2026 AI GOVERNANCE BENCHMARK REPORT

The AI Portfolio Explosion

When Activity Creates the Illusion of AI Value

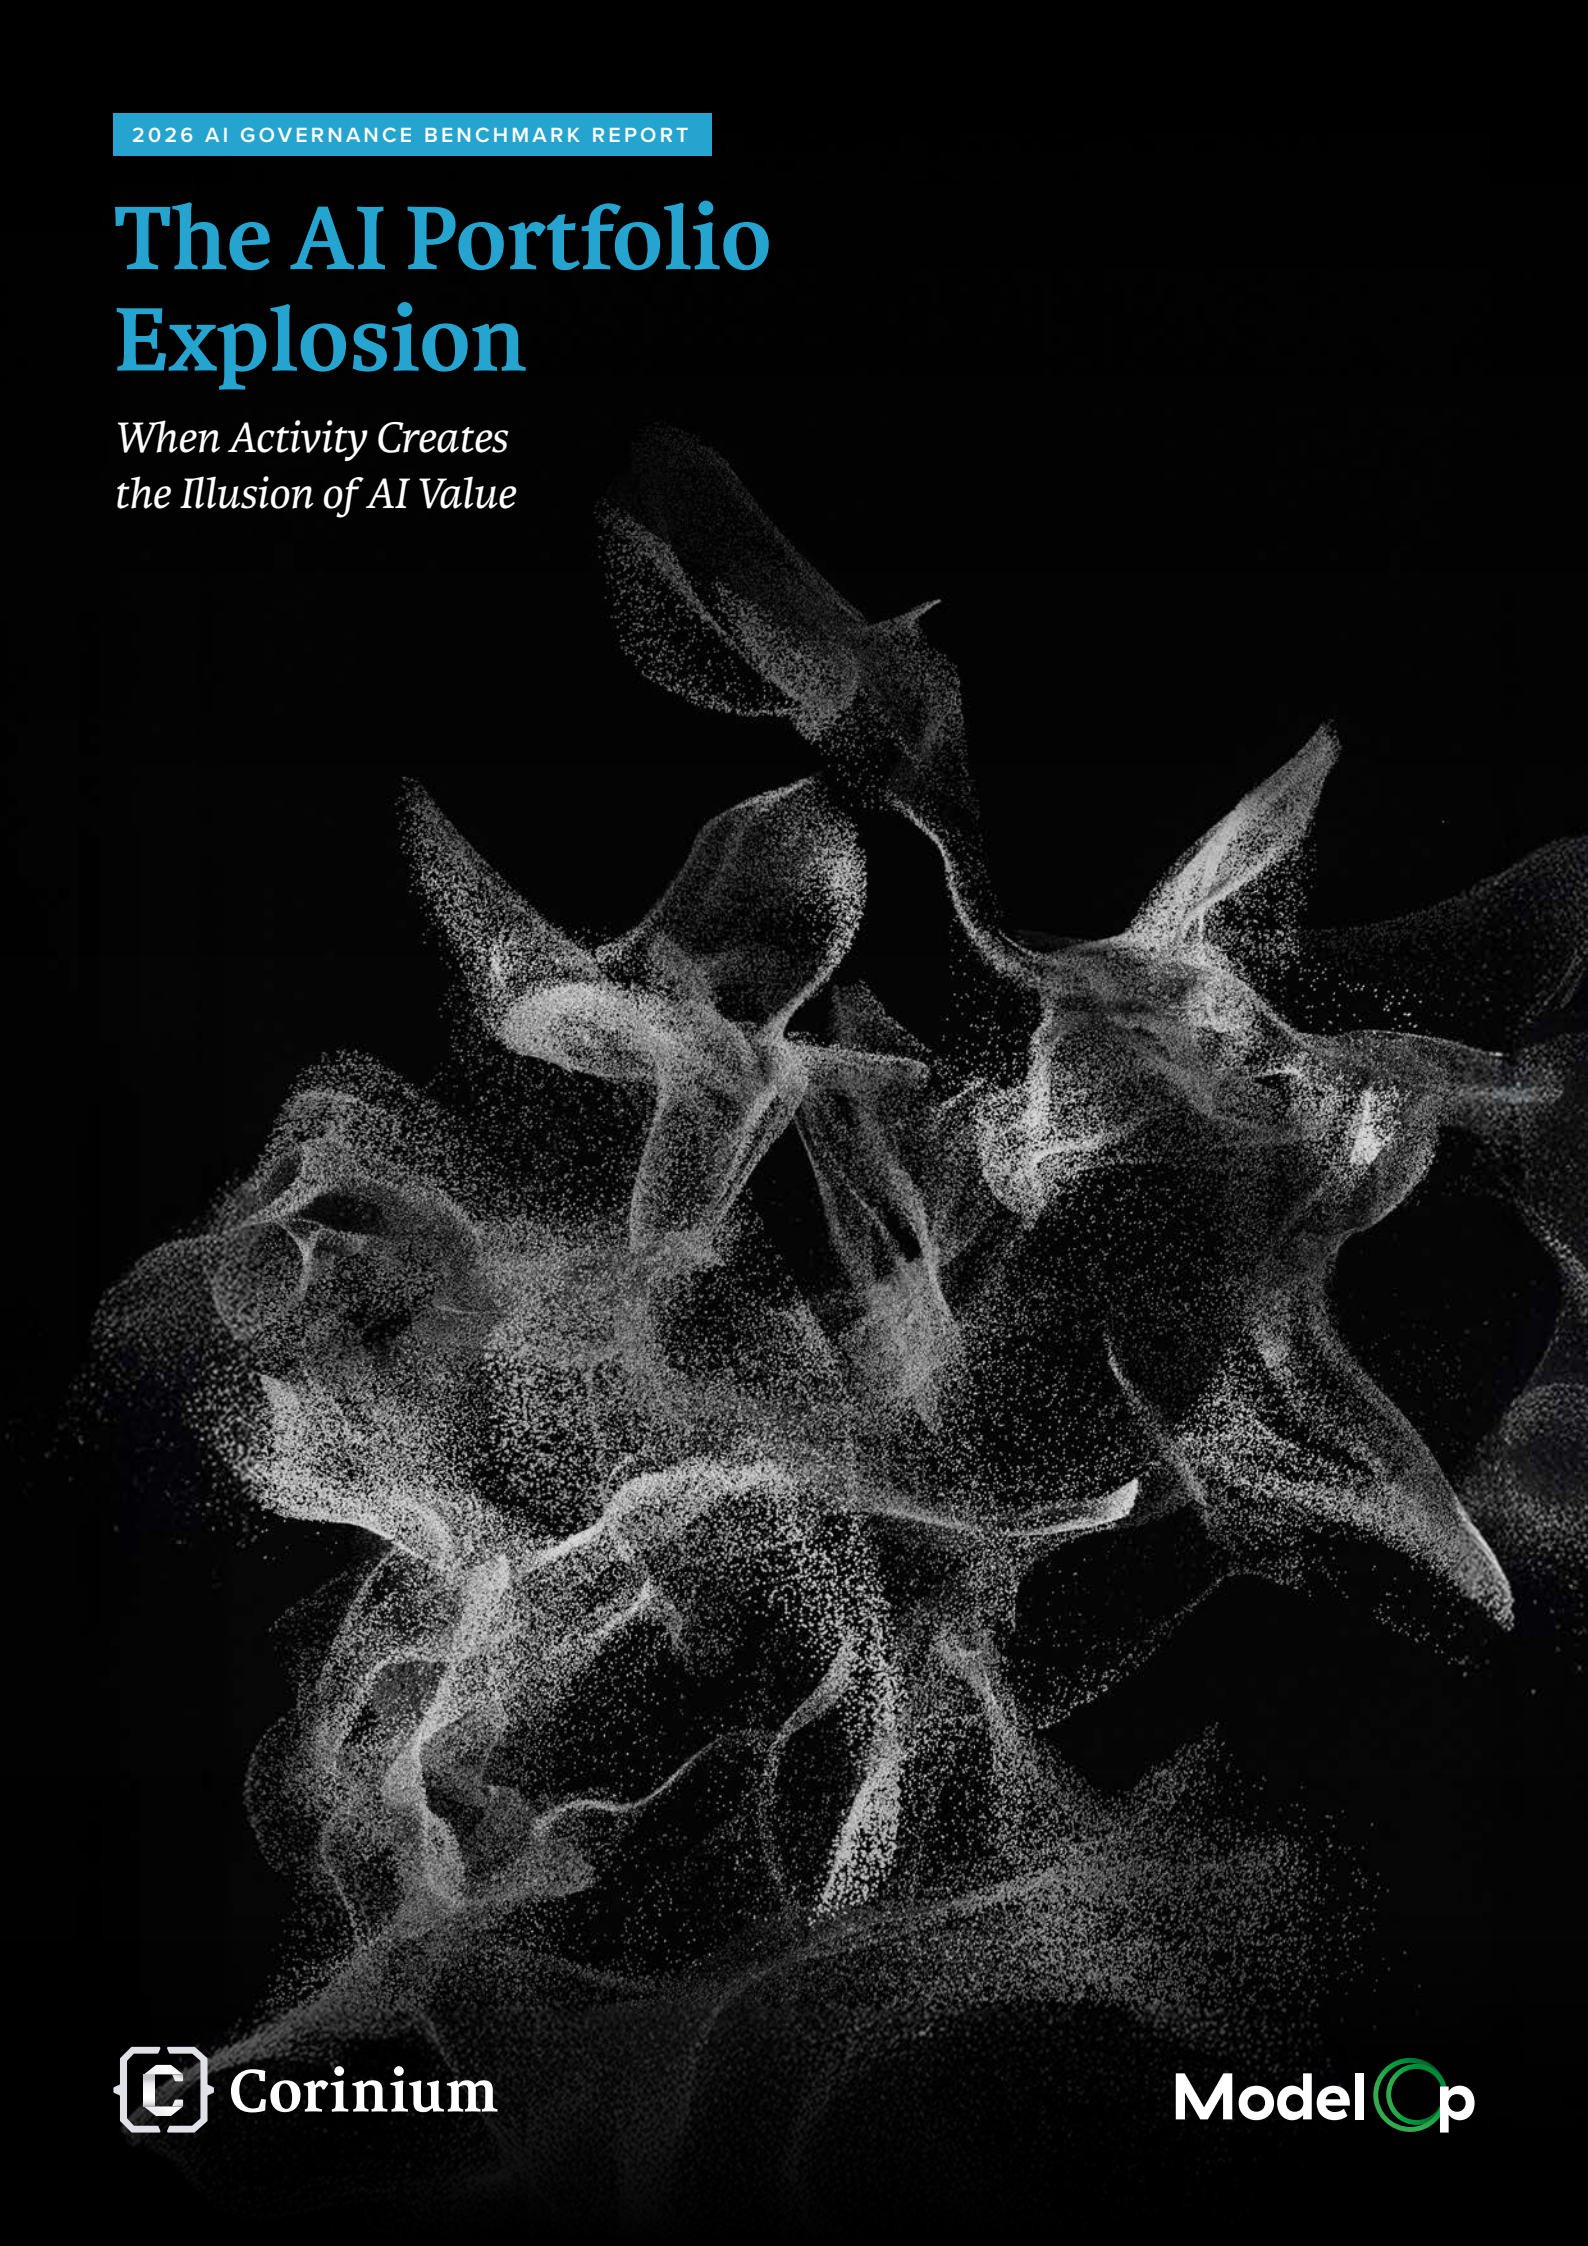

Corinium

ModelOp

Contents

Click below to navigate

2 *Key Findings*

3 INTRODUCTION
A new phase for enterprise AI

5 CHAPTER ONE
The AI portfolio explosion

11 CHAPTER TWO
Why is there an illusion of AI value?

16 CHAPTER THREE
Industrializing AI delivery

22 CONCLUSION
Scaling AI and realizing value means operating AI as a managed portfolio

23 DATA INSIGHTS
What's Blocking AI at Scale — And Why Agentic AI Changes Everything

Contributors



Jeannie Furlan
Vice President, Financial Data and Analytics Strategy, Mutual of Omaha



Farid Sheikhi
Senior Manager, Analytics Innovation and Data Enablement, Royal Bank of Canada



Ramila Peiris
Global Head Data, ML and AI Platform, MSAT, Sanofi



Barbara Widholm
Vice President, Emerging Technologies, State Street

Key Findings

1 AI portfolios are exploding

The most common enterprise AI portfolio now contains 101–250 proposed use cases—up sharply year over year.

2 Success at scale remains rare

Despite hundreds of proposed use cases, most enterprises operate fewer than 25 AI systems at production scale.

3 Time-to-market is accelerating for GenAI use cases that reach production

The majority of GenAI and agentic use cases can reach production in under 6 months—much faster than at the beginning of 2025 when more than half took 6 to 18 months to reach production.

4 ROI tracking is still manual

More than two-thirds of organizations rely on manual or projected ROI metrics, even for production AI systems.

5 Agentic AI multiplies third-party risk

Most enterprises connect agentic systems to 6 to 20 external tools or services, expanding dependency and cost exposure.

6 Lifecycle management adoption has surged

Use of commercial AI lifecycle management and governance platforms jumped from 14% to nearly 50% year over year.

INTRODUCTION

A new phase for enterprise AI

After heavy investment, hundreds of AI use cases are proposed, developed, and piloted each year in large enterprises. But progressing from AI pilots to production remains a huge challenge, and as a result, leadership attention is shifting to a consequential question: what measurable value is AI delivering once it is live, at production scale?

Across industries, executives face growing pressure to explain AI outcomes in financial and operational terms. Boards want transparency. Regulators expect accountability. AI leaders, CIOs, and CFOs demand clarity on costs, risk, and returns. In this environment, activity alone no longer signals progress. Leaders must be able to deliver AI use cases through a repeatable, governable, and economically accountable operation across the enterprise portfolio. ▶



This tension is intensified by the rise of generative and agentic systems. A single use case can now reach production in months, often accelerated by third-party platforms and prebuilt capabilities. While this speed lowers barriers to innovation, it also expands dependency, cost exposure, and third-party risk. Systems that operate continuously, interact with multiple tools, and evolve over time cannot be understood—or controlled—through periodic reviews or manual reporting.

The result is a growing value illusion. Activity creates confidence: more models deployed, more agents connected, more integrations added. But many use cases still fail to move past the pilot stage, and confidence collapses without operational intelligence. Visibility fades as costs accumulate without clear attribution and ownership fragments across teams and vendors. What initially appears successful during a pilot often becomes difficult to explain, measure, or defend months later.

Based on a global survey of 100 senior AI, data, and technology leaders, this report examines why AI value fails to scale even as use cases proliferate and deployment accelerates. It explores where friction can limit scale, why ROI remains difficult to track, and how enterprises become overwhelmed by expanding portfolios rather than empowered by them.

The problem is not a lack of ambition or innovation. Enterprises are not failing at building AI. They are failing at industrializing AI delivery and running AI as a managed portfolio. Leading organizations are responding by rethinking AI delivery itself—treating it as an operating discipline, not a deployment milestone. By embedding governance, monitoring, and accountability directly into delivery workflows, they move beyond the illusion of value and toward industrialized AI portfolio management.

This report explores that transition—and why those that fail to make it will remain bottlenecked by exploding portfolios, trapped in the illusion of value they believe AI is creating. ■



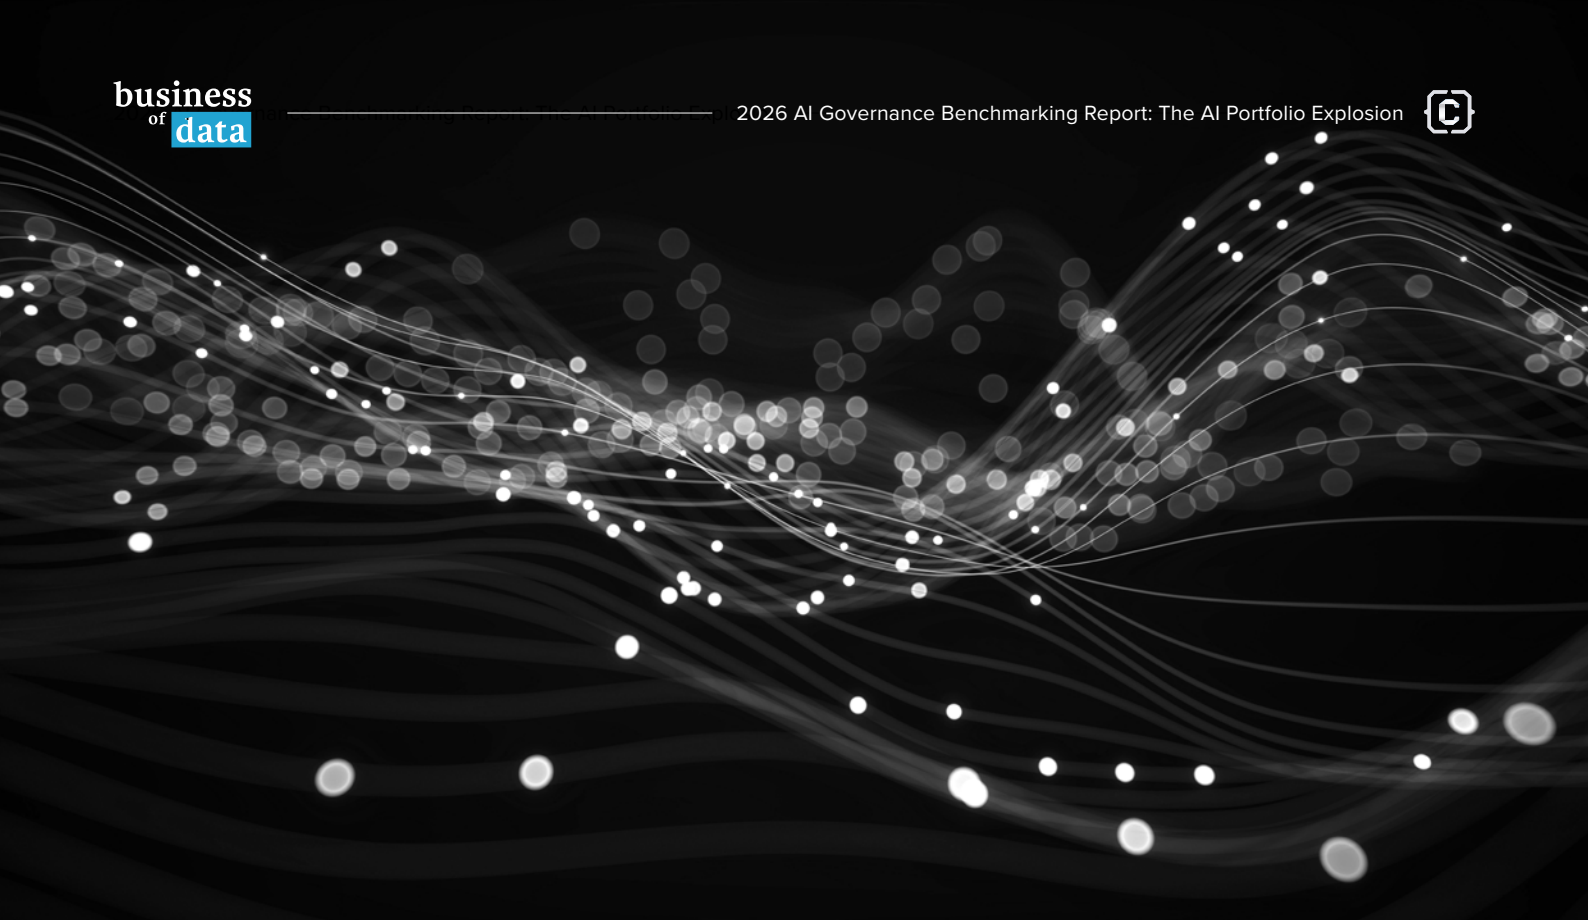

CHAPTER ONE

The AI portfolio explosion

Over the past few years, enterprises have launched a historic wave of AI investment. Across industries, AI initiatives have increased their breadth and depth, now incorporating traditional machine learning, generative and increasingly autonomous agentic systems.

Many organizations are moving models into production faster than they were a year ago, often by leveraging third-party platforms, tools, and prebuilt capabilities.

However, this acceleration may be mistaken for progress. AI portfolios are growing in volume and complexity, while enterprise control is not keeping pace.

As agentic systems connect to more tools, APIs, and external services, gaps in visibility and accountability become harder to contain. This is one factor that makes it difficult to translate AI activity to value.

Our research indicates a sharp expansion of enterprise AI portfolios over the past year. In 2025, most organizations (80%) reported fewer than 100 proposed AI and machine learning use cases. In 2026, portfolio sizes have shifted dramatically: 67% of organizations report between 101 and 250 proposed AI and ML use cases, while a growing share report more than 250. Only 21% now report fewer than 100 proposed use cases.

This explosion of activity is not confined to early-stage ideation. 86% of respondents now report 26 or more use cases in development, 50% report 26 or more in pilot, and 60% report between 6 and 25 AI systems in scaled production.

As a result, many organizations are managing hundreds of AI ideas, pilots, and production deployments at the same time, introducing levels of operational complexity that few organizations were designed to support. ▶

AI portfolios are exploding

How many AI/ML use cases does your organization currently have in each stage of the AI lifecycle?

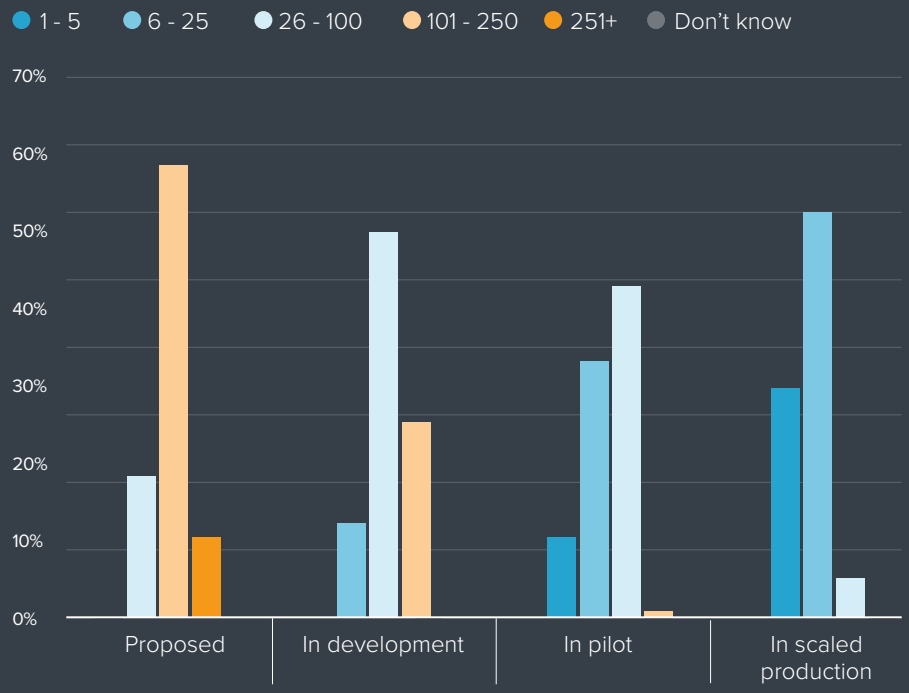

Year-over-year comparison

2025

Most enterprises reported fewer than 100 AI and machine learning use cases across their portfolios.

2026

Enterprise AI activity is spread across every stage of the lifecycle. Survey data shows that 67% of organizations reported 101–250 proposed AI and ML use cases, 86% reported 26 or more use cases in development, 50% reported 26 or more in pilot, and 60% reported 6-25 AI systems in scaled production. Only a small fraction, 6%, have 26 to 100 use cases in scaled production.

Agentic AI use cases

How many agentic AI use cases does your organization currently have in each stage of the AI lifecycle?

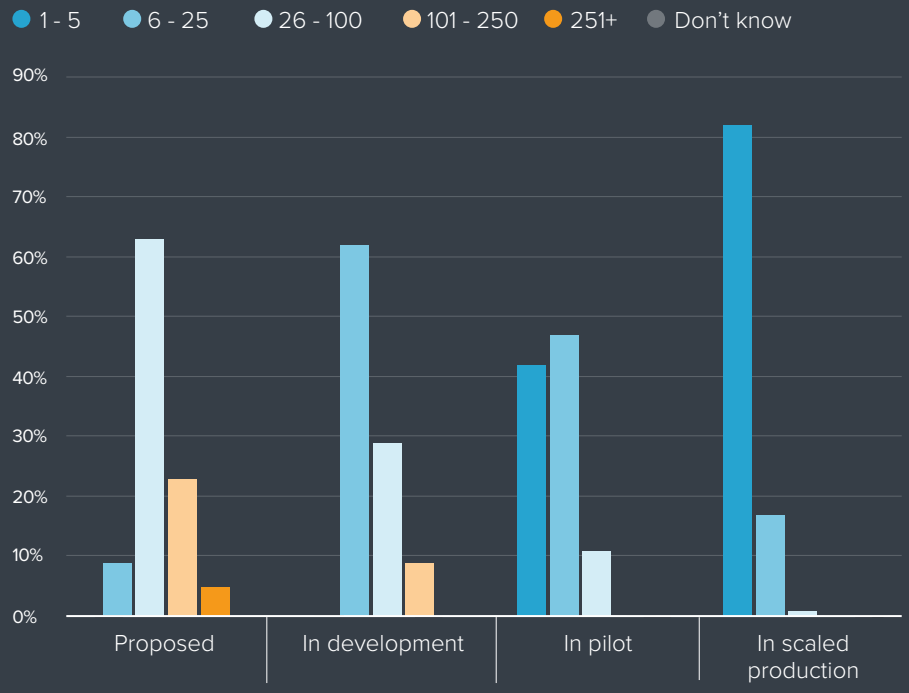

While nearly two-thirds of organizations have proposed 26–100 agentic AI use cases, 82% have scaled fewer than five.

Large AI portfolios, few use cases at production scale

Compared with last year's research, AI initiatives are reaching initial production environments more quickly.

In 2025, more than half of enterprises reported that GenAI initatives are taking between six and eighteen months to reach production. In 2026, respondents said that the majority of GenAI use cases now reach production within six months, with the largest share deploying in one to three months.

Similar acceleration appears across other AI categories. A significant portion of traditional machine learning models reach production within one to three months, while most complete deployment within six months. Agentic AI initiatives also move quickly, with the majority reaching production within three to six months and only a small minority extending beyond a year.

This acceleration is partly enabled by increased reliance on vendor-provided and third-party AI capabilities. These reduce development friction and infrastructure requirements, allowing organizations to move individual use cases into production environments more rapidly.

However, speed does not equal scale. Our research shows that faster deployment does not translate to scaled production adoption.

While enterprises can now move individual AI initiatives into production environments with greater speed, only a limited subset progress beyond initial deployment into sustained operational use. Furthermore, faster deployment may increase risk, cost opacity, and operational friction without an industrial-scale approach to delivering AI.

The constraint facing enterprises is no longer how quickly a single AI use case can be brought to production, but how many systems can be operated, governed, and sustained at once. More importantly, enterprises need to know which use cases are truly delivering value. ▶

Production timelines are accelerating

On average, how long does it take for AI and model initiatives to move from intake to production?

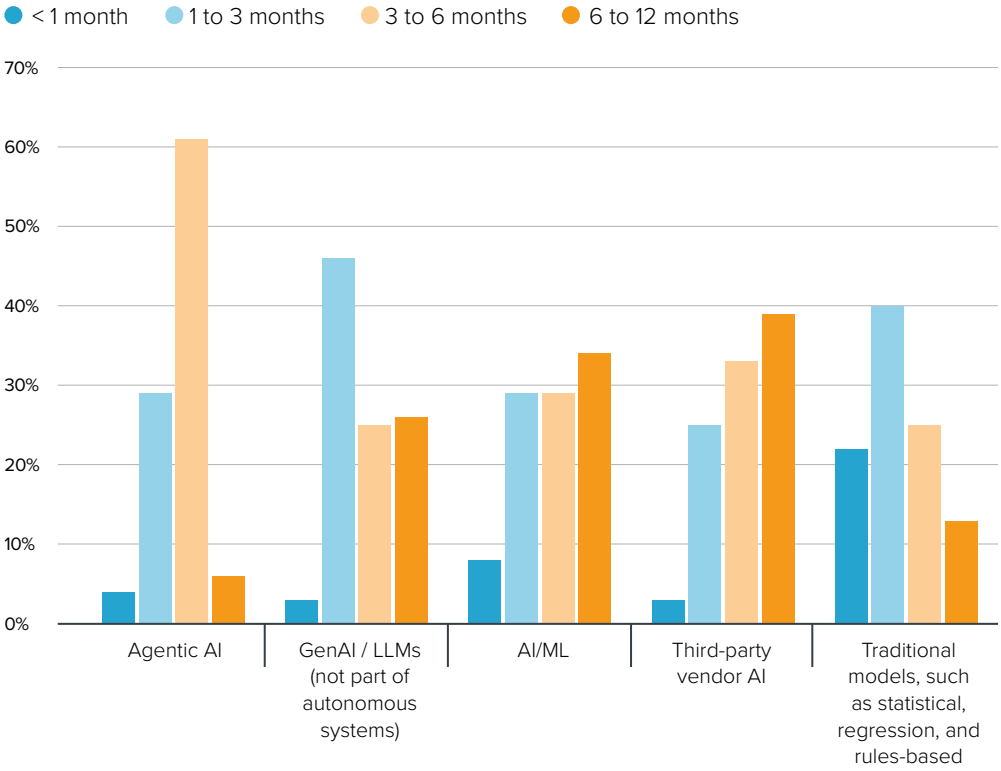

Year-over-year comparison

2025

More than half of enterprises reported GenAI taking 6–18 months to reach production.

2026

GenAI deployment timelines have shortened significantly, with the majority of use cases reaching production within six months, and the largest share reaching production in one to three months.

Source: Corinium Intelligence, 2026

Fragmented AI environments mean complexity at scale

As AI portfolios expand, the environments supporting them are increasingly heterogeneous.

Rather than deploying AI within unified platforms, most enterprises operate across a mix of cloud services, development frameworks, data platforms, and vendor tools. Different teams adopt different stacks, often optimized for speed and experimentation rather than long-term operational consistency and efficiency.

This expansion is accelerating as organizations adopt agentic AI. Our research shows that most enterprises connect agentic systems to between six and twenty third-party tools and services. Each additional connection introduces new dependencies, ownership boundaries, and potential points of failure.

> "Beyond app-level controls, the real governance challenge is coordination across data system owners. If someone changes a source system without communicating it or following standards, the downstream impact can break AI functionality"
>
> **Ramila Peiris**
> Global Head Data, ML and AI Platform, MSAT, Sanofi

As Ramila Peiris, Global Head Data of ML and AI Platform at Sanofi, explains, even minor changes can have outsized consequences once AI systems operate across multiple environments:

"Beyond app-level controls, the real governance challenge is coordination across data system owners. If someone changes a source system without communicating it or following standards, the downstream impact can break AI functionality."

Emerging standards such as Model Context Protocol (MCP) are enabling this expansion by simplifying how agents access external systems. While this increases flexibility and speed, it also significantly broadens the operational surface area that enterprises must manage.

At portfolio scale, these interconnected environments become difficult to manage. Teams may understand individual models or tools but lack a clear view of how systems interact once deployed. As a result, complexity compounds with each new use case instead of increasing linearly. ▶



Agentic AI development is concentrated around major platforms

What development tools are you using for Agentic AI?

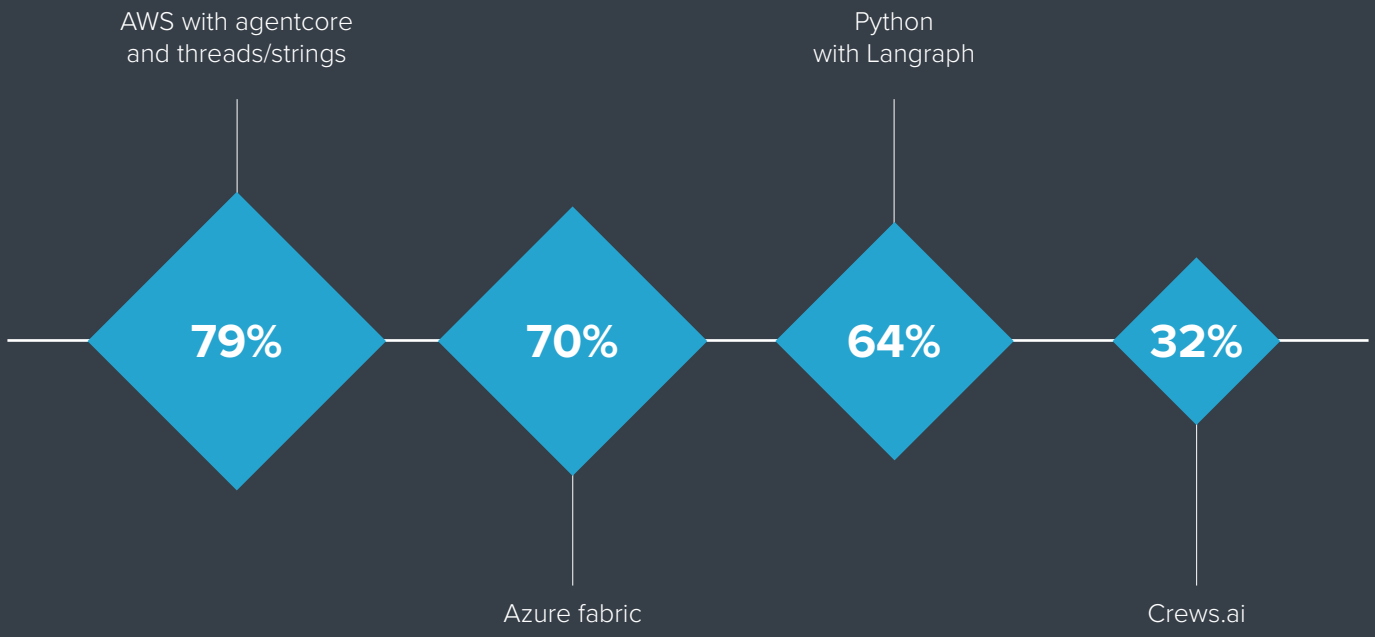

Agentic AI architectures depend on broad vendor ecosystems

How many of your third-party service vendors are you seeking to connect to via MCP as part of an agentic solution?

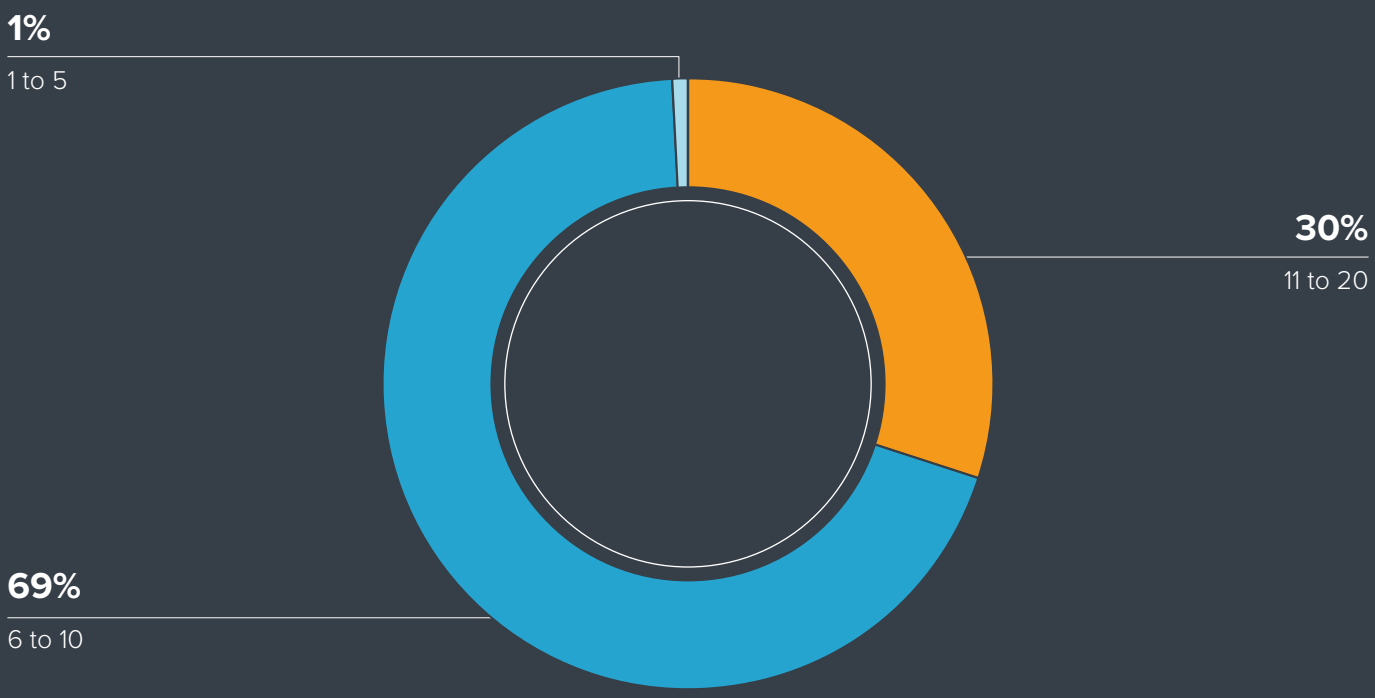

Source: Corinium Intelligence, 2026

The illusion of value

The research points to an enterprise AI landscape that is larger, faster moving, and more complex than ever before. Organizations are generating more AI use cases, moving models into production more quickly, and experimenting across a growing mix of platforms, tools, and vendors.

On the surface, this momentum suggests rapid progress. In practice, however, activity and interdependencies are scaling faster than enterprise control.

As Mutual of Omaha VP of Financial Data and Analytics Strategy Jeannie Furlan explains, the challenge is not a lack of ambition, but the widening gap between technological acceleration and organizational adaptability: "The pace at which these AI capabilities demand us to change is far faster than our ability to change. And that is both the people-side as well as the technology-side."

> "The pace at which these AI capabilities demand us to change is far faster than our ability to change. And that is both the people-side as well as the technology-side."
>
> **Jeannie Furlan**
> Vice President, Financial Data and Analytics Strategy, Mutual of Omaha

The result is an illusion of value. More models reach production. More agents are deployed. More integrations are added. Yet visibility in how these systems behave together across the enterprise does not increase at the same rate.

Understanding why this illusion has emerged, and why speed alone cannot close it, is critical to generating sustained value going forward. ■

Enterprise GenAI and agentic AI development relies on complex tool ecosystems

How many platforms, tools and systems does your enterprise use in total for the development and execution of both third-party vendor and in-house solutions related specifically to GenAI and Agentic AI?

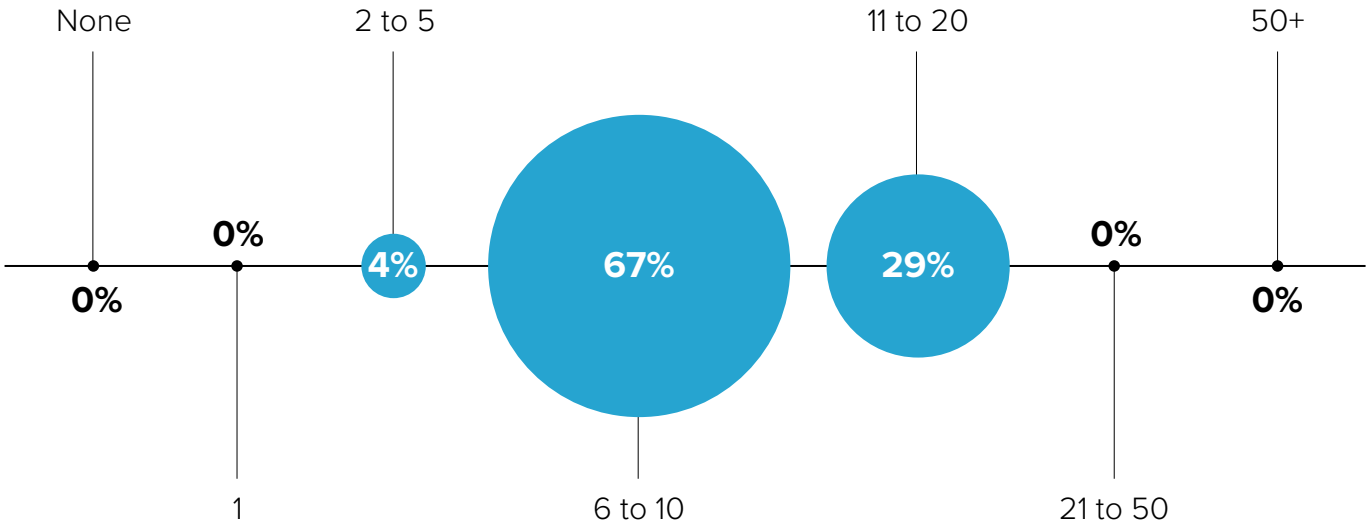

Source: Corinium Intelligence, 2026

CHAPTER TWO

Why is there an illusion of AI value?

Our research among 100 AI, data, and technology leaders shows that enterprises have no shortage of ideas about how to use AI, and as a result the number of use cases is exploding.

Across industries, teams are submitting more AI use cases and are able to launch singular use cases into production faster than ever before, supported by more tools, platforms and vendor ecosystems.

Our research shows that most enterprises continue to rely on manual input or projected assumptions to assess return on investment across traditional machine learning, generative AI, and agentic use cases. Fully automated tracking of realized value remains rare. ▶



As AI systems move into scaled production environments, this limitation becomes increasingly consequential. Without reliable, ongoing measurement of metrics tied to use cases and business outcomes, organizations cannot explain what systems actually delivered, how performance evolved over time, or whether cost and value assumptions remain valid once usage scales.

In this environment, ROI becomes a matter of guesswork rather than evidence, a vulnerability that is increasingly difficult to defend under executive and financial scrutiny. ▶

ROI measurement remains largely manual across AI initiatives

How does your organization currently measure ROI for each of the following types of AI initiatives?

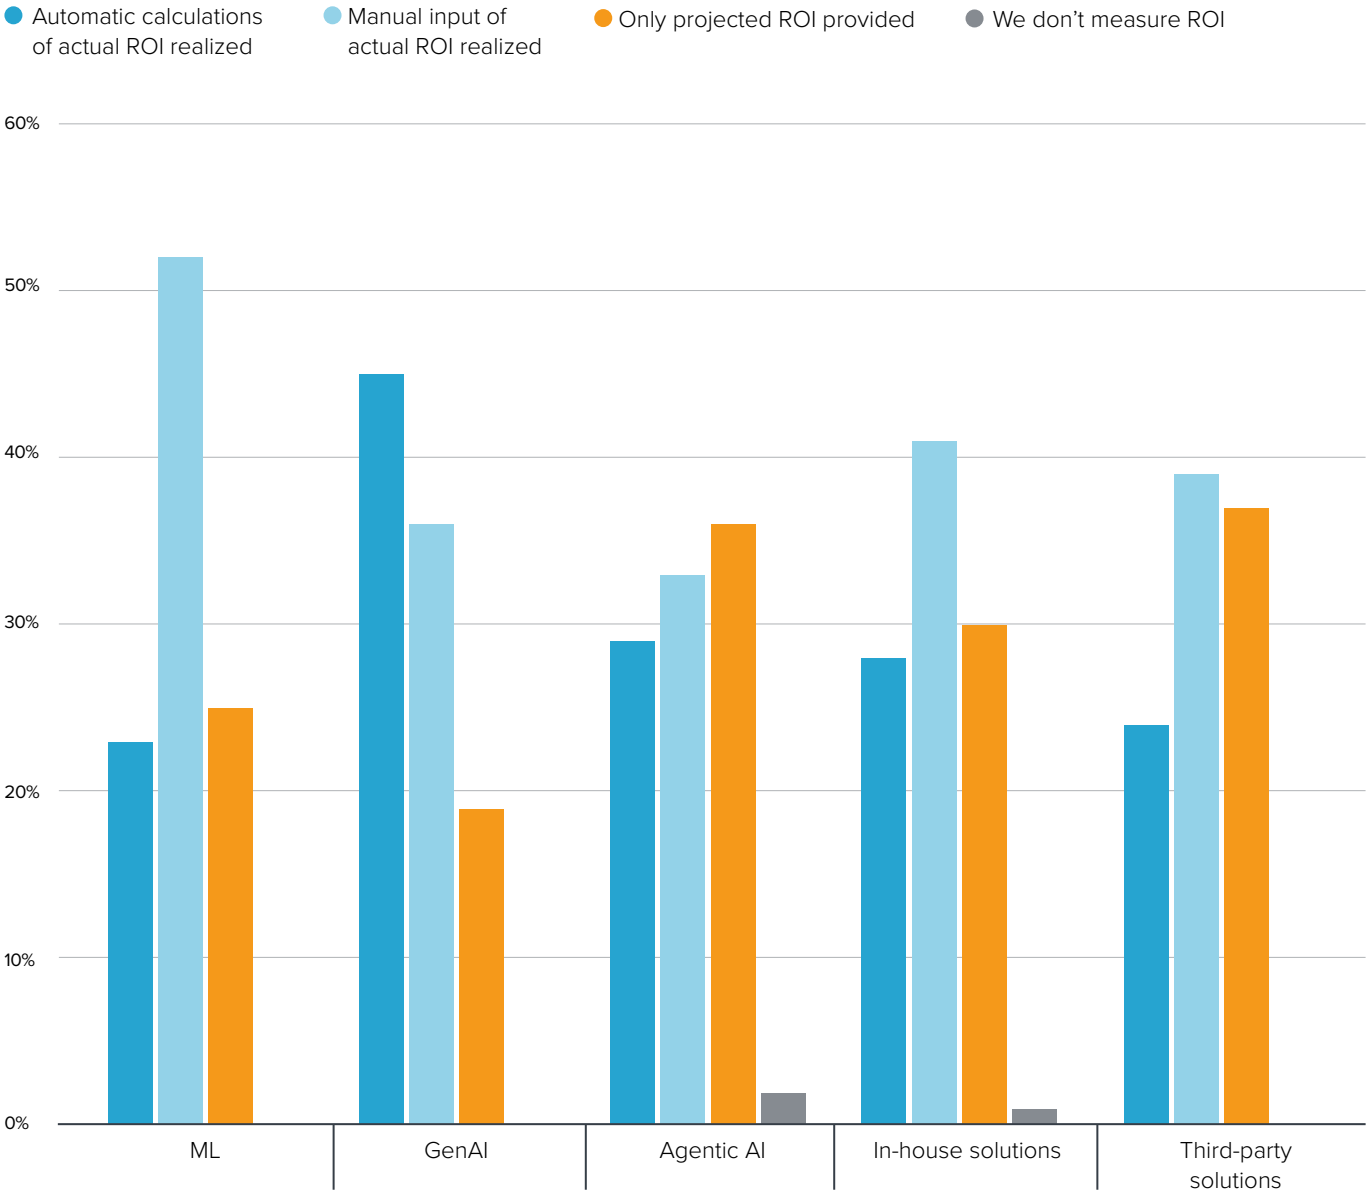

Source: Corinium Intelligence, 2026

Organizations express moderate confidence in tracking agentic AI costs

How confident are you in your organization's ability to accurately track the costs associated with agentic AI use cases?

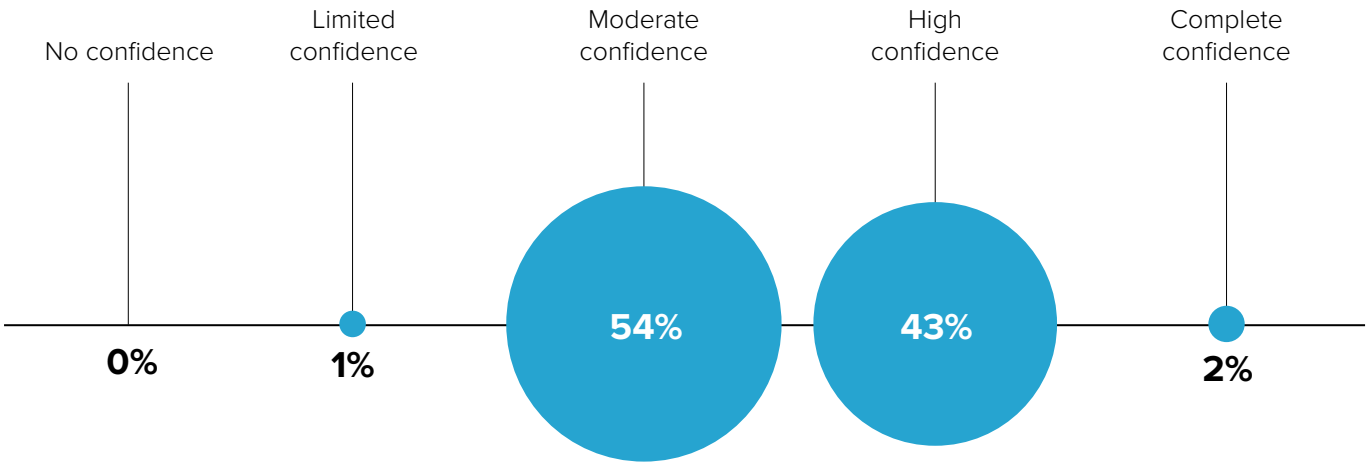

Source: Corinium Intelligence, 2026

Limited cost visibility compounds the value gap

Agentic and GenAI systems introduce ongoing, variable consumption patterns. Unlike traditional applications with predictable infrastructure spend, AI systems may invoke models, tools and external services dynamically based on context and workload.

This complicates any straightforward analysis of the costs and the benefits of AI initiatives.

While many organizations express moderate confidence in their ability to track AI costs, relatively few report complete confidence. Cost attribution is becoming yet more difficult as systems increasingly span multiple platforms, vendors and APIs.

Without clear insight into how costs accumulate across the full AI lifecycle, financial oversight weakens. Leaders may know that AI investment is increasing but struggle to explain why.

The implications extend beyond reporting. When ROI cannot be measured consistently, leaders lose the ability to compare AI initiatives, prioritize investment, or justify continued funding. AI portfolios become difficult to manage not because models are opaque, but because outcomes are. ▶

> "Business teams bring the use cases, IT brings the technical capability, and governance bodies ensure oversight across the entire ecosystem."

Farid Sheikhi
Senior Manager, Analytics Innovation and Data Enablement, Royal Bank of Canada

Fragmented ownership prevents accountability

Once AI systems move beyond pilot environments, ownership becomes increasingly fragmented. Models are often developed by data science teams, deployed through centralized platforms, consumed by business functions and supported by multiple technology and vendor partners.

This fragmentation is reflected directly in the survey data. While many organizations report moderate or high capability across individual AI governance policies, far fewer demonstrate full capability across the lifecycle.

"True maturity comes when data, AI, and business functions share accountability and move in sync, " Sheikhi concludes. "Responsible AI depends on synchronized ownership across all three domains."

For most governance activities, including third-party accountability, vendor due diligence, incident response and continuous monitoring, fewer than one in four organizations report mature abilities that allow comprehensive execution.

"Governance requires partnership across IT, cybersecurity, and the business. Business teams bring the use cases, IT brings the technical capability, and governance bodies ensure oversight across the entire ecosystem," Sheikhi says.

What practitioners say blocks AI from reaching scaled production

When asked to identify the single biggest barrier preventing AI use cases from progressing beyond pilot into scaled production, respondents did not point primarily to model performance or algorithmic limitations. Instead, their answers clustered around organizational, operational, and governance challenges.

Across AI/ML use cases, the most frequently cited barriers are related to fragmented and unreliable data foundations, integration challenges with legacy systems, and the absence of consistent governance and operational processes. Many respondents highlighted difficulties moving from experimentation into repeatable delivery, noting that models often function in isolation but fail when embedded into real business workflows. ▶

Full AI governance capability remains limited across organizations

What level of ability does your organization currently have to perform each of the following AI governance or policy actions?

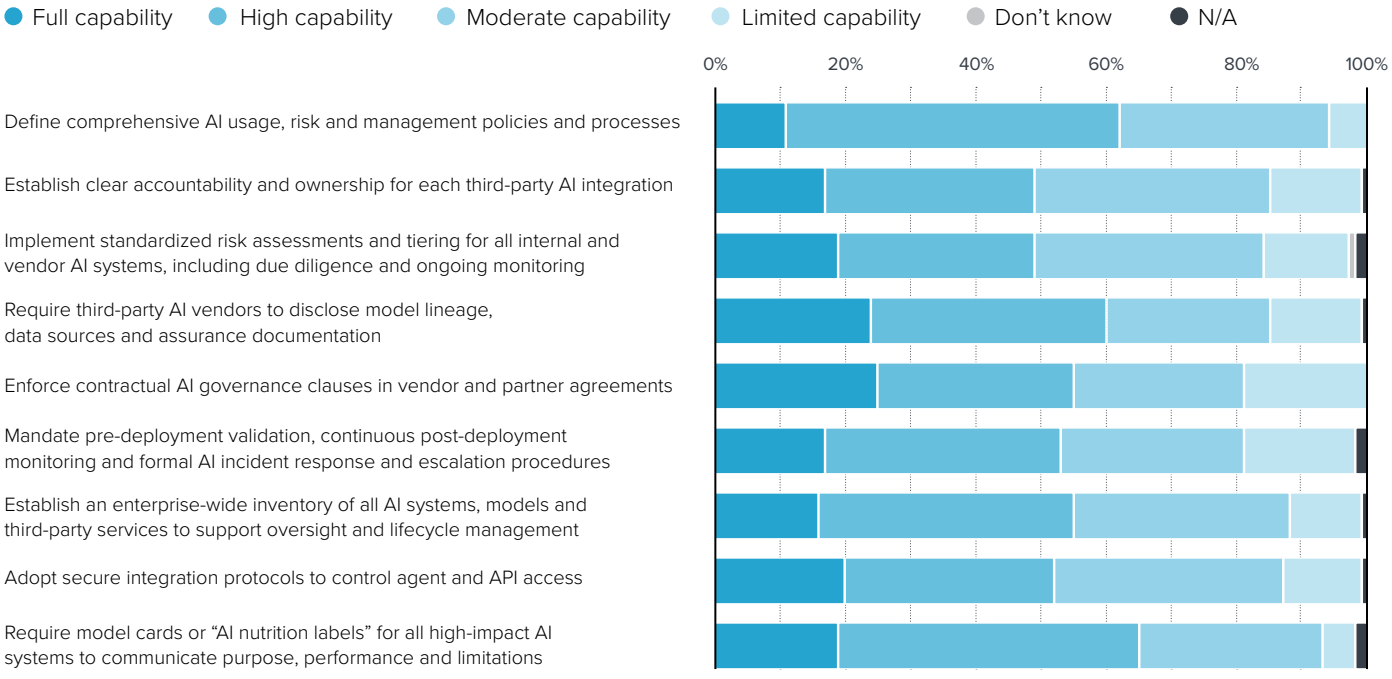

Source: Corinium Intelligence, 2026



For agentic AI use cases, these challenges intensified. Responses were heavily concentrated around governance, accountability, safety, and trust. Leaders cited unclear ownership for autonomous decision-making, limited explainability, and uncertainty around risk exposure as primary constraints on scaling.

It's clear that many AI initiatives struggle to reach sustained production not because the technology fails, but because enterprises lack the operational structures required to support industrial-scale AI delivery.

The AI value illusion is structural, not technical

AI initiatives are not failing due to a lack of technical capability, model performance, or investment. In fact, organizations are building and deploying AI faster than ever before.

These AI experiments create the perception of value, but the momentum can mask deeper structural weaknesses that limit the ability of these initiatives to create real business value.

Many pilots are not treated as enterprise initiatives. Ownership is often unclear, accountability fragmented, and business and IT priorities misaligned. As a result, pilots can demonstrate technical promise without establishing the organizational conditions required for scaled production.

> "True maturity comes when data, AI, and business functions share accountability and move in sync. Responsible AI depends on synchronized ownership across all three domains"
>
> **Farid Sheikhi**
> Senior Manager,
> Analytics Innovation and Data Enablement,
> Royal Bank of Canada

At the same time, business value remains difficult to articulate early. Organizations struggle to define measurable outcomes, link AI initiatives to clear KPIs, or quantify ROI in a way that will justify investment in the long term.

This disconnect defines the AI value illusion. It is not a failure of innovation speed, but of operational readiness. ■

CHAPTER THREE

Industrializing AI delivery

As AI portfolios mature, leading enterprises are shifting their focus from deploying individual AI pilots across fragmented teams to delivering and managing AI at an industrial scale as an operational discipline across the business.

The objective is not about additional oversight alone (although that remains important), but to enable portfolio-level decision making, investment prioritization, and executive accountability across the entire AI lifecycle. This shift reflects a growing recognition that AI value cannot be consistently measured and maintained through isolated tools, point-in-time reviews or manual coordination.

As Farid Sheikhi, Senior Manager for Analytics Innovation and Data Enablement at RBC, explains, traditional governance approaches were not designed for this operating reality:

"Organizations are moving away from static, checkbox governance toward lifecycle governance. With agentic systems, compliance at launch is no longer enough. You need continuous understanding, monitoring and direction of systems that adapt, interact with other systems and change their behavior over time." ▶

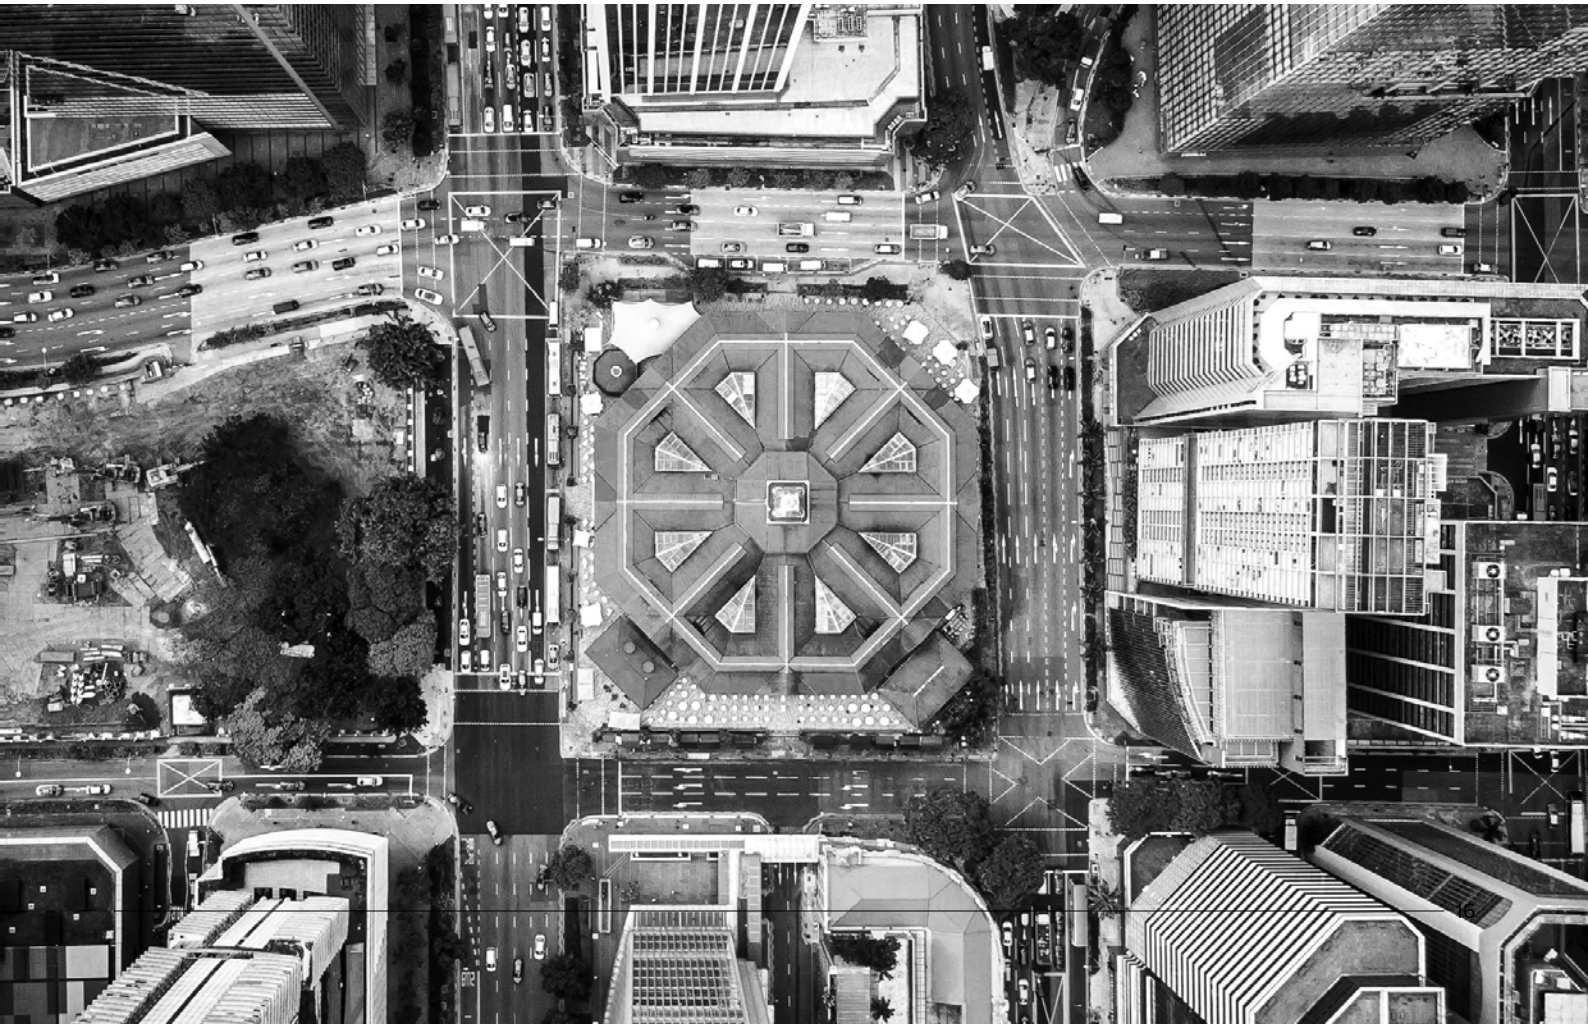

Commercial AI governance platforms have moved into the enterprise mainstream

Are you currently using or pursuing an AI lifecycle management or governance solution?

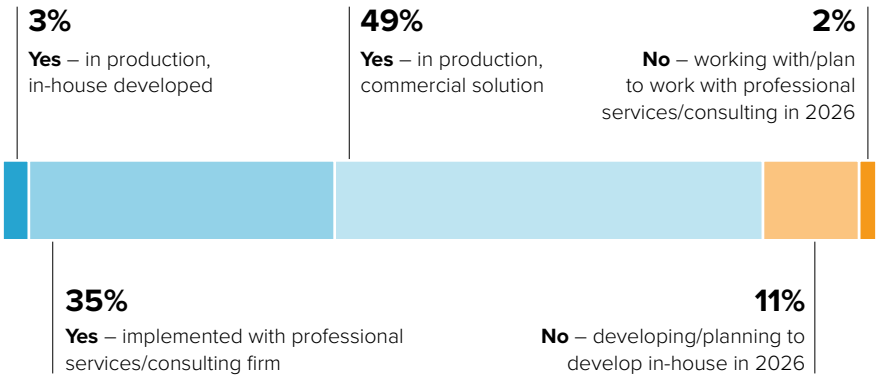

Year-over-year comparison

2025

Only 14% of enterprises had commercial AI governance software in production.

2026

Just under half of enterprises (49%) report using commercial AI lifecycle management or governance platforms.

Most organizations now allocate substantial budgets to AI governance

How much has your organization budgeted for AI lifecycle management or governance software in the upcoming year?

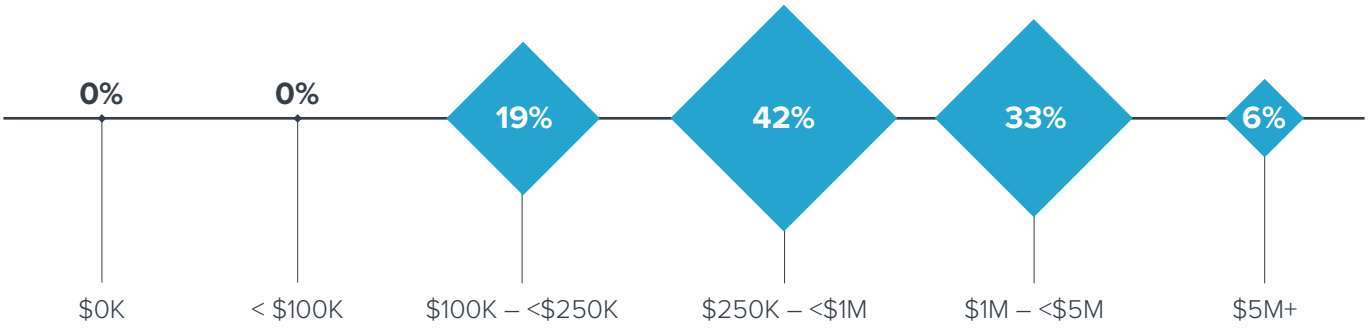

Source: Corinium Intelligence, 2026

Survey results indicate that this transition is already underway. In 2025, only 14% of enterprises reported having commercial AI governance software in production, with the majority, 74%, relying on professional services. In 2026, 49% of organizations report using commercial AI lifecycle management or governance platforms and only 35% reported relying on professional services, marking a significant shift in how governance and oversight are operationalized.

This year-over-year change reflects a broader move toward standardized, repeatable operating infrastructure designed to support AI at production scale.

Critically, lifecycle management enables governance to operate at the speed of AI itself. Controls are embedded into workflows, data is captured automatically and monitoring persists beyond launch. This reduces dependence on manual reporting cycles and creates the foundation required for consistent financial oversight, risk management and performance measurement.

For organizations seeking to move beyond experimentation, lifecycle management provides the mechanisms needed to transition from just deploying AI quickly to truly delivering AI at an industrial scale. ▶

Foundational capabilities take priority

As organizations invest in lifecycle management and governance platforms, the survey data shows a consistent pattern in how governance capabilities are being prioritized and funded.

Leading enterprises are focusing first on foundational capabilities that make portfolio-level oversight possible. Audit and traceability, risk management, real-time monitoring, and lifecycle workflows and orchestration rank as the highest governance priorities across respondents.

This prioritization is reinforced by continued investment in AI lifecycle management or governance software. Most organizations report annual spending on AI lifecycle management and governance software in hundreds of thousands of dollars, with more than 40% budgeting between $250,000 and $1 million annually, and nearly four in ten now allocating more than $1 million per year.

> "Organizations that are ready to scale have clear guardrails and escalation paths for autonomous actions."
>
> **Barbara Widholm**
> Vice President, Emerging Technologies, State Street



These investments signal a shift in mindset. AI lifecycle management and governance are implemented as operational infrastructure that can enable innovation and streamline compliance.

Many organizations are still working to establish a dependable view of their AI portfolios across teams, platforms and vendors. Without this baseline understanding, advanced controls risk creating the illusion of discipline without delivering meaningful oversight.

As Barbara Widholm, Vice President of Emerging Technologies at State Street, explains, effective governance depends on strong underlying structure rather than isolated controls:

"Organizations that are ready to scale have clear guardrails and escalation paths for autonomous actions. They integrate AI governance into existing enterprise processes rather than building separate structures." ▶

Audit, risk, and monitoring capabilities dominate AI governance priorities

Which capabilities of AI lifecycle management and governance are the highest priorities for your organization in 2026?

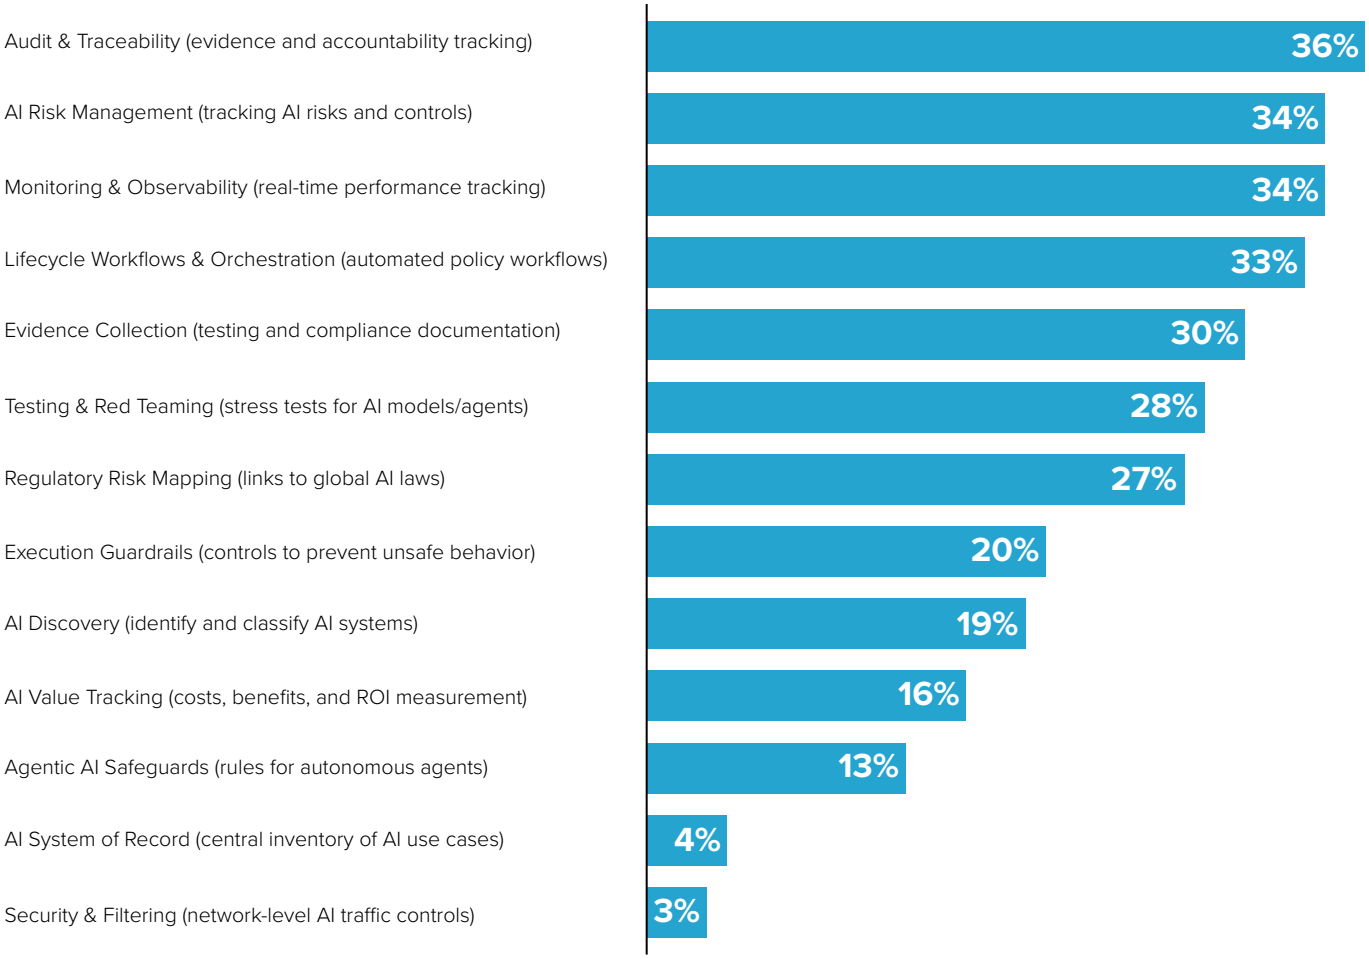

Source: Corinium Intelligence, 2026

The emphasis on foundational capabilities also reflects the operating realities we highlighted earlier. Fragmented tooling, decentralized development and reliance on third-party platforms make it difficult to apply governance consistently through manual coordination alone.

By establishing shared visibility and ownership across the portfolio and using workflow automation to embed governance by design, organizations create the conditions needed to introduce more sophisticated controls over time.

Across the survey results, one pattern consistently distinguishes organizations progressing toward profitable AI transformation: visibility.

Organizations investing in lifecycle management can see, in real time, what AI use cases and associated systems and assets exist across their enterprise, how they are being used, what they cost, and how value is realized.

Without this transparency, enterprises accumulate models faster than they can align them to business priorities. With it, they gain the ability to prioritize use cases, validate ROI, intervene when performance drifts and make informed investment decisions. ▶

Adoption challenges are driven more by organization than technology

What are the biggest challenges your organization faces in adopting new AI lifecycle management and governance software or platforms?

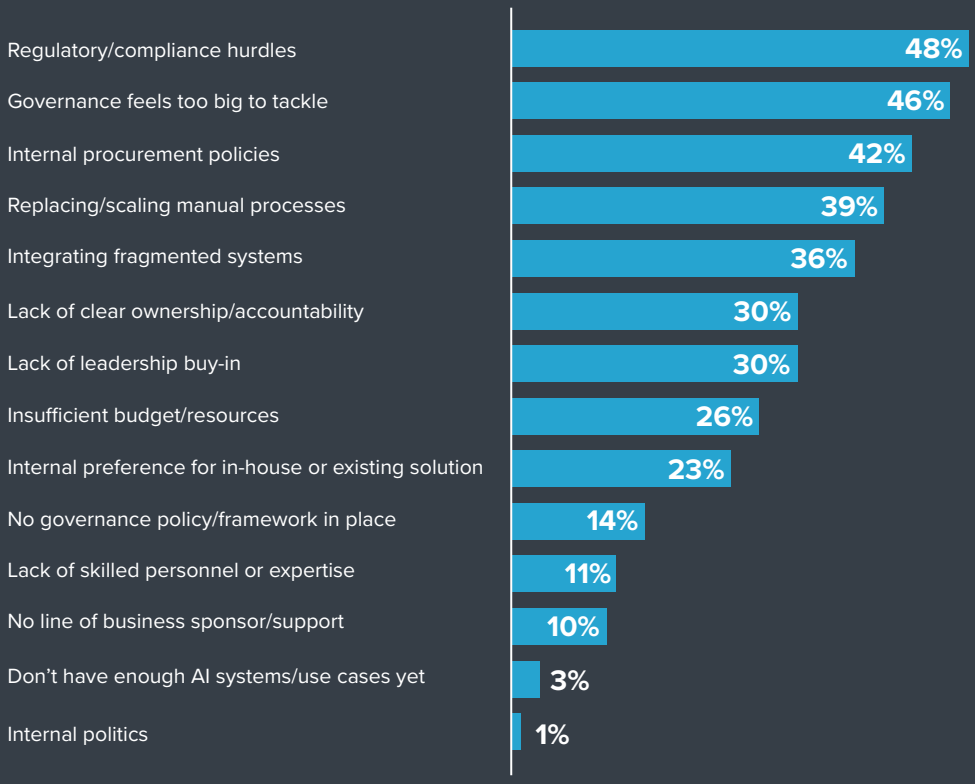

Year-over-year comparison

2025

58% of enterprises cited fragmented systems as a major barrier to AI governance and deployment.

2026

Fragmentation remains a significant challenge, though it has improved compared with 2025 and is now cited by just over one-third of organizations.

Source: Corinium Intelligence, 2026

AI lifecycle management is viewed as a strategic priority across organizations

How important is AI lifecycle management - the end-to-end management of an AI use case from intake through retirement while enforcing policies and applying governance - to your organizations overall AI strategy?

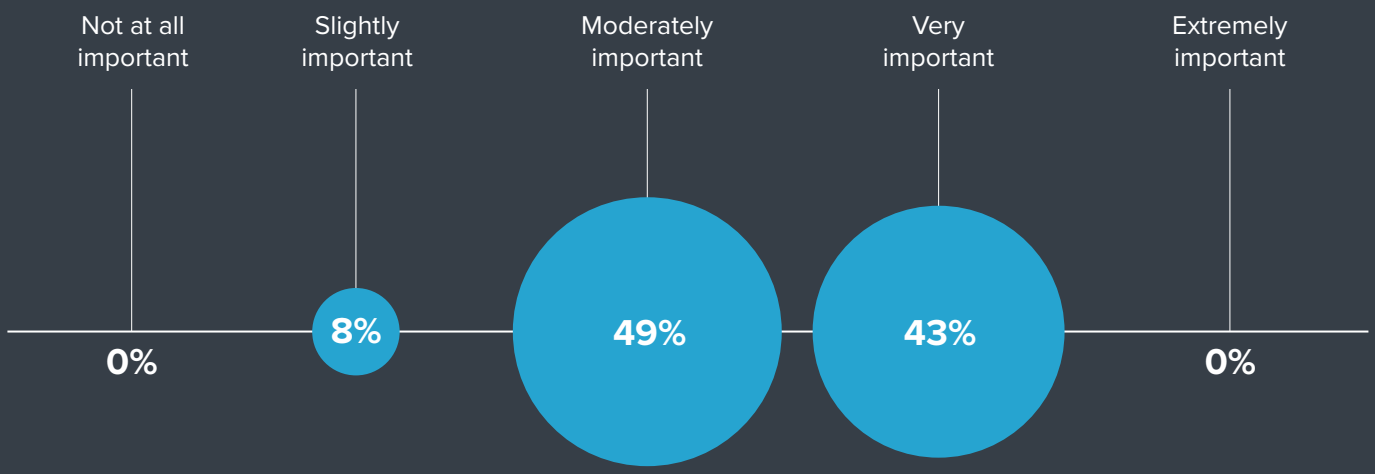

Source: Corinium Intelligence, 2026

Making AI governance work in practice

Enterprises are spending more money on AI governance than ever before, but there are perceived barriers to adopting new AI lifecycle management and governance platforms.

Nearly half of respondents (48%) named regulatory and compliance as a hurdle to adopting new tools, the largest reported hurdle, followed closely by the perceived scale and complexity of governance programs (46%), the difficulty of aligning new tooling with existing procurement policies (42%), and replacing or scaling manual processes (39%).

Integration challenges are also high on the list (36%), highlighting the reality that AI portfolios often span fragmented platforms, vendors, and development teams.

At the same time, the strategic importance of lifecycle management makes its successful implementation critical. Ninety percent of respondents rate AI lifecycle management as either moderately or very important to their organization's overall AI strategy.

While adoption challenges remain, organizations that treat lifecycle management as foundational are better positioned to cut through the illusion of AI value. ■

> "With agentic systems, compliance at launch is no longer enough. You need continuous understanding, monitoring and direction of systems that adapt, interact with other systems and change their behavior over time."

Farid Sheikhi
Senior Manager,
Analytics Innovation and Data Enablement,
Royal Bank of Canada





CONCLUSION

Scaling AI and realizing value means operating AI as a managed portfolio

The sheer volume of AI investment makes the illusion of value caused by an explosion in proposed use cases and pilots impossible to ignore. Organizations are deploying more AI solutions and models, adopting generative and agentic systems and moving faster from idea to production, yet many struggle to scale AI initiatives from pilots to production successfully and demonstrate sustained business impact.

This year's research shows that the challenge is no longer whether enterprises can build or deploy AI. It is whether they can make the jump from pilots to production and deliver it at an industrial scale with accountability and financial discipline. Leading enterprises are prioritizing end-to-end lifecycle management. Governance is becoming embedded into workflows, supported by standardized platforms and continuous monitoring rather than manual review and point-in-time controls.

As AI systems grow more autonomous and interconnected, this operating foundation becomes critical. Organizations that invest in visibility and automated lifecycle management are better positioned to connect AI initiatives to measurable outcomes and maintain confidence as portfolios expand.

Last year's research showed that speed to market was becoming a defining differentiator in enterprise AI adoption.

This year's findings suggest that speed alone is no longer sufficient. Enterprises are no longer satisfied with AI experiments, they want to see which AI projects are delivering value, where delivery is facing bottlenecks, and how to turn AI ideas into business value at scale.

The next phase of enterprise AI will not be defined by who deploys fastest, but by who can deliver AI at industrial scale and explain how their initiatives align to business use cases and real value. ■

Corinium

DATA INSIGHTS

What's Blocking AI at Scale — And Why Agentic AI Changes Everything

Insights from open-ended responses in our interviews with AI and agentic AI leaders



Top Barriers to Scaling AI and Machine Learning

1 Data Infrastructure and Data Quality

- Fragmented, siloed and inconsistent data
- Lack of unified platforms and reliable pipelines
- Difficulty managing lineage, provenance and versioning
- Limited ability to support continuous retraining and inference at scale

2 Legacy Systems and Integration Challenges

- Difficulty integrating models into legacy enterprise systems
- Rigid architectures not designed for ML workloads
- Fragmented MLOps tooling and limited orchestration
- Integration complexity across multiple platforms and vendors

3 Governance, Compliance and Risk Constraints

- Governance processes perceived as slow or bureaucratic
- Inconsistent lifecycle standards and policies
- Gaps in audit readiness, documentation and monitoring
- Regulatory friction and unclear compliance frameworks

4 Talent and Skills Gaps

- Shortage of ML engineers and MLOps practitioners
- Disconnect between prototype development and production engineering
- Heavy reliance on vendors due to limited in-house expertise

5 Organizational Alignment and Culture

- Pilots not treated as enterprise initiatives
- Unclear ownership, roles and accountability
- Business and IT misalignment
- Limited trust in AI outputs and resistance to automation

6 Business Value, ROI and Executive Sponsorship

- Difficulty demonstrating early, measurable value
- Challenges quantifying ROI to justify production investment
- Lack of clear KPIs or business context
- Insufficient funding or strategic prioritization

Top Barriers to Scaling Agentic AI

1 Governance, Accountability and Policy Gaps

- No established frameworks for autonomous decision-making
- Undefined accountability when agents act incorrectly
- Immature regulatory standards
- Lack of standard approaches for escalation, rollback and auditability

2 Trust, Transparency and Explainability Gaps

- Limited ability to trace agent reasoning or decision paths
- Unpredictable behaviors in complex environments
- Low interpretability in multi-agent interactions
- Stakeholders unable to understand or approve agent behavior

3 Cultural and Organizational Resistance

- Hesitation to delegate authority to machines
- Disruption to established approval hierarchies
- Executive discomfort with autonomous actions
- Limited readiness to adapt workflows and operating models

4 Security, Safety and Risk Exposure

- Autonomous API execution and action triggering
- Unknown propagation of errors across agents
- Inadequate cybersecurity frameworks for autonomy
- Absence of standardized safety controls and kill switches

5 Technology and Tooling Limitations

- Limited maturity of agent orchestration platforms
- Poor interoperability across agent frameworks
- Insufficient sandboxing and simulation environments
- Weak monitoring and observability for multi-agent behavior

6 Integration Challenges with Enterprise Systems

- Difficulty embedding agents into deterministic workflow systems
- Approval and governance models not designed for autonomy

AI vs Agentic AI — Key Differences

CATEGORY	TRADITIONAL AI / ML	AGENTIC AI
Primary Concern	Data infrastructure, pipelines and MLOps	Governance, accountability & safety
Top Technical Gap	Integration and legacy systems	Orchestration and observability
Top Organizational Gap	Business–IT alignment	Cultural readiness for autonomy
Primary Risk Focus	Compliance and audit friction	Liability and autonomous action
Trust Challenge	Trust in model outputs	Trust in autonomous reasoning
Maturity Level	Tooling exists but fragmented	Tooling largely immature or emerging

Key Insight

Traditional AI struggles to scale because of infrastructure limitations. Agentic AI struggles to scale because enterprises are not yet prepared for autonomy.

As AI systems become more autonomous, governance execution, accountability structures and organizational trust — not model performance — increasingly determine whether transformation is possible.



Discover Corinium Intelligence

Corinium is the world's largest business community of more than 250,000 data, analytics, customer experience and digital transformation leaders.

We're excited by the incredible pace of innovation and disruption in today's digital landscape. That's why we produce quality content, webinars and events to connect our audience with what's next and help them lead their organizations into this new paradigm.

Find out more: **www.coriniumintelligence.com**

Corinium

About ModelOp

ModelOp is the leading AI lifecycle management and governance platform, purpose-built for enterprises. ModelOp's platform provides a centralized AI system of record, automation from intake to retirement, and enforceable policies—helping enterprises bring ML, GenAI, Agentic AI, and vendor AI solutions into production 10X faster. ModelOp is used by the most complex and regulated institutions in the world—including major banks, insurers, regulatory bodies, healthcare organizations, pharmaceutical firms, and global CPG companies—because it delivers the structure, automation, and oversight necessary to operationalize AI at scale across the entire enterprise.

Gartner, Forrester, and IDC recognize ModelOp for its end-to-end AI lifecycle management and governance platform. In 2025, it was awarded the "Best AI Governance Software Award" from Netty Awards and received Business Intelligence Group's Artificial Intelligence Excellence Award.

To learn how ModelOp can help you accelerate and scale your AI initiatives through industrialized AI delivery, please contact **sales@modelop.com** or visit **www.modelop.com** to request a meeting.

ModelOp

About the Editor

Gareth Becker is an experienced editor and content marketer and produces B2B stories that focus on emergent trends in data and analytics, cloud computing, information security and more.

He works with world-leading brands to shine a light on fresh ideas and innovative products using a range of multimedia content.

To share your story or enquire about appearing in a Corinium report, blog post or digital event, contact him directly at **gareth.becker@coriniumgroup.com**



Gareth Becker
Content Strategist,
Corinium Global Intelligence

Corinium

Partner with *Business of Data* by Corinium

We'll develop industry benchmarking research, special reports, editorial content, online events and virtual summits to establish your brand as an industry thought leader.

FIND OUT MORE HERE







Discover Corinium Intelligence

Corinium is the world's largest business community of more than 250,000 data, analytics, customer experience and digital transformation leaders.

We're excited by the incredible pace of innovation and disruption in today's digital landscape. That's why we produce quality content, webinars and events to connect our audience with what's next and help them lead their organizations into this new paradigm.

Find out more: **www.coriniumintelligence.com**

Connect with Corinium

- **Join us at our events**
- **Visit our blog**
- **Read our white papers**
- **Follow us on LinkedIn**