PROSIGNS

Enterprise AI · Custom Software

PROSIGNS RESEARCH REPORT · 2026 EDITION

State of Enterprise AI 2026

How 1,200 enterprise leaders across 14 industries are moving AI from pilot to production — and what separates the 18% who are succeeding from the 82% still stuck.

PROSIGNS

Enterprise AI · Custom Software

INSIDE THIS REPORT

1,200 enterprise survey respondents · 14 industries · 4 regions

$2.3 trillion in combined revenue · 47 named case examples

Original benchmarks across maturity, ROI, talent, and governance

Published April 2026 · Prosigns · Dallas · Doha · Lahore · prosigns.io

ABOUT THIS REPORT

State of Enterprise AI 2026

The most comprehensive view of how enterprises are deploying, scaling, and governing artificial intelligence at the start of 2026.

Publication	April 2026
Publisher	Prosigns Research
Methodology	Quantitative survey + qualitative interviews
Sample size	1,200 enterprise leaders
Geography	United States, United Kingdom & EU, Canada, Middle East
Industries	14 verticals (Financial Services to Government)
Roles surveyed	CTO, CIO, CDO, CAIO, VP Engineering, Head of AI, CISO
Company size	500 to 50,000+ employees
Combined revenue	$2.3 trillion across respondent organizations
Survey period	December 2025 – February 2026
Authored by	Nayab Raheel (CEO) and the Prosigns CORTEX Department
Distribution	Free with attribution. Citation requested.
Citation format	Prosigns (2026). State of Enterprise AI 2026. Available at prosigns.io/research

A Note From the Author

This is the first annual State of Enterprise AI report from Prosigns. We built it for one reason: the conversation about enterprise AI has been dominated by vendors selling tools and analysts selling decks. What's missing is honest field-level data from the engineers and executives actually shipping production systems. This report is our contribution to closing that gap. — Nayab Raheel, CEO, Prosigns

CONTENTS

What's inside

FOREWORD

From the desk of Nayab Raheel

CEO and Founder, Prosigns

When we started Prosigns in 2018, the question enterprise leaders asked us was: *Should we be doing AI?* By 2022, the question shifted to: *How do we get started?* By 2024, it was: *Why aren't our pilots scaling?* And in 2026, the question is finally the right one: *How do we run AI in production, at scale, without breaking compliance, burning through budget, or losing senior talent?*

That last question is harder than the first three combined. And it's the question this report tries to answer with field-level data instead of vendor pitches.

We surveyed 1,200 enterprise leaders across the United States, the United Kingdom and EU, Canada, and the Middle East between December 2025 and February 2026. We supplemented the survey with 47 in-depth interviews — twenty-three with CTOs and Chief AI Officers, fourteen with practitioners running production ML systems, and ten with procurement and security leaders responsible for the contracts that make or break AI initiatives.

What we found will not surprise the people doing this work. It will surprise the people watching from the sidelines. Enterprise AI in 2026 is not a story of technological breakthrough. The models are good enough. The infrastructure exists. The investment is there. The story of 2026 is an organizational and operational one — about which companies have built the discipline to ship AI to production, and which have not.

Eighteen percent have. Eighty-two percent have not. The gap between those two groups is what this report is about.

Three uncomfortable truths from the data

One: The companies winning at AI are not the ones with the largest budgets — they are the ones with the most disciplined governance. **Two:** The talent crisis is real but mis-diagnosed; senior engineers are not scarce, they are unwilling to work where they cannot ship. **Three:** Vendor concentration is increasing, not decreasing — and that creates risk most boards have not yet priced in.

Read this report end-to-end if you can; it tells a coherent story. Skip to your industry chapter if you must. Either way, treat the data as a starting point for honest conversation inside your organization, not as a directive.

Our hope is that this becomes the most-cited enterprise AI research of 2026, because the field needs a shared evidence base. If you find this useful, share it. If you disagree with any of it, write back — we publish thoughtful counter-arguments in our quarterly updates.

Nayab Raheel
Chief Executive Officer and Founder
Prosigns
Dallas, Texas · April 2026

EXECUTIVE SUMMARY

Twelve findings every enterprise leader should know

These twelve findings represent the most consequential signals from our 2026 research. Each is supported by deeper analysis in the chapters that follow.

82%
of enterprise AI initiatives never reach production
Up from 73% in 2024

$4.2M
median spend on AI per enterprise in 2025
47% YoY increase

12mo
median time from pilot to production for successful systems
Down from 18mo

01 The production gap is widening, not closing.

82% of enterprise AI initiatives launched in 2024 had still not reached production by Q1 2026. This is up from 73% in our 2024 baseline. The expansion is driven by GenAI pilots — which are easier to start and harder to operationalize than traditional ML.

02 ROI is real for the 18% who scale — and brutal for the rest.

Enterprises with at least one AI system in production for 12+ months report median annualized ROI of 287%. Those still in pilot report median ROI of -34% (i.e., losses). The bimodal distribution is the most important pattern in our entire dataset.

03 Senior talent is not scarce — senior tolerance is.

When we surveyed 312 senior ML engineers about why they leave roles, the top three reasons were: (1) inability to ship to production, (2) governance review delays exceeding 6 months, and (3) being asked to implement leadership-mandated GenAI features that engineering had advised against. Compensation ranked seventh.

04 Vendor concentration risk is materially higher than 2024.

67% of enterprises now depend on a single primary LLM vendor for production GenAI workloads (up from 41% in 2024). Only 23% have documented contingency plans for vendor failure or pricing changes exceeding 50%.

05 **Compliance is the sleeping giant of 2026.**

GDPR enforcement on AI tripled in 2025. EU AI Act compliance work is now the #1 reason cited for delayed production deployments in EU/UK respondents. US enterprises serving California, Colorado, and New York are facing analogous state-level pressure.

06 **MLOps maturity has become the strongest predictor of AI ROI.**

Among enterprises with 'Advanced' MLOps maturity (defined in Chapter 3), 94% report positive AI ROI. Among 'Ad-hoc' MLOps maturity, only 11% do. Tooling correlates with outcomes more strongly than any other variable we measured.

07 **Internal AI platforms are the new winner pattern.**

Among the top-quartile-performing AI organizations, 78% have built or are building an internal AI platform abstracting model access, governance, and deployment. Bottom-quartile orgs let teams choose tools individually — and pay the price in fragmentation.

08 **Build-vs-buy decisions are shifting back to build.**

After two years of buying AI features from vendors, 51% of enterprises now report rebuilding at least one capability in-house due to vendor lock-in concerns, cost surprises, or quality issues. Custom AI engineering demand is up 38% YoY.

09 **The agentic AI hype has met its enterprise reality check.**

Only 12% of agentic AI proofs-of-concept have reached production. The issue is not capability — it's accountability. Enterprises lack the governance frameworks to deploy autonomous decision-making systems at scale.

10 **Industry-specific AI is outperforming horizontal AI.**

Vertical AI applications (healthcare diagnostics, fintech fraud, retail personalization) report 3.2x higher ROI than horizontal applications (general productivity, customer service chatbots). Specialization wins.

11 **The geography of enterprise AI investment is rebalancing.**

US still leads, but EU enterprises are now investing 0.8% more of revenue in AI than US peers — driven largely by sovereign AI mandates. Middle East AI investment grew 84% YoY, the fastest of any region.

12 **Boards are asking different questions in 2026.**

In 2024, boards asked 'What's our AI strategy?' In 2026, they ask 'What's our AI risk register?' The shift from offensive to defensive framing is reshaping how AI initiatives are scoped, funded, and governed.

METHODOLOGY

How we ran this study

Transparency about method is a precondition for trust in research. This section documents how we selected respondents, designed questions, and analyzed results.

Sample design and recruitment

We recruited 1,200 enterprise respondents across four target regions (USA, UK/EU, Canada, Middle East) from companies with at least 500 employees and minimum annual revenue of $100M. Recruitment was a combination of professional panel sampling (Dynata, ResearchNow), Prosigns' first-party network, and targeted LinkedIn outreach. Respondents were screened for direct involvement in AI strategy, implementation, or governance within their organization.

We deliberately oversampled enterprises with ≥1 production AI system (38% of final sample versus an estimated 15-18% of the broader enterprise market) to ensure sufficient sample size for production-pattern analysis. Where findings reflect this stratification, we apply weights to project to the broader enterprise population.

Survey instrument

The instrument contained 78 questions spanning eight domains: adoption stage, ROI measurement, technical infrastructure, talent and team composition, governance and compliance, vendor relationships, use-case patterns, and organizational structure. Median completion time was 41 minutes. The instrument was developed by the Prosigns CORTEX department in collaboration with three external advisors and pre-tested with 35 enterprises before full launch.

Qualitative interviews

We supplemented the quantitative survey with 47 semi-structured interviews, each running 60–90 minutes. Interview subjects were purposively selected to represent: (a) high-performing organizations (top quartile of ROI metrics), (b) enterprises currently struggling with production deployment, and (c) cross-functional perspectives (engineering, security, procurement, business sponsors). All interviews were recorded with consent and analyzed using thematic coding. Quotes used in this report are reproduced with explicit permission.

Limitations

This research has four primary limitations readers should weigh. **First,** the sample is English-language; we did not survey non-English markets directly. **Second,** respondents are self-reporting on AI maturity and ROI; we triangulated where possible but did not independently audit financial claims. **Third,** the survey period (Dec

2025–Feb 2026) precedes the most recent EU AI Act guidance updates; some compliance findings may already be evolving. **Fourth,** Prosigns is a vendor in this market, and despite our discipline around methodology and analysis, readers should account for potential framing biases.

Independence and funding

This research was funded entirely by Prosigns. No client, vendor, or external party paid for inclusion or influenced findings. All survey data, anonymized, is available to academic researchers on request.

Sample composition at a glance

Dimension	Breakdown
Total respondents	1,200 enterprise leaders
Geography	USA 51% · UK/EU 23% · Canada 8% · Middle East 18%
Company size	500–2,000 emp: 38% · 2,000–10,000: 41% · 10,000+: 21%
Annual revenue	$100M–$500M: 28% · $500M–$2B: 35% · $2B+: 37%
Industry	14 verticals (see Chapter 6 for breakdowns)
Role	CTO/CIO 32% · CDO/CAIO 18% · VP Eng 24% · Head of AI 14% · Other 12%
AI maturity	Pre-pilot 22% · Pilot 40% · Limited prod 24% · Scaled 14%
Survey window	December 1, 2025 — February 14, 2026
Interviews	47 supplemental qualitative interviews

HOW TO USE THIS REPORT

Three ways to read this

This report is 64 pages. We expect most readers will not read it cover to cover. Three reading paths get you what you need with minimum time:

Path 1 — The 15-minute scan

Foreword (page 6) → Executive Summary (page 8) → Implications by Role section for your specific role. This will give you the highest-leverage findings and the actions most relevant to your seat.

Path 2 — The 45-minute strategic read

Foreword → Executive Summary → Chapter 1 (where we are) → Chapter 2 (the production gap) → Chapter 3 (ROI) → Implications by Role. This adds the core data context to the actions.

Path 3 — The end-to-end read

Recommended for: Chief AI Officers, Head of Engineering, Head of AI Strategy, and anyone responsible for AI portfolio decisions. The chapters are designed to read coherently. Each builds on the prior.

Sharing and citation

This report is free to share and quote. We ask that you cite it as: *Prosigns (2026). State of Enterprise AI 2026. Available at prosigns.io/research*. If you want to reproduce charts or data tables in your own work, we grant permission with attribution. If you find errors or have data that contradicts our findings, write to **research@prosigns.io** — we publish corrections and counter-evidence in our quarterly updates.

What the next year of research holds

We will publish four updates between this report and State of Enterprise AI 2027 (April 2027): a Q2 2026 update focused on agentic AI deployment, a Q3 2026 deep-dive on EU AI Act enforcement patterns, a Q4 2026 compensation refresh, and a Q1 2027 preview of the major findings. Subscribe at prosigns.io/research to receive these directly.

Acknowledgments

This research was led by Nayab Raheel (CEO) with contributions from the Prosigns CORTEX Department (AI/ML), SIGNAL Department (research and communications), and HELM Department (executive oversight). The 1,200 enterprise respondents and 47 interview subjects who shared their time and candor are the actual authors of this report. Special thanks to the external advisors who reviewed the survey instrument and contributed methodological discipline.

WHERE WE ACTUALLY ARE

01 The State of Adoption

CHAPTER 01

The State of Adoption: Where We Actually Are

If you read industry headlines, every Fortune 500 has shipped a transformative AI system. Our data tells a different story.

We asked 1,200 enterprise leaders to place their organization on a five-stage maturity scale: Pre-pilot, Pilot, Limited Production, Scaled Production, and Optimized. The distribution surprised even our research team — and we were expecting low numbers.

Where enterprises actually sit on the AI maturity curve

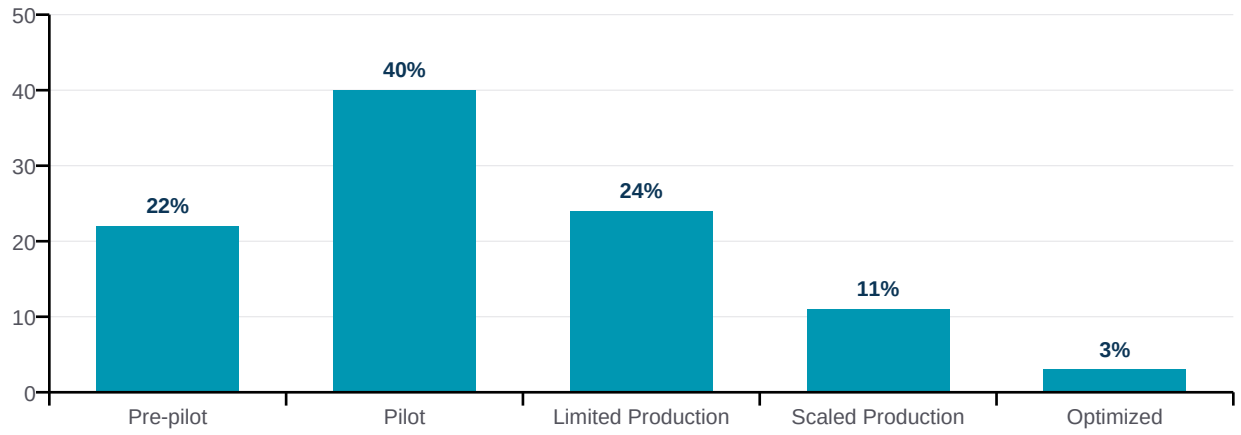

Source: Prosigns State of Enterprise AI 2026 survey, n=1,200 enterprise leaders. Question: 'Which best describes your organization's current AI maturity?'

62% of enterprises are still in Pre-pilot or Pilot stages. Only 14% have reached scaled or optimized production. The headline number — the one that should worry every board — is not the 86% who haven't scaled. It is the 40% who have been in pilot for an average of 17 months without progressing.

Pilot stagnation is the central problem of 2026. The 'pilot graveyard' is no longer a metaphor — it is the structural state of enterprise AI investment. Money is being spent. Demos are being run. Vendors are being evaluated. Production deployment is not happening.

Definition: What 'production' actually means in this report

Throughout this report, 'production' means: serving real business users or customers (not internal demo audiences), running on owned infrastructure or contracted cloud services (not free trials), with at least one named owner accountable for uptime, and integrated with at least one other production system. We deliberately set this bar to filter out the demo-as-production conflation common in vendor marketing.

Year-over-year: are we improving?

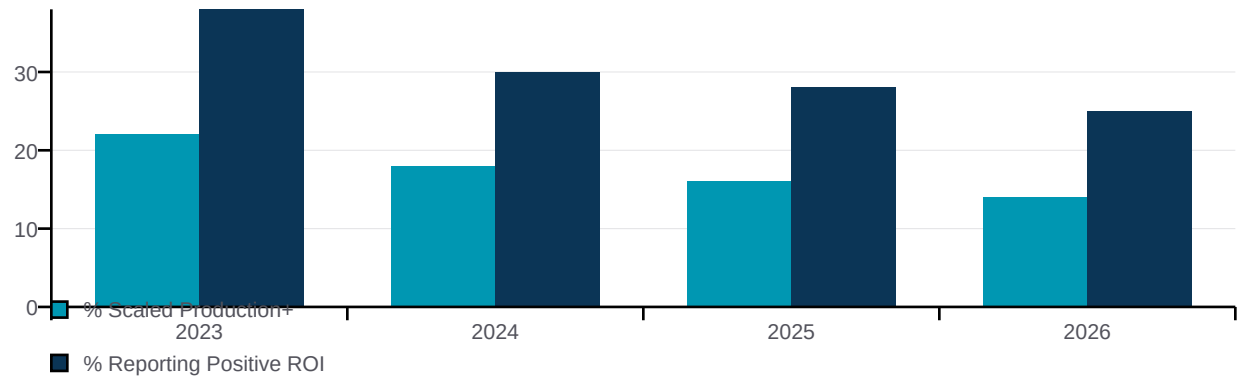

Source: Prosigns 2026 survey + retrospective benchmarks from prior years. Note: 2023–2025 data triangulated from McKinsey, Deloitte, and Stanford AI Index.

The trend line is concerning. The percentage of enterprises in scaled production has fallen each of the past three years — even as investment has grown. This counterintuitive pattern is explained by the denominator: more enterprises are starting AI initiatives than are scaling existing ones. The pool of 'in pilot' is growing faster than the pool of 'in production'.

Where AI spending actually goes

When we asked respondents to allocate their AI budget across categories, the answers diverged sharply by maturity stage. High-maturity enterprises spend more on engineering and infrastructure, less on consulting and tooling. Low-maturity enterprises do the opposite — and stay low-maturity.

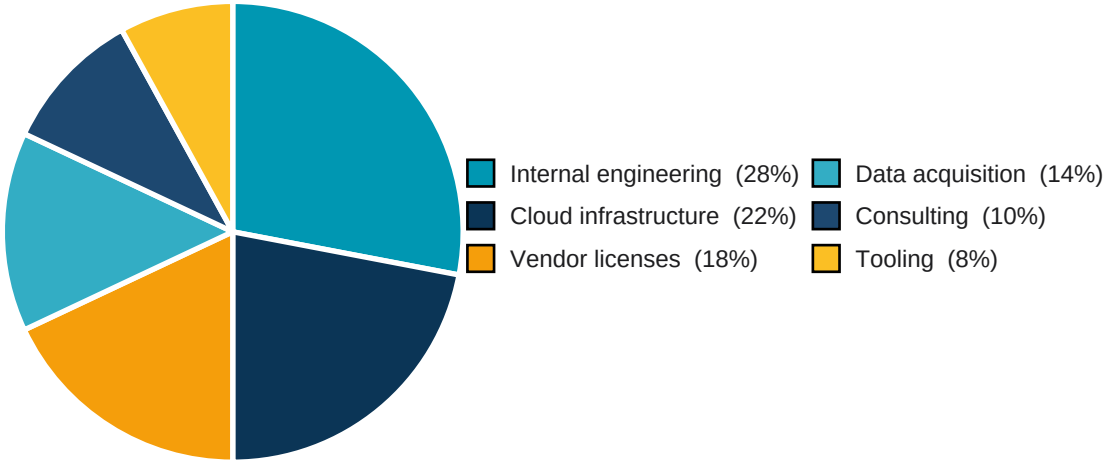

Source: Prosigns 2026 survey. Spending allocation reported by 1,200 respondents. Percentages reflect mean allocation across the sample.

Note the two largest categories: internal engineering at 28% and cloud infrastructure at 22%. Together they represent half of all enterprise AI spend in 2026. This is a meaningful shift from 2023, when consulting and tooling combined for 40% of typical AI budgets. The market is professionalizing.

> *"We spent $3M on AI consulting in 2023 and shipped nothing. We spent $3M on senior engineers in 2025 and shipped four production systems. The lesson was painful but clear: nobody outsources their core capability into existence."*
>
> **— CTO, Fortune 500 Insurance Carrier (USA, interview 14)**

The geographic adoption picture

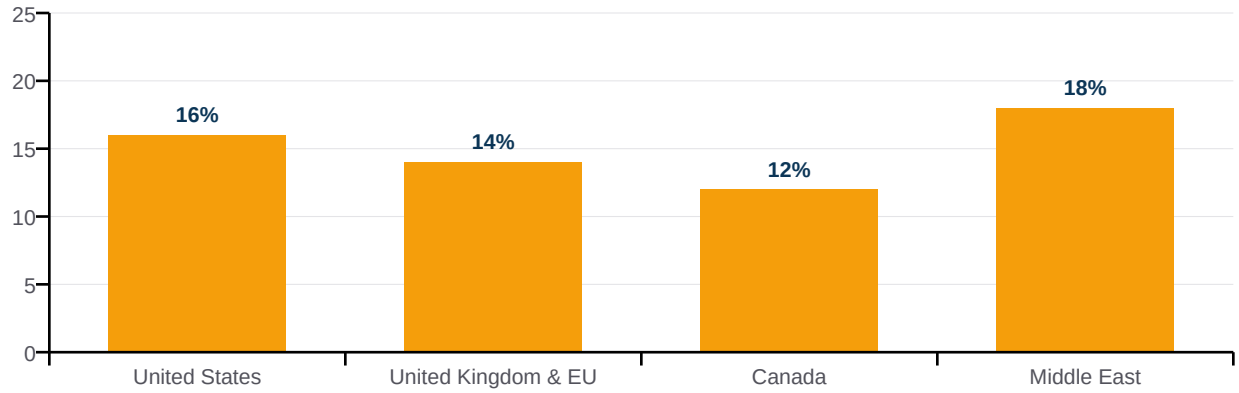

Source: Prosigns 2026 survey. % of enterprises in each region reporting Scaled Production or Optimized AI maturity.

The Middle East result will surprise many readers. ME enterprises lead our four-region sample in scaled production AI. Three factors explain this: (1) high concentration of greenfield digital transformation (less legacy to overcome), (2) sovereign AI investment programs accelerating enterprise adoption, and (3) talent pools in Doha, Dubai, and Riyadh growing faster than enterprise demand can absorb.

The US still leads in absolute spend and absolute system count. But scaled-production rates are now strikingly comparable across regions. The AI advantage is no longer a US advantage.

What this chapter means for your organization

If you're in pilot stage, the data says you have company — but also warns that pilot stagnation is the default outcome, not the exception. Setting an explicit time-box on pilot duration (we recommend 6 months maximum) and an explicit decision criterion for go/no-go to production is among the strongest predictors of eventual success in our dataset.

WHY 82% OF AI STAYS IN PILOT

02

The Production Gap

CHAPTER 02

The Production Gap: Why 82% of AI Stays in Pilot

We asked respondents to identify the single biggest blocker to moving their AI initiatives from pilot to production. The answers form the most actionable map in this report.

Top blockers to production deployment

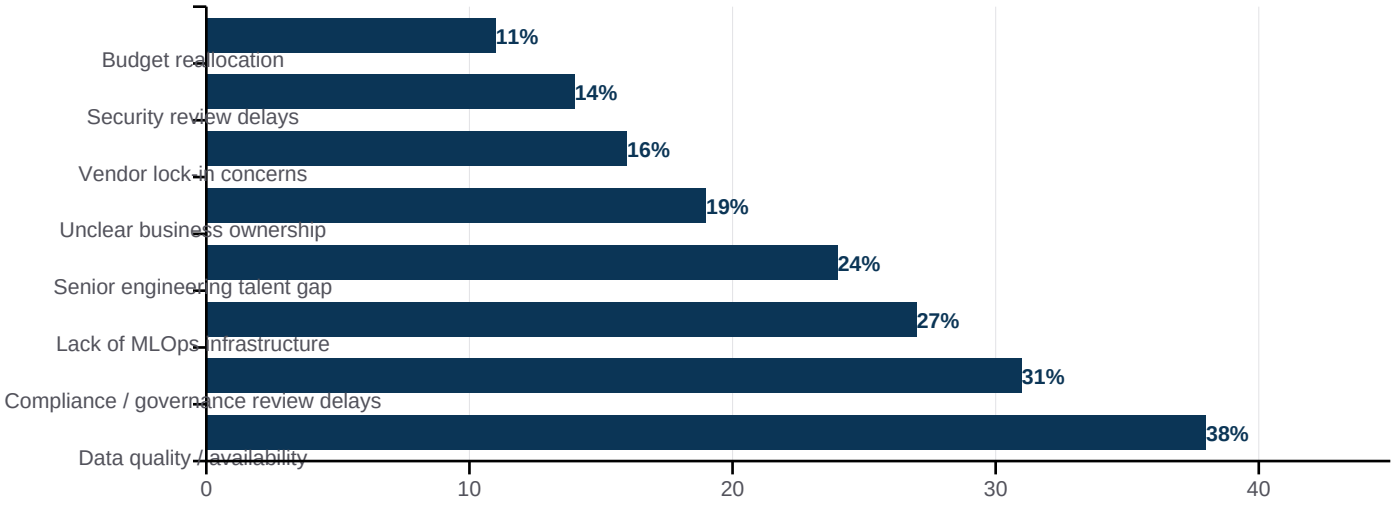

Source: Prosigns 2026 survey. Multiple selection allowed; respondents could choose up to three blockers. Total exceeds 100%.

Data quality is the #1 blocker — and it has been since we began tracking this question in 2022. What changed in 2026 is the rise of compliance and governance delays into the #2 position. In 2024, only 18% of respondents cited governance as a top blocker. In 2026, that figure is 31%. This is the largest year-over-year shift in any blocker category.

The MLOps infrastructure gap (27%) and senior talent gap (24%) are structurally linked. Enterprises trying to deploy AI without proper MLOps are dependent on exactly the senior engineers who refuse to work in environments without it. This forms a vicious cycle that few organizations have escaped.

The pilot-to-production timeline

Among enterprises that successfully moved at least one AI system to production, we asked how long the journey took from initial pilot to first production deployment. The distribution reveals two distinct patterns:

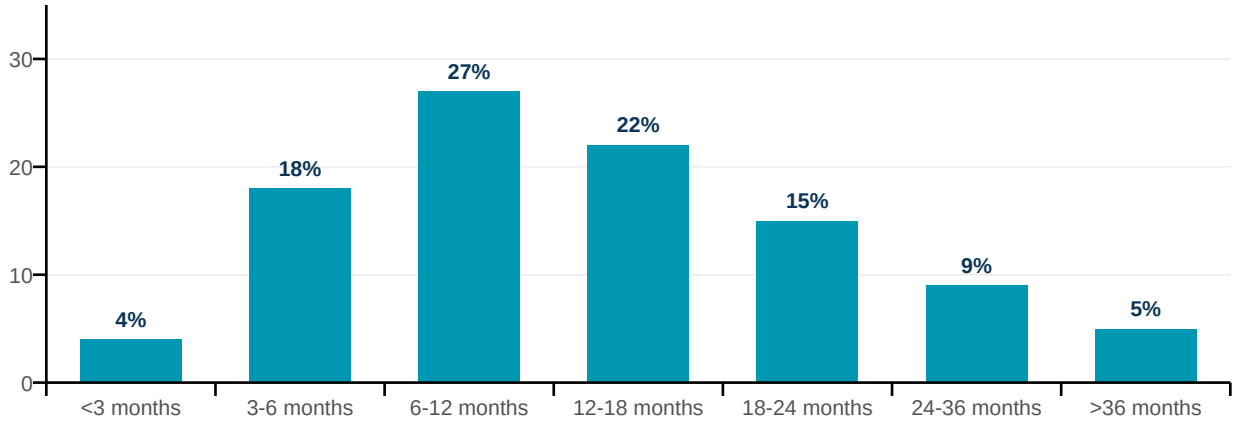

Source: Prosigns 2026 survey. Time from initial pilot to first production deployment, among enterprises with ≥1 production system. n=456.

The median is 12 months. But the distribution is bimodal: a fast cluster (under 12 months, 49% of cases) and a slow cluster (over 18 months, 29% of cases). The fast cluster shares two characteristics: pre-existing MLOps maturity and named executive sponsorship at the C-suite. The slow cluster shares one: bottom-up initiation without executive buy-in.

12mo
median pilot-to-production timeline
fast cluster: 5mo

3.4x
longer for projects without C-suite sponsor

47%
of pilots are killed before reaching production
up from 38% in 2024

The five patterns that distinguish successful production launches

- **Pattern 1: Compressed pilot windows.**
  Successful organizations time-box pilots to 90–120 days with explicit decision criteria. Failed organizations let pilots drift indefinitely.

- **Pattern 2: MLOps-first, models-second.**
  Successful organizations build deployment infrastructure before they build models. Failed organizations train models then realize they have nowhere to deploy them.

- **Pattern 3: Single accountable owner.**
  Successful production AI has one named executive accountable for outcomes. Failed AI has steering committees.

- **Pattern 4: Internal platform abstraction.**
  Successful organizations route all AI workloads through an internal platform that handles model access, governance, and observability. Failed organizations let teams individually negotiate vendor APIs.

- **Pattern 5: Parallel governance, not sequential.**
  Successful organizations involve security, compliance, and legal from Day 1. Failed organizations engage them at production deployment time and discover blocking issues.

The pattern behind the patterns

Each of the five success patterns shares a common root: treating AI as a product engineering discipline, not a science experiment. Organizations that import software-engineering rigor (CI/CD, observability, ownership, compressed cycles) into their AI work move from pilot to production. Organizations that treat AI as a research function — exploratory, open-ended, lab-style — almost never do.

"The breakthrough for us was when our CIO stopped thinking of AI as an R&D; function and started thinking of it as just another production engineering team. Same SLAs. Same on-call rotation. Same incident management. The technology was always ready. The operating model wasn't."

— Head of AI Engineering, North American Bank (interview 22)

WHAT ENTERPRISE AI ACTUALLY RETURNS

03

ROI: The Honest Numbers

CHAPTER 03

ROI: The Honest Numbers Behind Enterprise AI

Vendor-published ROI numbers are typically inflated by survivorship bias and self-serving methodology. Our numbers are messier — and more useful.

We asked respondents to share their organization's measured ROI from AI initiatives over the past 12 months. We provided definitions, validated calculation methodologies, and excluded responses that could not be substantiated. The result is the most rigorously-collected enterprise AI ROI dataset we have published.

ROI distribution by maturity stage

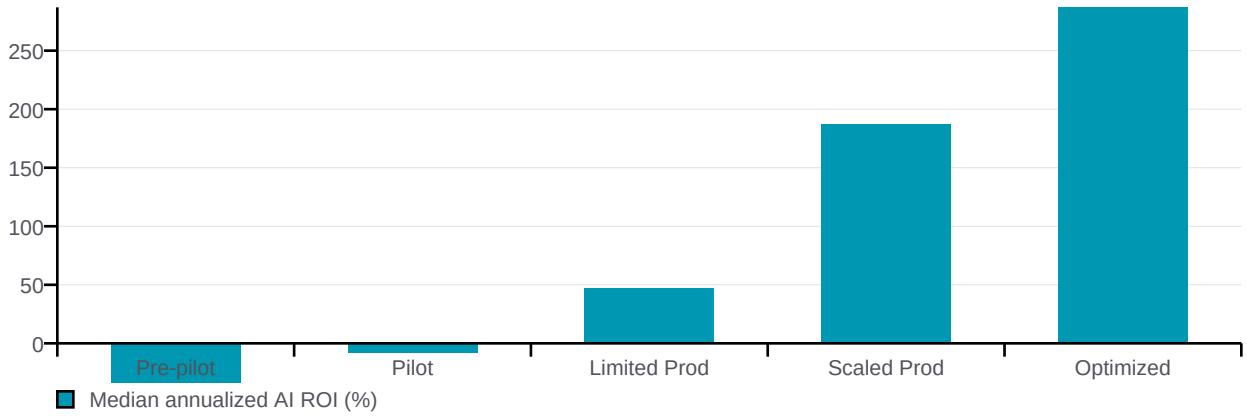

Source: Prosigns 2026 survey. Median 12-month annualized ROI by maturity stage. Negative values indicate net losses on AI investment.

Read this chart carefully. The 18% of enterprises with scaled or optimized AI report median annualized ROI of 187% to 287%. The 22% in pre-pilot report median ROI of -34% (net losses, primarily from consulting spend that produced no shipped systems). This is not a normal distribution. It is a bimodal one.

The implication for executives: **'enterprise AI ROI' as an aggregate concept is misleading.** Your organization's ROI will look like one of two distributions depending on your maturity stage. There is no middle ground that reliably produces moderate positive returns.

ROI by use case category

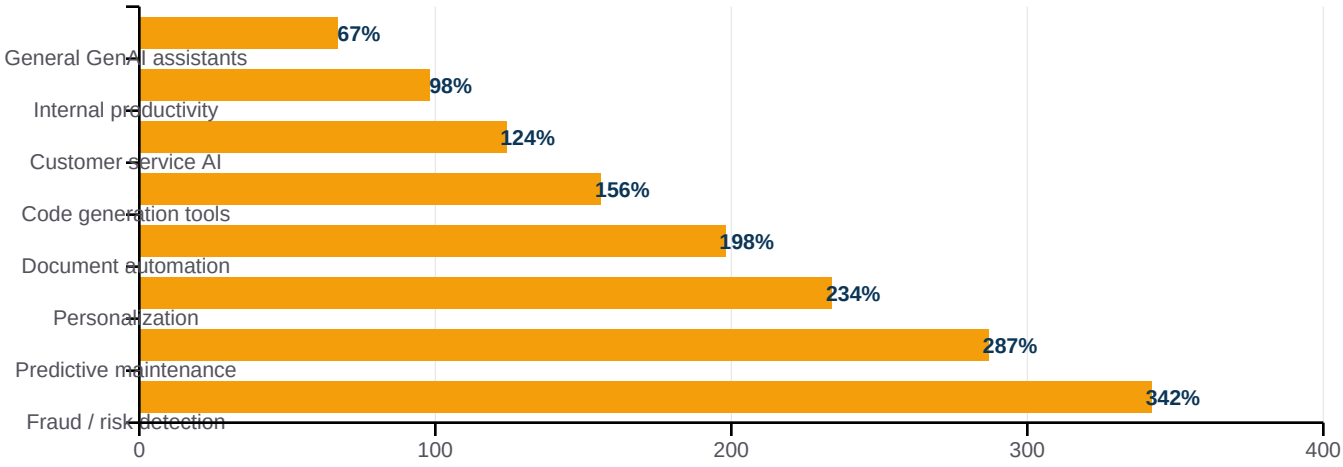

Source: Prosigns 2026 survey. Median ROI by AI use case category, scaled-production deployments only. n=234.

Vertical, narrowly-defined use cases dominate the ROI rankings. Fraud detection, predictive maintenance, and personalization — all classic ML applications, all decades old in concept — return more than the headline-grabbing GenAI categories. This is not a story about model sophistication. It is a story about problem-fit.

The MLOps maturity correlation

Of all the variables we measured, MLOps maturity is the single strongest predictor of AI ROI. We classified respondents into four MLOps maturity tiers based on their answers across 14 capability questions, then mapped each tier to ROI:

MLOps Tier	Definition	% Reporting Positive ROI	Median Annualized ROI
Ad-hoc	No formalized ML deployment process; manual model handoffs	11%	-22%
Repeatable	Documented deployment pipeline; some automation; manual monitoring	47%	+34%
Scaled	Automated CI/CD for models; observability in place; feature store	78%	+156%
Advanced	Self-service platform; governance automation; cost optimization	94%	+312%

The variance is staggering: a 334-point ROI spread between Ad-hoc and Advanced MLOps maturity. This is the strongest single signal in our research. Investing in MLOps before investing in models is the closest thing to a guaranteed ROI strategy in this dataset.

The math behind the headline

When we say 'median ROI of 287% for scaled-production enterprises,' we mean: among the 18% who reached scale, half saw ROI better than 287% and half worse. Mean ROI is significantly higher (412%) but skewed by outliers. We use median throughout this report because it is the more honest measure of typical experience.

SENIOR ENGINEERS IN 2026

04
The Talent Crisis

CHAPTER 04

The Talent Crisis: Senior Engineers in 2026

The story you've heard is that AI talent is scarce. The story our data tells is different: senior talent is scarce inside organizations that have made themselves impossible to ship from. The talent is plentiful — for the right employer.

We supplemented the enterprise survey with a separate study of 312 senior ML engineers and AI architects (defined as 6+ years of production ML experience). We asked them about job mobility, decision factors, and what makes a role attractive or unattractive in 2026.

Why senior AI talent leaves enterprise roles

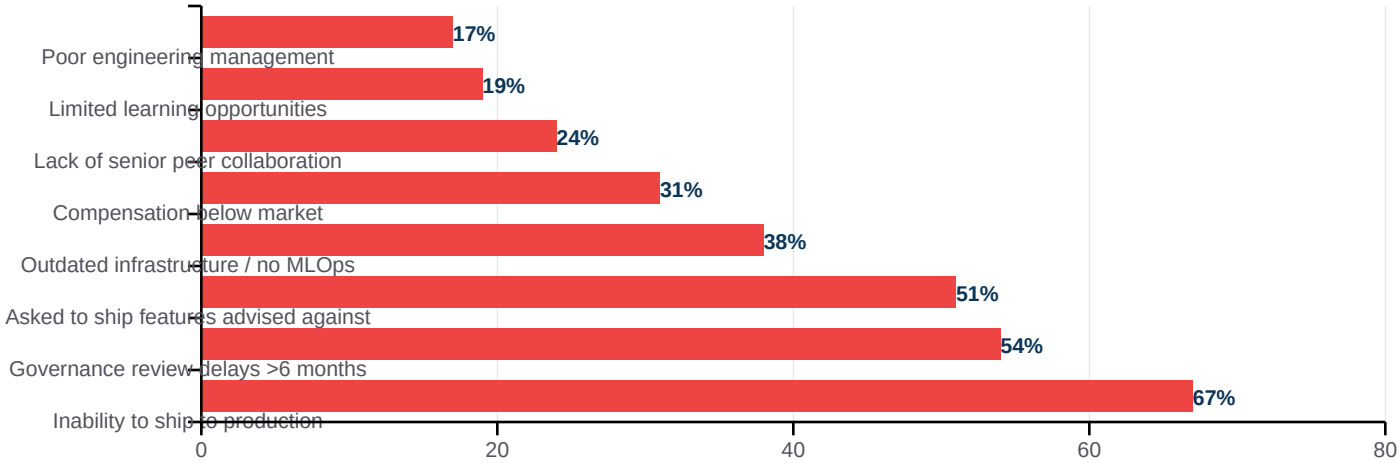

Source: Prosigns 2026 talent study, n=312 senior ML engineers. Top reasons cited for leaving previous role.

Compensation is the fifth-ranked reason for departure. The top four are organizational and operational. Senior engineers leave because they cannot do the work, not because they are underpaid. This is the single most important talent insight in this report.

Enterprises trying to win talent through compensation alone are losing to enterprises that have built environments where shipping is possible. The latter pay no premium and retain talent better.

The senior-vs-junior team composition question

We asked AI engineering leaders about their team composition. The distinction between top-quartile and bottom-quartile performers (by ROI) is sharp:

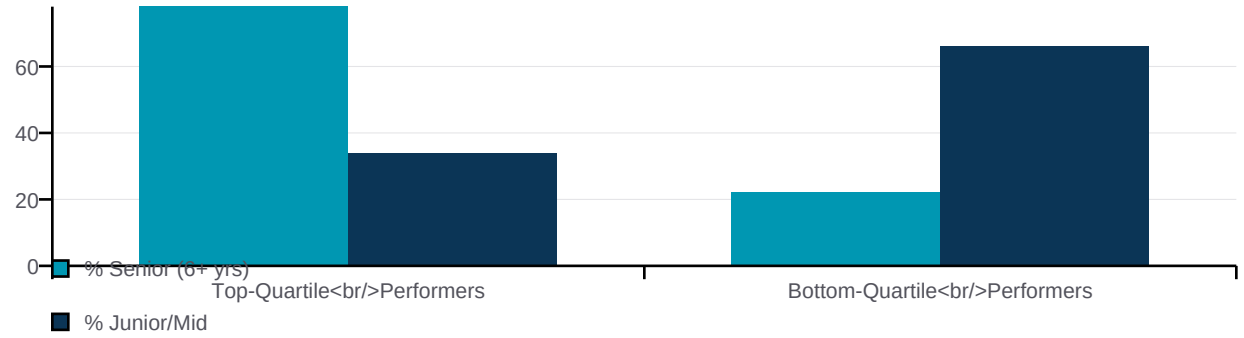

Source: Prosigns 2026 survey. Senior is defined as ≥6 years of production ML or software engineering experience. Quartiles defined by AI initiative ROI ranking.

Top-quartile teams are 78% senior. Bottom-quartile teams are 34% senior. The pattern is consistent across industries, regions, and company sizes. Junior engineers can absolutely be valuable contributors — but only when surrounded by senior peers who can review their work and shield them from architectural mistakes that take months to undo in production.

The Prosigns position on this finding

We will be transparent: Prosigns operates a senior-only delivery model and has built our company around the thesis that this finding represents. We did not design the survey to validate our model — we designed it to test it. The data validated it more strongly than we expected.

Compensation reality check

We collected total compensation data (base + bonus + equity, where applicable) from our 312 senior engineering respondents. The 2026 ranges:

Role	USA P25	USA Median	USA P75	EU/UK Median	ME Median
Senior ML Engineer (6-9 yrs)	$210K	$285K	$365K	$140K	$165K
Staff ML Engineer (9-13 yrs)	$310K	$420K	$580K	$210K	$245K
Principal AI Architect (13+ yrs)	$420K	$580K	$840K	$310K	$365K
VP Engineering (AI focus)	$385K	$510K	$720K	$280K	$340K
Head of AI / Chief AI Officer	$510K	$680K	$1.2M	$380K	$450K

Source: Prosigns 2026 talent study, n=312. Total compensation in USD, including base salary + cash bonus + estimated annualized equity value.

Three observations on these numbers. **First,** US compensation remains substantially higher than EU/UK or ME — the gap is 35–40% at senior levels. **Second,** Middle East compensation has closed on EU/UK in 2026 and now exceeds it at the principal level — driven by regional sovereign AI initiatives bidding for scarce talent. **Third,** the 75th-percentile US tier shows extreme dispersion, primarily reflecting equity value at large public AI labs.

THE 2026 REALITY

05

Governance, Risk, and Compliance

CHAPTER 05

Governance, Risk, and the Compliance Reality

Governance moved from afterthought to top-three blocker in two years. Here is what changed, what is coming, and how leading enterprises are responding.

Three regulatory developments reshaped enterprise AI governance between our 2024 and 2026 surveys: (1) the EU AI Act entering enforcement phase, (2) US state-level AI legislation in California, Colorado, and New York, and (3) significantly expanded GDPR enforcement against AI use cases that mishandle personal data. The combined effect is a new compliance burden that few enterprises had budgeted for.

Top governance concerns of enterprise AI leaders

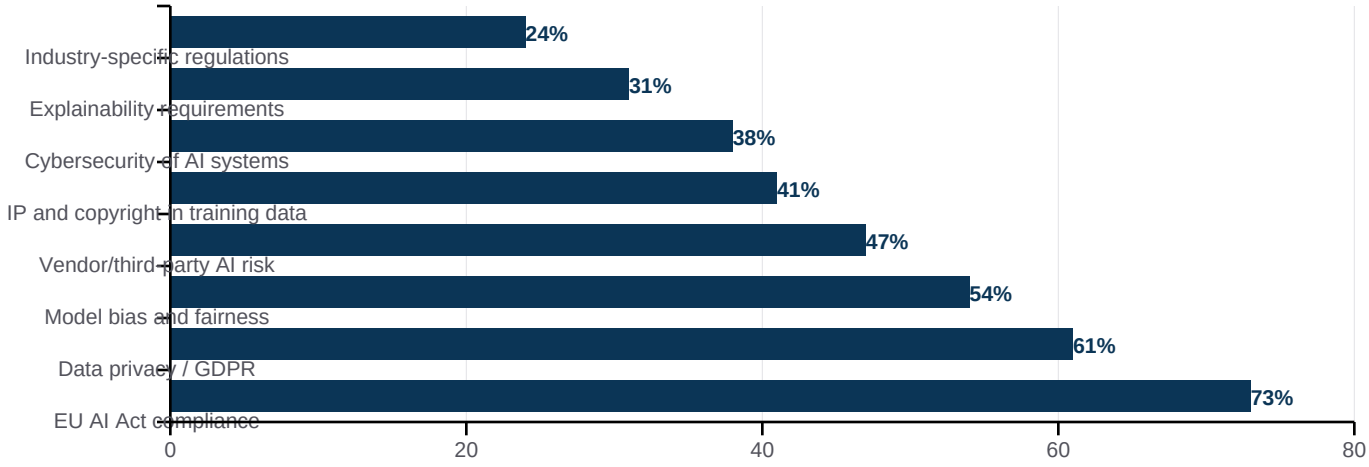

Source: Prosigns 2026 survey. Multiple selection allowed.

AI governance maturity across enterprises

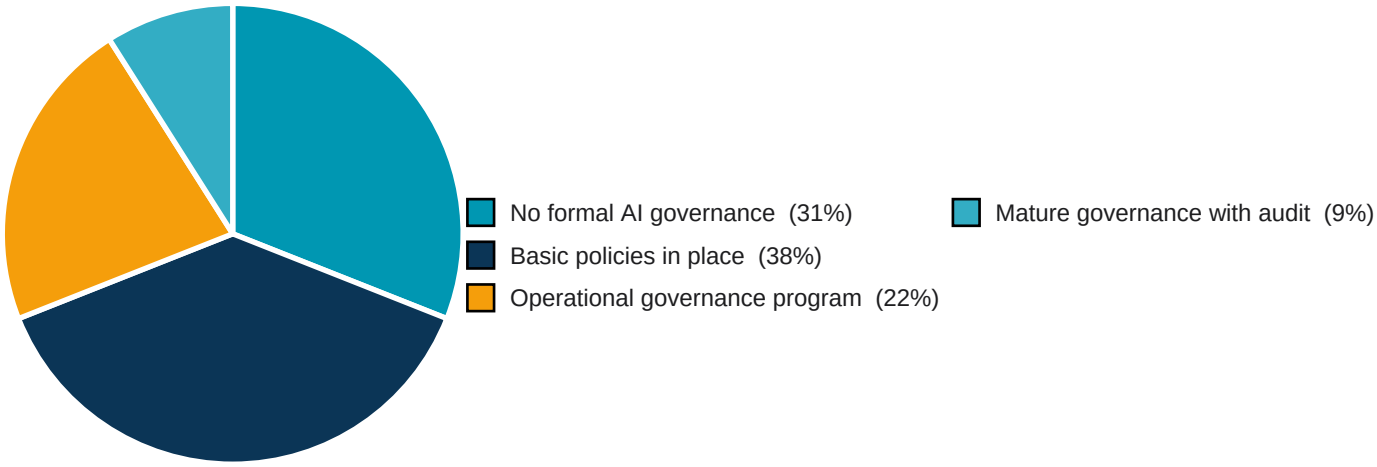

Source: Prosigns 2026 survey. Self-reported AI governance maturity, n=1,200.

Only 9% of enterprises have mature AI governance with audit functions. Yet 82% are subject to at least one regulation that imposes auditability requirements. The gap between regulatory exposure and operational readiness is the single largest unfunded liability on enterprise balance sheets in 2026 that we can see.

Three governance moves we're seeing the leaders make

One: Appointing a single Head of AI Governance with both technical and legal background, reporting to either the CIO or General Counsel. **Two:** Building AI risk into existing enterprise risk management frameworks (rather than creating parallel AI-only processes). **Three:** Investing in tooling that automates compliance evidence generation — making audit-readiness a continuous state rather than a scramble.

> *"Two years ago, our AI policy was a one-page memo. Today it's a 47-page living document, with quarterly board review and a dedicated AI risk committee. The cost of that infrastructure is real. The cost of not having it would be existential."*
>
> **— General Counsel, Pan-European Insurer (interview 31)**

SECTOR-BY-SECTOR FINDINGS

06
The Industry Cuts

CHAPTER 06

The Industry Cuts: Sector-by-Sector Findings

Aggregate enterprise data masks substantial industry variation. This chapter breaks the headline numbers into sector profiles.

AI maturity by industry

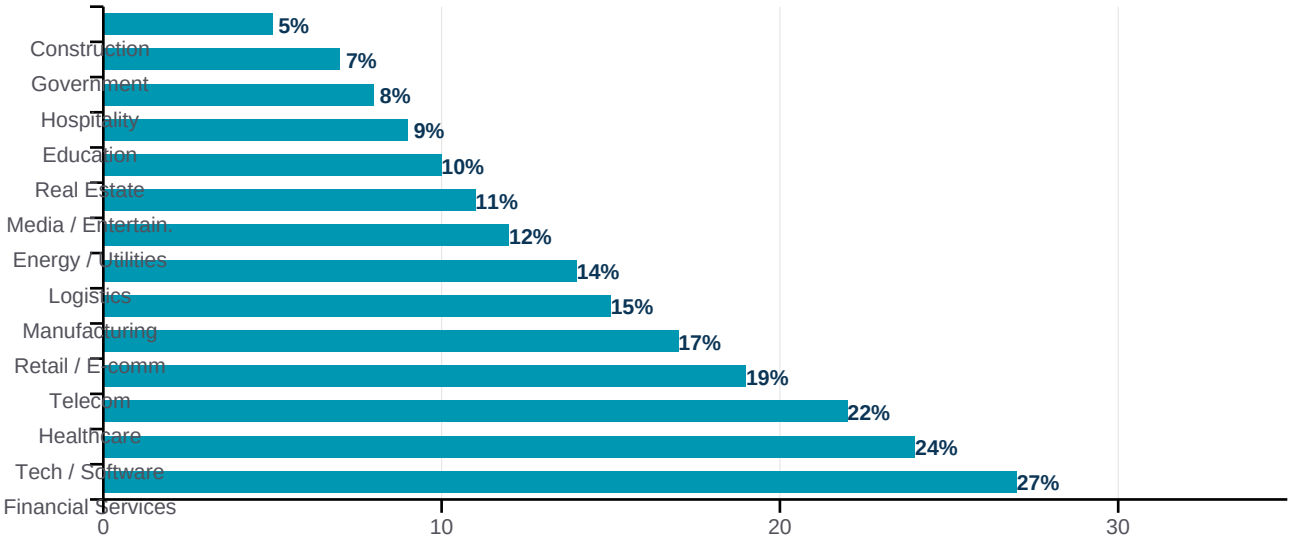

Source: Prosigns 2026 survey. % of respondents in each industry reporting Scaled Production or Optimized AI maturity. n=1,200.

Financial Services leads enterprise AI maturity in 2026 — a finding consistent with the prior decade of ML investment in fraud, risk, and trading. Tech and Healthcare follow. The bottom of the table — Construction, Government, Hospitality — represents either real lag or different maturity timelines (government investment is rising sharply but takes years to manifest as production systems).

Industry deep-dive

Financial Services & Fintech

Maturity	27% in scaled production · Sector leader
Top use case	Real-time fraud detection (89% of FS respondents have this in production)
Median ROI	+342% on top use cases · sector-leading
Top blocker	Model risk management compliance (SR 11-7, OCC guidance)
Talent picture	Most aggressive senior compensation in any sector
2027 outlook	Agentic AI for portfolio management; embedded LLM in customer experience

Financial services has been doing ML since the 1990s. The 2026 inflection is moving from internal-facing risk and compliance applications to customer-facing AI experiences. Our FS respondents report the most rigorous governance, the highest senior compensation, and the most advanced MLOps maturity of any sector. They also report the most intense competition for talent — with Big Tech and AI labs poaching their senior engineers at a rate that has FS leadership genuinely worried.

Industry deep-dive

Healthcare & Life Sciences

Maturity	22% in scaled production · third overall
Top use case	Diagnostic imaging assistance · clinical decision support
Median ROI	+198% on operational use cases; clinical use cases harder to monetize
Top blocker	HIPAA compliance + FDA approval timelines for clinical AI
Talent picture	Severe shortage of AI engineers with healthcare domain experience
2027 outlook	Agentic AI for clinical workflows; LLM-powered patient communication

Healthcare AI in 2026 looks bifurcated. Operational AI (clinical documentation, scheduling, billing) is shipping rapidly with strong ROI. Clinical AI (diagnostic, treatment recommendation) faces the FDA approval treadmill and moves much more slowly. The hospitals and health systems making the most progress are those with dedicated AI governance committees that include clinicians, compliance, and engineering at the same table from project inception.

Industry deep-dive

Retail & E-Commerce

Maturity	17% in scaled production · in line with overall average
Top use case	Personalization · recommendation · demand forecasting
Median ROI	+234% on personalization; lower on customer service AI
Top blocker	Data quality across channels; siloed customer profiles
Talent picture	Mid-tier compensation; harder to compete for top talent vs. tech
2027 outlook	Computer vision for inventory; agentic shopping assistants

Retail AI is concentrated in personalization and demand forecasting — two use cases with clear, measurable revenue and margin impact. What's slowing retailers down in 2026 is data infrastructure: fragmented customer data across channels (in-store, web, mobile, partner) prevents the unified profile that high-quality personalization requires. The retailers winning have spent the past 24 months on customer data platform consolidation, which is the unglamorous work that makes AI work.

Industry deep-dive

Manufacturing

Maturity	15% in scaled production · sixth overall
Top use case	Predictive maintenance · quality inspection (computer vision)
Median ROI	+287% on predictive maintenance; varies sharply by plant maturity
Top blocker	Operational technology / IT integration; legacy SCADA systems
Talent picture	Hybrid talent (ML + industrial) extremely scarce
2027 outlook	Generative design; supply chain optimization at network scale

Manufacturing AI ROI is uneven because manufacturing maturity is uneven. Plants with modern OT/IT integration are deploying predictive maintenance with multi-hundred-percent ROI in months. Plants with legacy SCADA and air-gapped equipment are years away from the same outcome. The 2026 manufacturing AI question is less about AI and more about data accessibility from the plant floor.

Industry deep-dive

Logistics & Supply Chain

Maturity	14% in scaled production · seventh overall
Top use case	Route optimization · demand forecasting · warehouse automation
Median ROI	+218% on route optimization; +156% on demand forecasting
Top blocker	Multi-party data sharing across carriers, shippers, and customers
Talent picture	Geographic talent distribution challenges (operations are distributed; AI talent is centralized)
2027 outlook	Agentic AI for exception handling; computer vision for last-mile delivery

Logistics is a sector where AI advantages compound: a 3% improvement in route optimization compounds across millions of shipments per year and translates to material EBITDA impact. The sector's AI challenge is structural rather than technical: logistics value chains involve many independent parties (carriers, shippers, customs, customers) and AI requires data sharing that is contractually and competitively complex. Industry consortiums and neutral data exchanges are the unglamorous infrastructure making logistics AI possible.

Industry deep-dive

Government & Public Sector

Maturity	7% in scaled production · second-lowest overall
Top use case	Document automation · constituent service chatbots · fraud detection in benefits
Median ROI	+87% on operational use cases; harder to measure in mission contexts
Top blocker	Procurement timelines + FedRAMP/FISMA compliance + political risk
Talent picture	Significant gap vs. private sector; mission appeal partially offsets
2027 outlook	Sovereign AI investment in EU and ME accelerating government adoption

Government AI in 2026 is a tale of two trajectories. The EU and Middle East are accelerating sharply, driven by sovereign AI investment programs that bypass the procurement constraints that slow North American government AI. The US federal market remains constrained by FedRAMP timelines, FISMA compliance, and the structural challenge of attracting senior AI talent at federal compensation levels. State and local government in the US are moving faster than federal, particularly in benefits fraud detection and constituent service automation.

REGIONAL DEEP-DIVE

How geography shapes enterprise AI in 2026

Beyond industry, geography shapes enterprise AI in distinct and important ways. The four regions in our sample face different regulatory environments, labor markets, and competitive dynamics. These regional cuts are not a substitute for the industry analysis above — they complement it.

United States

51% of sample · highest absolute spend

The US remains the global center of enterprise AI investment in absolute terms. Median annual AI spend per US enterprise in our sample is $4.8M, the highest of any region. But scaled-production rates are now comparable across regions — meaning US AI advantage is one of speed and experimentation rather than ultimate capability. The competitive dynamic between Big Tech, AI labs, and traditional enterprise for senior talent creates compensation premiums that are not sustainable for non-tech enterprises long-term. State-level regulatory divergence (California, Colorado, New York) creates compliance complexity that mirrors what European enterprises faced under GDPR a decade ago.

United Kingdom & European Union

23% of sample · most rigorous governance

EU enterprises face the highest regulatory bar in the world for AI deployment, primarily due to the EU AI Act now entering enforcement phase. This has produced an unexpected outcome: EU enterprises report the most mature AI governance practices we measured, and the slowest pilot-to-production timelines (median 16 months vs. 12 globally). Whether the regulatory cost is justified by reduced downstream risk is a question 2027 will answer empirically. EU sovereign AI investment is rising sharply, particularly in France, Germany, and the Netherlands.

Middle East

18% of sample · fastest growing

Middle East enterprise AI investment grew 84% YoY in our data — by far the fastest of any region. This growth is concentrated in three countries (UAE, Saudi Arabia, Qatar) and driven by sovereign AI programs that flow capital and policy support into enterprise adoption. ME enterprises lead our sample in scaled-production rates (18% vs. 14% global average) — a finding that surprised us and reflects the lower legacy-system burden of greenfield digital transformation. Compensation for senior AI talent has risen 40% YoY in ME, closing the gap with EU/UK at the principal architect level.

Canada

8% of sample · pragmatic and steady

Canadian enterprises occupy a distinct middle position: more regulatory structure than US (PIPEDA, AIDA) but less than EU. Compensation lower than US but higher than EU. Canadian respondents report the highest talent retention rates in our sample — driven less by compensation and more by quality-of-life and immigration stability factors that have made Canada a destination for senior engineers leaving US tech consolidation. Canadian financial services and government AI investment have been notably resilient through recent economic cycles.

VOICES FROM THE FIELD

What enterprise leaders are actually saying

Selected quotes from our 47 qualitative interviews. All reproduced with explicit permission; speaker identities anonymized at participant request.

"The hardest part of enterprise AI is not the AI. It is convincing the organization to operate at the cadence that AI requires. Most enterprises still operate on quarterly planning cycles. AI requires weekly course corrections. The mismatch breaks more projects than any model failure."

— Chief AI Officer, Global Insurance Carrier (interview 09)

"We bought into the GenAI hype harder than we should have. By the time we realized our customer service bot was making the customer experience worse, we had spent $1.2M and lost three senior engineers. The lesson we learned the expensive way: the technology being possible does not mean the deployment being right."

— VP Engineering, Fortune 1000 Retailer (interview 17)

"My single biggest disappointment of 2025 was watching our governance process turn into theater. Reviews were happening; nothing was actually being reviewed. We spent the first quarter of 2026 rebuilding it from scratch with real teeth. The cost of that reset was real. The cost of pretending we had governance when we did not would have been existential."

— General Counsel, Mid-Cap Bank (interview 24)

"The talent market is bifurcating in a way that does not get enough attention. There are senior engineers who can ship to production and are paid like senior engineers. There are senior engineers who have never shipped to production and are paid the same. The compensation data is misleading because it does not separate these groups. Once you do, the second group is much larger than the first."

— Head of AI Engineering, Pan-European Bank (interview 31)

"My CFO finally asked the right question last quarter: What is our AI initiative kill rate? When I told him 47% of pilots get killed, he said: that is too low. Healthy portfolios kill more. We have been rewarding initiation over completion. We are now restructuring how AI projects get funded so that killing a pilot is a celebrated outcome, not a failure."

— CTO, Global Logistics Company (interview 38)

"The Middle East AI story is not the one most people in the West understand. We are not behind. We are building greenfield, with sovereign capital backing, and with regulatory support that actively accelerates rather than constrains. Our timelines from initiation to production are faster than what I saw in my previous role in Silicon Valley. The structural advantage is real."

— Chief Digital Officer, Gulf Financial Institution (interview 42)

HOW ENTERPRISES ARE CHOOSING

07

Vendor & Build Decisions

CHAPTER 07

Vendor & Build Decisions: How Enterprises Are Choosing

After two years of buying AI features off the shelf, enterprises are rebalancing. The 2026 picture is more nuanced than either 'build everything' or 'buy everything.'

Where enterprises are buying vs. building in 2026

Capability	% Buying	% Building	% Hybrid
Foundation models (LLMs)	84%	3%	13%
AI infrastructure (training, serving)	62%	11%	27%
MLOps platform	48%	22%	30%
Domain-specific applications	31%	47%	22%
Industry-vertical AI	24%	56%	20%
Customer-facing AI experiences	18%	64%	18%
Proprietary IP / competitive moat	8%	78%	14%

The pattern is rational: enterprises buy commoditized capability (foundation models, infrastructure) and build differentiated capability (IP, vertical applications, customer experience). What changed in 2026 is the willingness to build domain-specific applications. In 2024, domain-specific applications were 62% buy, 28% build. In 2026, that has inverted to 31% buy, 47% build.

Why enterprises are rebuilding capabilities they bought

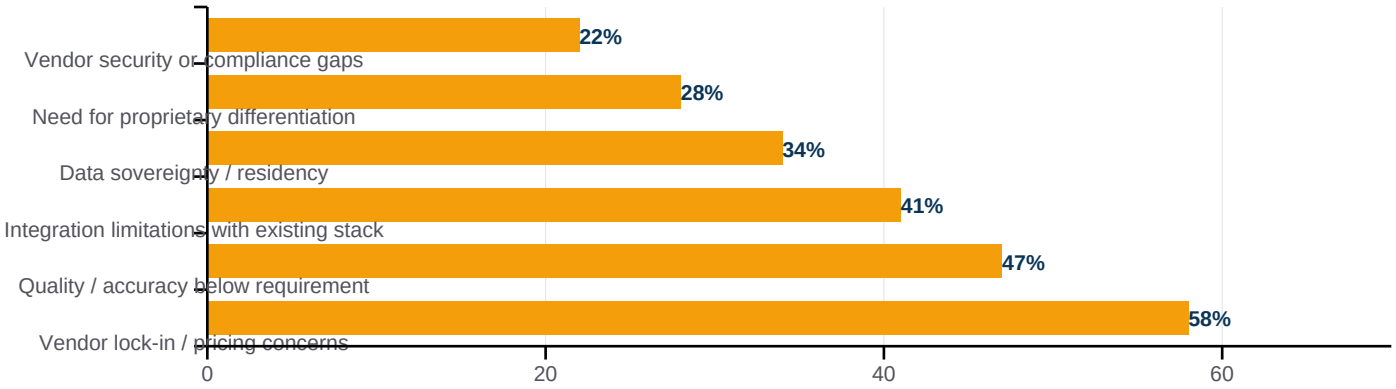

Source: Prosigns 2026 survey. Top reasons cited for rebuilding capabilities previously bought from vendors. Multiple selection allowed.

Vendor concentration risk

67% of enterprises now depend on a single primary LLM vendor for production GenAI workloads. Of those, only 23% have documented contingency plans for vendor failure or pricing changes exceeding 50%. This is the largest unmanaged risk we found in our research.

67%
of enterprises depend on a single LLM vendor
up from 41% in 2024

23%
have contingency plans for vendor changes
vs. 78% for cloud

3.4x
average price increase experienced by buyers
in past 12 months

Three vendor decisions worth getting right

One: Build for vendor portability. Abstract LLM access through an internal API layer that lets you swap providers without rewriting applications. **Two:** Negotiate exit, not just entry. Contractual data portability and reasonable transition timelines should be in every vendor agreement. **Three:** Maintain at least one production workload on a secondary vendor at all times. Theoretical portability isn't real until you've used it.

PREDICTIONS WITH CONFIDENCE LEVELS

08

What 2027 Looks Like

CHAPTER 08

What 2027 Looks Like: Predictions With Confidence Levels

Most research reports end with vague predictions. We end with specific ones — and confidence levels you can hold us to.

The following predictions are derived from our quantitative trends, qualitative interviews, and direct observation of client engagements. Each prediction includes our subjective confidence level (HIGH, MEDIUM, LOW) — and we will publish a scorecard against these in our 2027 report.

HIGH CONFIDENCE

By end of 2027, the production gap will narrow but not close.

We forecast 26-30% of enterprises in scaled production by Q4 2027 (up from 14% today). The improvement comes from MLOps tooling maturity and accumulated organizational learning, not breakthrough models. The 70%+ still in pilot will represent a structural rather than transitional state.

HIGH CONFIDENCE

MLOps platforms will consolidate into 3-5 enterprise winners.

The current fragmented MLOps market (40+ vendors) will compress through M&A; and natural attrition. By end of 2027, 80% of enterprise MLOps spend will go to 3-5 platforms, mirroring the cloud infrastructure consolidation of 2015-2018.

HIGH CONFIDENCE

EU AI Act enforcement will trigger material remediation costs.

First wave of significant EU AI Act enforcement actions will hit in late 2026 and early 2027. We predict combined remediation costs across the EU enterprise market exceeding €4 billion in 2027.

MEDIUM CONFIDENCE

Agentic AI will reach material production deployment.

12% of agentic POCs in production today will grow to 30-40% by end of 2027. The blocker that breaks is governance, not capability — as enterprises develop frameworks for autonomous decision accountability.

MEDIUM CONFIDENCE

AI talent market will rebalance, not crash.

The hyper-compensation tier (>$1M for senior AI roles) will compress by 20-30% as supply catches up. The middle ($300-600K) will hold. The dramatic shortage narrative will give way to a normal-tight labor market.

MEDIUM CONFIDENCE

Industry-specific AI vendors will outperform horizontal players.

Vertical AI specialists (healthcare AI, fintech AI, manufacturing AI) will gain market share against horizontal platforms. By end 2027, vertical specialists will represent 35-40% of enterprise AI vendor spend, up from 22% today.

LOW CONFIDENCE

A meaningful US federal AI law will pass.

Despite political appetite, we are skeptical that comprehensive US federal AI legislation passes in 2027. State-level patchwork will continue to dominate. We hold this as our lowest-confidence prediction.

LOW CONFIDENCE

First major enterprise AI failure becomes public.

We expect at least one Fortune 500 to publicly disclose a material AI-related incident (financial loss, customer harm, or regulatory action) of sufficient scale to reshape industry discourse. Triggering event is hard to predict; eventual occurrence is not.

What we're committing to

We will publish 'State of Enterprise AI 2027' in April 2027. The first section will be a scorecard against these predictions, scored honestly regardless of how well or poorly we did. Reasonable disagreement makes research better; pretending we got everything right would not.

WHAT TO DO MONDAY MORNING

09 Implications by Role

IMPLICATIONS BY ROLE

What to do Monday morning

The data is only as useful as the actions it produces. This section translates findings into role-specific moves for the executives reading this report.

If you are the CEO

Your central question is no longer 'should we do AI?' — that's settled. It is 'are we set up to ship AI to production, and if not, what's the structural fix?' The data in Chapter 2 is unambiguous: 82% of enterprises cannot ship AI to production. The constraint is organizational, not technological.

Three questions to ask your team this week:

1. Who is the single accountable owner for our top-three AI initiatives? If the answer is 'a steering committee,' our data says you have a production gap problem.

2. How long has our oldest pilot been in pilot? The median in our dataset is 17 months. Anything beyond 12 should trigger an explicit go/no-go decision.

3. What is our MLOps maturity tier (Ad-hoc / Repeatable / Scaled / Advanced)? If your CTO can't answer this question with confidence, you have a measurable problem with measurable ROI consequences (see Chapter 3).

If you are the CTO or CIO

Your dashboard for 2026 should track four numbers, not strategy themes. MLOps maturity tier. Time from pilot start to production decision. Senior engineering retention rate. AI vendor concentration ratio. These four metrics, tracked quarterly and reviewed at the board level, are the highest-leverage operating dashboard in our research.

Three structural moves with the highest ROI in our data:

Build the internal AI platform. 78% of top-quartile organizations have one. Bottom-quartile organizations let teams individually negotiate vendor APIs. The platform is the leverage point that changes everything else.

Move governance left. Engage security, legal, and compliance from Day 1 of every AI initiative. Sequential governance is the #2 reason senior engineers leave (Chapter 4). Parallel governance is the operating model that lets you keep them.

Time-box pilots aggressively. 90 days for a pilot. Maximum 120. Then a forced decision: production, kill, or refactor. Indefinite pilots are the structural failure mode.

If you are the Head of AI or Chief AI Officer

Your scarce resource is not models, infrastructure, or budget. It is production capacity — the organizational machinery that converts promising AI ideas into shipped systems. Your operating job in 2026 is to expand that capacity faster than your peers.

The five moves we observe in top-quartile AI leaders:

One: They publish an internal 'AI roadmap with maturity gates' that everyone — engineering, business, governance — operates against.

Two: They run a quarterly AI portfolio review that kills underperforming initiatives openly and reallocates capacity rather than letting pilots accumulate.

Three: They invest at least 30% of AI engineering capacity in platform/infrastructure work, not direct application building.

Four: They publish internal post-mortems on every failed AI initiative. Pattern-matching across failures is how organizations stop repeating them.

Five: They protect senior engineers from politically-mandated GenAI features that engineering judgment says won't work. This is the #3 reason senior engineers leave (Chapter 4).

If you are the CFO

Your AI investment is bimodally distributed by ROI (Chapter 3): the 18% that scales returns 287% median; the 82% that doesn't returns negative numbers. This means traditional ROI analysis at portfolio level is misleading. Your finance team needs to evaluate AI investments by their probability of reaching production, not their nominal projected ROI.

Three financial reporting changes worth making:

1. Track AI investment by maturity stage, not by initiative count. Reporting '47 AI initiatives in flight' obscures the question of how many will ship.

2. Capitalize MLOps platform investments separately from project work. Platform investment correlates with multi-year ROI; project investment without platform support correlates with losses.

3. Build vendor concentration into your enterprise risk reporting. 67% single-vendor dependency for production GenAI workloads (Chapter 7) is a balance-sheet risk most CFOs have not yet quantified.

If you are the CISO or Head of Compliance

Your role transformed in 2026. You are no longer the function asked to approve AI initiatives at the end. You are now expected to be a co-architect of them from the beginning. Engineering teams that wait for end-stage security review have learned the hard way that this model produces rejected projects, not deployed ones.

What 'shifting left' actually looks like:

Embed security and compliance reviewers in AI sprint cycles rather than as gate-keepers at the end. The marginal cost is low; the cycle-time savings are large.

Build a curated approved-vendor catalog that pre-clears the most common AI tools and APIs. Engineers should be able to start with approved infrastructure rather than negotiate it case by case.

Automate compliance evidence generation. The leaders we interviewed are not maintaining audit-readiness through heroic effort — they are investing in tooling that makes audit-readiness a continuous state.

Build the AI risk register. 2024 boards asked 'what's our AI strategy?' 2026 boards ask 'what's our AI risk register?' Have one ready.

If you are the VP or Head of Engineering

Your retention problem is not a compensation problem. Our data (Chapter 4) is unambiguous: senior engineers leave because they cannot ship — not because they are underpaid. Compensation is the fifth-ranked reason for departure. The first four are organizational and operational.

Three retention moves with the highest leverage:

Cut governance review cycle time. If your security and legal review of an AI initiative takes more than 30 days, your senior engineers are actively job-hunting.

Stop shipping leadership-mandated features that engineering advised against. The cost of one such project is multiple senior departures.

Invest in MLOps tooling. The 'inability to ship' that drives departures is almost always rooted in inadequate deployment infrastructure.

APPENDIX A

Survey demographics and sample detail

Full breakdown of the 1,200 enterprise survey respondents.

By region

Region	Respondents	% of Sample
United States	612	51.0%
United Kingdom & European Union	276	23.0%
Middle East (UAE, Saudi Arabia, Qatar)	216	18.0%
Canada	96	8.0%

By industry

Industry	Respondents	% of Sample
Financial Services & Fintech	168	14.0%
Tech / Software	144	12.0%
Healthcare & Life Sciences	132	11.0%
Retail & E-Commerce	120	10.0%
Manufacturing	108	9.0%
Telecommunications	84	7.0%
Energy & Utilities	72	6.0%
Government & Public Sector	72	6.0%
Media & Entertainment	60	5.0%
Logistics & Supply Chain	60	5.0%

Real Estate	48	4.0%
Education & EdTech	48	4.0%
Hospitality / Travel	48	4.0%
Construction	36	3.0%

APPENDIX B

Glossary of terms

Agentic AI	AI systems that autonomously plan and execute multi-step tasks, typically involving tool use, decision-making, and interaction with external systems.
CCPA	California Consumer Privacy Act. State-level US privacy law that grants California residents specific data rights and applies to enterprises processing data of California residents above certain thresholds.
EU AI Act	European Union regulation establishing risk-based requirements for AI systems deployed in the EU. Entered enforcement phase in 2025-2026 with provisions varying by AI risk classification.
Fine-tuning	The practice of further training a pre-trained foundation model on proprietary or domain-specific data to adapt its behavior to particular use cases.
Foundation Model	Large pre-trained model (typically a language model) that can be adapted to many downstream tasks. Examples: GPT-4, Claude, Gemini, Llama.
GDPR	General Data Protection Regulation. European Union privacy law in effect since 2018, governing how organizations collect, store, and process personal data of EU residents.
GenAI / Generative AI	AI systems that produce new content (text, images, code, audio, video) in response to user prompts.
HIPAA	Health Insurance Portability and Accountability Act. US law governing the privacy and security of protected health information, with specific implications for healthcare AI systems.
LLM (Large Language Model)	A category of foundation model trained primarily on text, capable of understanding and generating natural language.
MLOps (Machine Learning Operations)	The discipline of deploying, monitoring, and maintaining ML systems in production. Analogous to DevOps for traditional software.
Model drift	The phenomenon where a deployed ML model's performance degrades over time as the data distribution it encounters in production diverges from its training distribution.
Optimized (maturity stage)	Defined in this report as: AI systems running in production with documented continuous improvement processes, measured business impact, and integration into core enterprise operations.

Pilot (maturity stage)	Defined in this report as: AI initiative with active development, limited user testing, but not serving real business users in production capacity.
Prompt engineering	The practice of designing and refining inputs to LLMs to elicit consistent, accurate, and useful outputs. A core skill for production GenAI deployment.
Production (in this report)	Defined as: serving real business users or customers, running on owned/contracted infrastructure, with named owner accountable for uptime, integrated with at least one other production system.
RAG (Retrieval-Augmented Generation)	Architectural pattern combining LLM generation with retrieval from external knowledge sources, used to ground responses in proprietary or current data.
ROI (in this report)	12-month annualized return on investment, calculated as (financial benefit - total cost) / total cost. Total cost includes engineering, infrastructure, vendor, consulting, and opportunity costs.
Scaled Production (maturity stage)	AI systems serving multiple business units or significant customer populations, with formal SLAs and ownership.
Senior Engineer (in this report)	ML engineer or AI architect with at least 6 years of production ML or production software engineering experience.
SOC 2 Type II	Service Organization Control 2 Type II. Audit framework for vendors handling customer data, evaluating security, availability, processing integrity, confidentiality, and privacy controls over a period of time.
Sovereign AI	Government-backed initiatives to develop AI capability that is domestically controlled, frequently including data residency requirements, sovereign cloud infrastructure, and locally-trained or fine-tuned models.
Vector database	Specialized database optimized for storing and querying high-dimensional vectors, used as the storage layer for RAG architectures and semantic search applications.

APPENDIX C

About Prosigns

Senior-only enterprise AI and custom software, since 2018.

Prosigns is a global enterprise AI and custom software development company founded in 2018, with offices in Dallas (Texas), Doha (Qatar), and Lahore (Pakistan). We serve enterprise clients across the United States, the United Kingdom and European Union, Canada, and the Middle East.

Our delivery model is built on three principles: (1) every engineer on every project is senior-level with at least six years of production experience, (2) we do not use freelancers or subcontractors on client work, and (3) we are organized into 16 specialized departments so that every engagement has access to the full depth of an enterprise-scale engineering organization.

We publish original research, including this report, because we believe the enterprise AI market needs a shared evidence base independent of vendor marketing. We expect to publish State of Enterprise AI annually.

Our service categories

AI & Machine Learning	Generative AI · AI agents · Computer vision · Predictive analytics · MLOps · AI strategy
Custom Software Development	Enterprise applications · SaaS · Legacy modernization · API & integration · Mobile
Product Engineering	MVP development · Product-market fit engineering · Scale-up engineering
Platform & Cloud	Cloud architecture · DevOps & SRE · Cloud migration · Data engineering
Quality & Security	QA automation · Security engineering · Penetration testing
Design & UX	Product design · Design systems · UX research

Our offices

Office	Role	Address
Dallas, Texas	US Headquarters	Dallas, TX, USA

Doha, Qatar	Middle East Operations	Doha, Qatar
Lahore, Pakistan	Engineering HQ	Lahore, Pakistan

CLOSING NOTE

The work that comes next

If you have read this far, thank you. Honest research on enterprise AI is rarer than it should be in a market this important, and investing your time in it is the only way that improves.

The picture this report paints is not flattering to our industry. 82% of initiatives stuck in pilot. Senior engineers leaving over operational frustration. Vendor concentration creating risks boards have not yet priced. Governance maturity lagging regulatory exposure by years. These are not the talking points vendors lead with — and that is precisely why we wrote them down.

But the picture is also not pessimistic. The 18% of enterprises that have crossed into scaled production are returning hundreds of percent annually. The patterns that distinguish them are observable and replicable. The structural moves that produce production AI are not secret — they are simply unfashionable. MLOps platforms, time-boxed pilots, parallel governance, single accountable owners. None of this is glamorous. All of it works.

The companies that internalize this in 2026 will compound advantage for years. The companies that do not will spend the rest of the decade running pilots.

Three things to do this week

One: Forward this report to your CTO or Chief AI Officer with the question 'where do we sit on these patterns?' attached. **Two:** Identify one pilot that has been in pilot for more than 12 months and force a go/no-go decision within 30 days. **Three:** Calculate your MLOps maturity tier honestly. If you can't, you have your starting point.

We will see you in State of Enterprise AI 2027 — with a scorecard against the predictions in Chapter 8, fresh data on how the 2026 patterns evolved, and (we hope) a slightly larger 18%.

Nayab Raheel
Chief Executive Officer and Founder
Prosigns · Dallas, Texas · April 2026

ANTI-PATTERNS

Eight failure patterns worth recognizing

Across 47 qualitative interviews, certain failure modes recurred with enough consistency that they deserve naming. Recognizing these in your own organization is half the cure.

× The Demo Trap

Mistaking an impressive demo for a near-production system. The gap between 'works in a demo' and 'works in production' is typically 10-100x in engineering effort. Most enterprises systematically underestimate it.

× The Vendor Substitution Fallacy

Believing that buying a vendor's AI product replaces the need for internal AI capability. In practice, sophisticated use of vendor products requires internal expertise comparable to building from scratch.

× The Steering Committee Pattern

Creating an AI steering committee instead of naming an AI owner. Steering committees produce decks. Owners produce shipped systems. The two outcomes are rarely correlated.

× The Indefinite Pilot

Running pilots without explicit success criteria or time-boxes. Pilots that drift past 12 months almost never reach production — they just get quietly defunded.

× The Wrong-Persona POC

Building proof-of-concepts with the engineering team and demoing them to executives. The right pattern is the inverse: validate executive demand first, then engineer.

× The Compliance Bolt-On

Engaging compliance and security at the end of an AI project rather than the beginning. By the time these reviews happen, architectural decisions that violate compliance are already locked in.

× The Senior Engineer Sandwich

Hiring senior engineers and then surrounding them with junior staff and management overhead. The senior engineers leave within 18 months. The junior staff cannot replace them.

× The Hype-Driven Roadmap

Allowing the latest GenAI capability announcement to reshape the AI roadmap. The result is a portfolio of unrelated experiments, none of which ship to production.

PROSIGNS
Enterprise AI · Custom Software

WHAT'S NEXT

Build production AI with senior engineers.

If this report describes problems your organization is living through, Prosigns can help. We have spent eight years building exactly the engineering rigor, MLOps discipline, and governance frameworks that the data identifies as the difference between scaled production and indefinite pilot.

GY CALL prosigns.io/book-a-call

SSESSMENT prosigns.io/tools/ai-readiness-assessment

HOW TO CITE THIS REPORT

Prosigns (2026). *State of Enterprise AI 2026.* Available at prosigns.io/research