdown prevents the necessary feedback loops that turn a successful model trial into a robust system. The lesson here is simple but hard-won: Success is measured by uptime and throughput, not just accuracy. We have to acknowledge the critical difference between the "AI chef" (the Data Scientist) and the "AI plumber" (the MLOps Engineer). You cannot scale without the plumber. The road to scale requires trading the fleeting high of the easy prototype for the enduring value of disciplined, collaborative, production-ready engineering.

Andy Skipper
Founder, CTO Craft



In partnership with

Ten10 THE SCALE FACTORY

Methodology

This report is based on a structured, primary research effort designed to capture the unfiltered perspectives and real-world challenges faced by senior technology decision-makers responsible for Artificial Intelligence (AI) and Machine Learning (ML) adoption within their businesses.

The methodology encompassed two primary phases: data acquisition via a targeted industry survey, and rigorous qualitative and quantitative analysis for report generation.

Survey Design and Data Acquisition

Survey Objective and Scope

The core objective of the survey was to move beyond anecdotal evidence and quantify the key systemic and organisational barriers preventing successful AI scaling. The survey was meticulously structured to investigate five crucial dimensions of the AI lifecycle:

1. **Success Quantification:** Establishing the actual rate of Proof-of-Concept (PoC) transition into measurable, production-ready business value.
2. **Bottleneck Identification:** Pinpointing the most critical technical and non-technical challenges.
3. **MLOps Maturity Assessment:** Benchmarking the industry against established MLOps maturity models (ranging from Manual to Fully Automated).
4. **Organisational Health:** Assessing the effectiveness of cross-functional collaboration essential for successful deployment.
5. **Strategic Intent:** Understanding the primary focus areas for future investment and the single most pressing impediment leaders would solve instantly.

The survey employed a mix of quantitative metrics and qualitative, open-ended questions designed to capture nuanced, actionable insights.

Contribution of the CTO Craft Community

The research was conducted in collaboration with the CTO Craft community, a highly engaged and relevant professional network of senior technology leaders. This partnership served two critical functions:

Firstly, it ensured the high relevance and quality of the respondents. By targeting a known, established community of practitioners, the data collected reflects insights from individuals who hold direct budgetary and strategic responsibility for the subjects discussed.

Secondly, the community's insight was vital during the survey design and validation phase. Input from community leaders helped refine the terminology, ensure questions were framed around real-world challenges, and eliminated ambiguity, thus optimising the utility of the final dataset.

Respondent Demographics and Data Analysis

The demographics confirm a highly relevant respondent group critical to understanding AI scaling challenges:

- Senior Leadership Focus: Over 85% held senior titles (CTO, VP/Director of Engineering, Head of Data/AI), ensuring insights come directly from decision-makers.
- Scale Challenges: Responses were concentrated in mid-to-large enterprises (over 500 staff). This is crucial as scaling AI in organisations of this size presents the most severe organisational alignment, security, and integration challenges.
- Sector Diversity: Responses covered Tech, Finance, Retail, Healthcare, and Media, broadening the findings' strategic applicability beyond a single industry.

Analytical Approach

The analysis used a mixed-methods approach, combining quantitative data with qualitative synthesis (thematic analysis of open-ended responses) for depth. The report structure is designed to guide senior leaders through diagnosis, prognosis, and prescription, focusing on actionable strategy related to organizational design and resource shifting.

Mind the Gap: Bridging from Sandbox to Scale

Brought to you by CTO Craft, Ten10 and The Scale Factory

We would like to extend a heartfelt thank you to everyone who took the time to complete the survey. Your valuable insights and honest feedback make this report possible and help us continue to support and empower the tech leadership community.

This report wouldn't be what it is without your contributions, and we truly appreciate your willingness to share your experiences. By participating, you've played an essential role in helping us better understand the evolving landscape of tech leadership.

A special thanks to our partners at Ten10 and The Scale Factory for their continued support and collaboration.



ctocraft.com



ten10.com



THE SCALE FACTORY

Part of Ten10

scalefactory.com

Behind the Report:

Andy Skipper
Founder
CTO Craft



Claire Bobie
Stakeholder Relationship Manager
CTO Craft



Nick Bates
COO
CTO Craft



Matt Hunt
Head of Marketing
CTO Craft



Jenny Briant
Academy Director
Ten10



Mike Mead
Chief Product & Technology Officer
The Scale Factory



In partnership with