

Redefining enterprise intelligence with autonomous AI



Preface

“Redefining enterprise intelligence with autonomous AI” is an MIT Technology Review Insights report sponsored by Uniphore. The findings and perspectives presented here draw on research and in-depth interviews with industry experts and executives. Adam Green was the author of the report, Laurel Ruma was the editor, and Nicola Crepaldi was the publisher. The research is editorially independent, and the views expressed are those of MIT Technology Review Insights.

We would like to thank the following individuals for their time and insights:

Prasad Jayaraman, Principal, KPMG

Cristopher Kuehl, Chief Data, Information, and AI Officer, Continent 8 Technologies

Dan Morris, Global Head of Industry Solutions for Marketers, Databricks

Gerry Murray, Research Director of Enterprise Marketing Apps and Agents, IDC

Tushar Shah, Chief Customer Officer, Uniphore



CONTENTS

01 Executive summary 4

02 Why enterprise AI struggles to scale 5

 Process design in an AI-speed world 6

03 Architecture decisions for long-term AI value 7

 The vertical integration vs. composability trade-off 7

 Governance as architectural infrastructure 9

04 Organizational choices for AI scale 11

 Deciding where to build, buy, or partner 11

 Intelligence that flows across functions..... 12

05 Foundations for the agentic AI shift..... 13

 The competitive stakes of agentic AI delay 14

 Conclusion..... 14



01

Executive summary

Enterprise AI is no longer a future ambition. It is in full operational flight. Model capabilities are advancing faster than most organizations can absorb, while the cost of performance continues to fall. Globally, AI investment is set to reach \$2.5 trillion in 2026, up 44% from the previous year.¹²

For many enterprises, this investment has produced fragmentation. Intelligence can accumulate in silos so that sales agents are unaware of open support tickets, for instance, or marketing systems are personalizing content without visibility into what finance already knows about a customer. Each function may perform well in isolation, but the enterprise as a whole learns little and has less information to act upon.

The shift from AI as a tool to AI as an operating model – what we call the “agentic shift” in this report – demands something more fundamental than better models or faster infrastructure. It requires connecting people, processes, and data in real time, along with the governance and control to act on that intelligence reliably.

This means rethinking both architecture and operating models simultaneously. First, rebuilding data infrastructure for accessibility rather than volume. Second, replacing fixed tech stacks with composable architectures that can evolve as models and tools change. And, lastly, resolving questions of AI sovereignty, including where intelligence runs, who controls it, and how it operates across organizational and jurisdictional boundaries.

Key findings include the following:

- Enterprise AI’s scaling problem is structural.** Process-first companies are pulling ahead. Global AI spending is rising sharply and model capabilities are advancing faster than most organizations can integrate them. Yet the majority of enterprises are still not growing revenue through AI or fundamentally rethinking how they operate. The companies generating sustained returns share a common discipline. They treat process redesign as the work that precedes model selection, building for how the technology will evolve rather than retrofitting roles and workflows after deployment. For them, the agentic shift begins with the operating model.
- Data readiness, not data abundance, is what makes AI compoundable.** Most enterprises discover too late that having data and having AI-ready data are very different things. A sovereign, composable foundation – one that queries and prepares data where it resides, without migration or centralization – can convert raw data estates into intelligence that AI agents can act upon. As data residency laws, multi-cloud environments, and structural complexity make centralization increasingly impractical, sovereign control over where models run and data lives is what keeps that adaptability intact.
- Intelligence confined to a function stops compounding.** Organizations generating meaningful AI returns embed it across strategy, workflows, people, and culture simultaneously. The breadth of integration is what allows value to accumulate rather than stall at departmental boundaries. This is the point at which composable stops being an architectural preference and becomes a direct business imperative: The ability to connect agents, models, and data across the enterprise without rebuilding from the ground up each time.
- Agentic AI makes the sovereign, composable bet decisive.** AI agents acting autonomously across systems – placing orders, escalating cases, and updating records in real time, for example – must work on current data wherever it lives and move fluidly across the heterogeneous tools that characterize most enterprises. That requires sovereign infrastructure and composable architecture built together, not bolted on after the fact.

02 Why enterprise AI struggles to scale

The returns from enterprise AI are real, but unevenly distributed and rarely compounding. Only 6% of organizations qualified as AI high performers in 2025, reporting significant financial returns, according to McKinsey.³ Despite billions invested in generative AI, research from MIT Media Lab's project NANDA suggests just 5% of integrated pilots deliver measurable business value. And a study by S&P Global found that nearly one in two companies abandon AI initiatives before reaching production.⁴


The proof of concept (PoC) model is partly to blame. Pilots succeed because they are designed to. Data is curated, scope is narrow and the teams running them are typically senior and motivated. The moment a working concept confronts real production conditions like fragmented data, legacy integration, compliance

reviews, and operational handoffs, a different reality surfaces. "It's not that 95% of AI pilots never reach production, it is that they are set up for failure," says Gerry Murray, research director of enterprise marketing apps and agents at IDC.

What looks like an AI problem is almost always an infrastructure and process problem wearing AI's clothes. These structural barriers could include fragmented data trapped across systems and formats, generic models that lack the contextual grounding to operate reliably in complex enterprise environments, and AI tools built for developers rather than the business users who need to act on the intelligence they produce.

Research by Informatica on generative AI deployments highlights this pattern. The biggest obstacles preventing initiatives from reaching production are data quality and readiness, cited by 43% of respondents, tied with gaps in technical maturity, followed by skills shortages at 35% (see Figure 2).⁶

Figure 1: Most companies are struggling to move AI pilot projects into production



of organizations have been unable to transition even half their generative AI pilots to production

Source: Compiled by MIT Technology Review Insights, based on data from Informatica, 2026⁵

"Until organizations treat data architecture and operational readiness as first-class citizens, production AI will continue to fail where PoCs succeed."

Cristopher Kuehl, Chief Data, Information, and AI Officer, Continent 8 Technologies



“Until organizations treat data architecture and operational readiness as first-class citizens, production AI will continue to fail where PoCs succeed,” says Christopher Kuehl, chief data, information, and AI officer at Continent 8 Technologies.

The key to scaling AI, then, is about data readiness, governance, and architecture, says Tushar Shah, chief customer officer at Uniphore. “Enterprises that close the gap do so by treating their data foundation as a prerequisite, not an afterthought. That includes automating data discovery and preparation across complex estates, building sovereign control over where data lives and models run and replacing brittle centralized pipelines with composable infrastructure that queries intelligence where it already resides.

“That is the foundation on which the agentic enterprise is built,” he adds. “Without it, the shift from pilot to production remains out of reach.”


Process design in an AI-speed world

Even organizations that acknowledge the PoC-to-production reality must contend with the pace of innovation. Enterprise technology evolved through institutional processes built for stability: budgeting cycles, procurement reviews, and staged deployments. The AI ecosystem moves at a vastly different cadence. “AI is moving so fast that even consulting firms, venture capitalists, startups, and researchers are still figuring it out,” observes Prasad Jayaraman, principal at KPMG. “Nobody has all the answers yet, and those answers will keep changing.”

“You can build as many AI and agentic workflows as you want, but their accuracy is only as good as the data and the model behind them.”

Tushar Shah, Chief Customer Officer, Uniphore

Figure 2: Top barriers preventing generative AI pilots from reaching production



Source: Compiled by MIT Technology Review Insights with data from Informatica, 2026¹¹

Kuehl characterizes the dynamic as the “tech world running at two different speeds.” One speed is the enterprise’s own: deliberate, governance-heavy, designed to minimize disruption. The other is the AI ecosystem’s, producing capability shifts monthly, with best practices still forming and the competitive consequences of falling behind compounding quickly.

That places process redesign at the center of adoption. The organizations reporting the greatest value from AI initiatives are treating process redesign as a variable they can control, even before selecting models.⁸ Instead of trying to predict where the technology will land, these organizations redesign architectural choices, workflows, and governance structures so they can evolve as the technology does.

Most workflows were designed for software that supported human judgment. Now, AI systems can increasingly act within those workflows. “AI is the first generation of technology that can take action on its own,” explains Jayaraman. “We are moving from systems of record and systems of intelligence to actual systems of action. This has never been contemplated, let alone done before.”

03 Architecture decisions for long-term AI value



Organizations often assume that because data already exists, it is ready for AI. However, that is not always the case. The distinction between data abundance and data readiness is a consequential and often underestimated gap in enterprise AI strategy. As Jayaraman puts it, leaders should be asking: “Is our data organized for the AI world? Can we capture institutional knowledge and tacit knowledge?”

Those questions matter more as AI moves from analytics to action. A reporting dashboard can tolerate a stale data point or a partial view, an AI agent making decisions in real time cannot. “You can build as many AI and agentic workflows as you want,” observes Shah, “but their accuracy is only as good as the data and the model behind them.”

This dependency runs through every layer of an effective AI stack, and the layering itself is deliberate. Shah describes an architecture that begins with accessible, governed data at the foundation and then

builds upward through a knowledge layer that captures shared context, domain expertise, and institutional meaning. Models selected and sized to specific tasks for accuracy, performance, and cost efficiency are added on top. And, finally, an agentic layer that enables systems to act autonomously with full observability.

Each layer depends on the integrity of the one below. This is why data weaknesses are so frequently misdiagnosed as model failures, and why model upgrades alone rarely fix what ails a struggling AI program. This architecture is also why sovereignty is not a compliance footnote but a structural requirement. As data residency laws tighten, multi-cloud environments multiply, and enterprise data estates grow more complex, centralizing data can become both technically impractical and legally risky.

The alternative is querying and preparing data where it already resides, without migration or replication. This can preserve the integrity of each layer while keeping the enterprise in control of its own intelligence.

“AI is the first generation of technology that can take action on its own. We’re moving from systems of record and systems of intelligence to actual systems of action. This has never been contemplated, let alone done before.”

Prasad Jayaraman, Principal, KPMG

“It’s not that 95% of AI pilots never reach production, it is that they are set up for failure.”

Gerry Murray, Research Director of Enterprise Marketing Apps and Agents, IDC

Dan Morris, global head of industry solutions for marketers at Databricks, pushes back on one common misconception that delays progress. “The idea that you have to have data in a pristine state before you can do anything causes unnecessary delays and missed opportunities,” he explains. In practice, organizations can begin extracting value earlier than they think, provided they establish “a single source of truth with a unified layer of governance, regardless of how your data and systems are organized.”

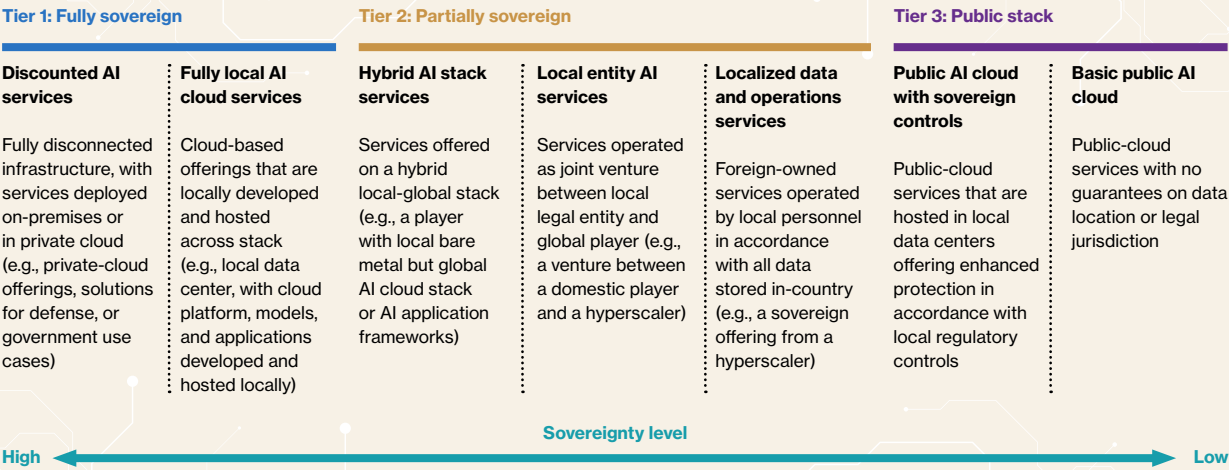
The vertical integration vs. composability trade-off

Two dominant architectural approaches sit at opposite ends of a spectrum. Vertically integrated stacks – where a vendor tightly couples model access, orchestration, data infrastructure, and deployment environments into a unified platform – can reduce

operational complexity. Composable architectures – systems built from interchangeable, independently upgradeable components connected through open interfaces – can trade some of that operational complexity for long-term adaptability. They allow enterprises to leverage technology advances by swapping a model, retrieval layer, or orchestration engine without dismantling the whole.

“Closed ecosystems succeed early because they deliver velocity. And for narrow use cases, that speed creates a real sense of confidence. In IT, confidence is everything,” says Kuehl. “But over time, you inherit their ecosystem, for better and for worse.” As organizations navigate a rapidly evolving AI ecosystem, the trade-off increasingly favors composable architectures that let models, tools, and agent capabilities evolve without large-scale rewrites.¹⁰

Figure 3: A tiered approach to AI workloads match sovereignty needs with the right level of control for each use case



Source: Compiled by MIT Technology Review Insights, based on data from McKinsey, 2026⁹


At the center of this shift sits the AI gateway, a control layer between applications and models that orchestrates multiple providers through a single interface. “The AI gateway allows you to swap out different models over time as the ecosystem evolves,” explains Morris, noting that the goal is flexibility similar to multi-cloud. That composability allows organizations to match AI capabilities to specific needs. In practice, that means aligning model scale with actual task demands. Smaller, purpose-built language models can handle well-defined, repetitive functions more efficiently, while more capable models earn their cost where broader reasoning is genuinely required.

The composable infrastructure market reflects this architectural shift. It is set to grow to \$21.9 billion by 2031 from \$8.4 billion in 2025, according to Mordor Intelligence, representing a compound annual growth rate of around 17%¹¹ (see Figure 4).


“The AI gateway allows you to swap out different models over time as the ecosystem evolves.”

Dan Morris, Global Head of Industry Solutions for Marketers, Databricks

Figure 4: Composable infrastructure market size



Source: Compiled by MIT Technology Review Insights with data from Mordor Intelligence, 2026¹²



Organizational choices for AI scale



Across industries, AI rollouts have followed a familiar arc. Leadership sets ambitious targets for productivity or cost reduction, teams move quickly to deploy foundation models, and departments begin building their own tools. Jayaraman, who has tracked the cycle across clients in multiple sectors at KPMG, describes the integration pattern that follows vertical adoption in a question: “How do we avoid 100 siloed tools, and make them work together?”

A Zapier survey found that three in four enterprises have already felt the consequences of disconnected AI, with compliance exposure, redundant spending, and lost employee time among the most common.¹³

Largely, this is a challenge of coordination that requires a governance architecture built for scale that allows AI deployments to talk to each other. This includes shared data products, common policy frameworks, and observability systems that monitor how models behave and how data moves through them, providing visibility without requiring sign-off at every decision point.

Deciding where to build, buy, or partner

Among organizations with sustained AI performance, a practical build-versus-buy principle is emerging: Buy foundational components that evolve quickly and require deep specialization and build the differentiating layer where competitive advantage lives.¹⁴ “Your competitiveness isn’t in the platform,” Kuehl argues. “It is in the last mile: your data signals, your workflows, your risk posture, your customer guardrails. If it reflects your brand identity, build it. If it can be commoditized, buy it.”

The vendor landscape increasingly supports that strategy. Gartner estimates that by 2026, more than 80% of independent software vendors will embed generative AI capabilities in enterprise applications, up from less than 1% in 2023.¹⁵

Shah advises organizations to avoid long-term lock-in in favor of multi-vendor designs with shorter commitment windows. “Avoid five- or six-year contracts,” he says. “A multi-vendor design with shorter commitment windows gives you the flexibility to continually evaluate and adopt the best technologies.”

“The idea that you have to have data in a pristine state before you can do anything causes unnecessary delays and missed opportunities.”

Dan Morris, Global Head of Industry Solutions for Marketers, Databricks

“Because data and business processes are always changing, the AI solutions built on top of them must evolve continuously. Sovereign architecture is what ensures that intelligence remains the enterprise’s own to build on.”

Tushar Shah, Chief Customer Officer, Uniphore

Recent research underscores the financial stakes: 84% of organizations report AI costs eroding margins by 6% or more, and 80% miss their AI forecasts by at least 25%.¹⁶ Securing capabilities with appropriate optionality – ensuring that procurement decisions made now do not become architectural constraints over time – has become as important as selecting the right tool in the first instance.

Intelligence that flows across functions

Even well-architected AI systems fail to deliver enterprise-wide value when intelligence stays confined to the function that generated it. Murray sees this most clearly in customer-facing operations, where functional excellence at one stage of the customer journey can be quietly dismantled downstream by teams with no visibility into what came before. “Customer experience only works when it spans every touchpoint,” he says.

Shah draws the same conclusion at an organizational scale. “It is not just engineering, it is not just customer experience or customer support. You want to transform legal, finance, and all adjacent functions, and they are all using different systems.” The ability to connect across those systems is what determines whether AI investments accumulate into organizational intelligence or remain productive only within the function that deployed them.

Deloitte research from 2025 confirms the pattern: Organizations reporting meaningful returns treat AI as an enterprise-wide transformation.¹⁷ Embedding integration across strategy, workflows, people, and culture simultaneously allows value to compound rather than stall.



05 Foundations for the agentic AI shift



Agentic AI imposes a diverse set of architectural demands. An AI agent placing an order or updating a customer record across systems must act on the current context or it risks compounding errors autonomously. This requires continuous, real-time data access across functions, reliable mechanisms for carrying context across handoffs between systems and people, and the capacity to trace, audit, and correct autonomous actions as they occur.

“Eventually, each agent needs to be able to hand off whatever it provided for the customer to the next agent, system, or person that’s going to encounter that customer,” observes Murray. That continuity of context distinguishes enterprise AI that compounds over time from AI that merely executes.

Already, more than one in three organizations run agentic AI as a primary AI model type, according to a 2025 Cloudera survey.¹⁸ But most organizations’ current governance cannot support AI agents operating at scale, according to Bain & Company’s 2025 analysis of enterprise AI architecture. The report identifies the shift from centralized platform ownership to distributed coordination as one of the defining challenges.¹⁹

Figure 5: Preparing IT architecture for agentic AI means considering the full stack



Source: Compiled by MIT Technology Review Insights with data from Uniphore, 2026²⁰

“Your competitiveness isn’t in the platform. It is in the last mile: your data signals, your workflows, your risk posture, your customer guardrails. If it reflects your brand identity, build it. If it can be commoditized, buy it.”

Cristopher Kuehl, Chief Data, Information, and AI Officer, Continent 8 Technologies

The competitive stakes of agentic AI delay

When it comes to rolling out agentic AI, organizations are feeling the pressure. One 2025 study by PwC finds that nearly three quarters (73%) of executives believe that how they deploy AI agents will determine competitive advantage in the next 12 months. Nearly half (46%) fear their organization may fall behind competitors.²¹ Tech-forward enterprises that moved early report operating profit gains of 10% to 25% after scaling AI across core workflows, according to Bain & Company.²²

As AI systems move from supporting decisions to executing them autonomously – transacting, orchestrating, and resolving, for example – the structure of competitive advantage shifts with them. The workflows, approval chains, and response cycles that shaped outcomes in a human-driven operating model increasingly act as bottlenecks once agents operate in real time.

The change is critical in preparing for agentic AI commerce as interactions are increasingly mediated by agents, says Gerry Murray, research director of enterprise marketing apps and agents at IDC. “In

agentic commerce, brands lose journey data but gain rich context data,” Murray explains. “Agentic marketplaces will function like stock exchanges with agents declaring programmatic intent. Brands must develop new sense-and-respond capabilities that function in real time at scale. The brand with the best offer, the fastest, wins the customer – and that customer is immediately out of market for everybody else.” The same dynamic appears across enterprise functions. The organization with the most relevant response, generated the fastest from the richest context, captures the outcome. “Brands that lack the infrastructure to keep up will be hard pressed to establish market share in agentic channels,” Murray adds.

In response, 5% to 10% of enterprise technology spending could be directed toward foundational agentic capabilities over the next three to five years, estimates Bain & Company. A figure that may look modest against their longer-term projection that up to half of all enterprise technology spend could ultimately run through agentic AI.²³

Interoperability is where that complexity is most acute in practice. Without a shared interface layer to mediate the heterogeneous systems that characterize most enterprises, AI agents may encounter friction that narrows both the scope of their actions and the impact those actions can have.

“The bigger issue comes down to how do I get it past security, legal, and other protocols,” Shah explains. “That’s a big hurdle for enterprise CIOs who want to move quickly.” New vendors require three to four months of internal approval – a cycle that could push CIOs toward incumbent vendors whose primary advantage is familiarity rather than suitability for agentic deployment.

Conclusion

Architectures that support growing intelligence share common properties. Data is treated as an appreciating asset rather than a byproduct of operations: democratized, governed, and continuously refined. Composable infrastructure allows models, tools, and agentic AI capabilities to evolve without rebuilding from the ground up. Sovereignty over that data and those models ensures that the intelligence an organization accumulates remains its own, deployable across any environment and able to compound over time.

That interdependence – between democratized data, composable infrastructure, and sovereign control – determines what enterprise AI can do. Organizational models determine what it does: How fast governance can move, where differentiation is built versus bought, and whether intelligence spreads across the enterprise or stays confined to the functions that generated it.

Endnotes

1. Rowland Manthorpe, "AI is developing so fast it is becoming hard to measure, experts say," Sky News, February 26, 2026, <https://news.sky.com/story/ai-is-developing-so-fast-it-is-becoming-hard-to-measure-experts-say-13512068>
2. "Gartner Says Worldwide AI Spending Will Reach \$2.5 Trillion in 2026," Gartner, January 15, 2026, <https://www.gartner.com/en/newsroom/press-releases/2026-1-15-gartner-says-worldwide-ai-spending-will-total-2-point-5-trillion-dollars-in-2026>
3. "The state of AI in 2025: Agents, innovation, and transformation," McKinsey & Company, November 5, 2025, <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>
4. Alex Johnston, "Generative AI shows rapid growth but yields mixed results," S&P Global, October 27, 2025, <https://www.spglobal.com/market-intelligence/en/news-insights/research/2025/10/generative-ai-shows-rapid-growth-but-yields-mixed-results>
5. "CDO insights 2025," Informatica, 2025, https://www.informatica.com/lp/cdo-insights-2025_5039.html
6. Preetam Kumar, "The Surprising Reason Most AI Projects Fail – And How to Avoid It at Your Enterprise," Informatica, March 31, 2025, <https://www.informatica.com/blogs/the-surprising-reason-most-ai-projects-fail-and-how-to-avoid-it-at-your-enterprise.html#a-lack-of-data-in-a-time-of-data-overload?>
7. "CDO insights 2025," Informatica, 2025, https://www.informatica.com/lp/cdo-insights-2025_5039.html
8. "The state of AI: How organizations are rewiring to capture value," McKinsey & Company, March 12, 2025, <https://www.informationweek.com/data-management/how-much-data-is-too-much-for-organizations-to-derive-value>
9. "The sovereign AI agenda: Moving from ambition to reality," McKinsey & Company, December 18, 2025, <https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/tech-forward/the-sovereign-ai-agenda-moving-from-ambition-to-reality>
10. "Enterprise Technology Report: from Pilot to Production," MACH Alliance, January, 2026, <https://machalliance.org/enterprise-technology-report>
11. "Composable Infrastructure Market Size & Share Analysis," Mordor Intelligence, <https://www.mordorintelligence.com/industry-reports/composable-infrastructure-market>
12. "Composable Infrastructure Market Size & Share Analysis," Mordor Intelligence, <https://www.mordorintelligence.com/industry-reports/composable-infrastructure-market>
13. Dylan Reber, "Tool sprawl limits AI integration for 70% of enterprises," Zapier, December 3, 2025, <https://zapier.com/blog/ai-sprawl-survey/>
14. "The Evolution of Build vs Buy," KPMG, <https://kpmg.com/uk/en/insights/ai/the-evolution-of-build-vs-buy.html>
15. "Gartner CFO & Finance Executive Conference: Day 1 Highlights," Gartner, September 11, 2024, <https://www.gartner.com/en/newsroom/press-releases/2024-09-11-gartner-cfo-and-finance-executive-conference-london-day-1-highlights>
16. "2025 State of AI Cost Management Research Finds 85% of Companies Miss AI Forecasts by >10%," PR Newswire, September 10, 2025, <https://www.prnewswire.com/news-releases/2025-state-of-ai-cost-management-research-finds-85-of-companies-miss-ai-forecasts-by-10-302551947.html>
17. Richard Horton et al., "AI ROI: The paradox of rising investment and elusive returns," Deloitte, October 22, 2025, <https://www.deloitte.com/uk/en/issues/generative-ai/ai-roi-the-paradox-of-rising-investment-and-elusive-returns.html>
18. "Enterprise AI and Data Architecture in 2025: From Experimentation to Integration," Cloudera, September 25, 2025, <https://www.cloudera.com/blog/business/enterprise-ai-and-data-architecture-in-2025-from-experimentation-to-integration.html>
19. Paul Gautheron et al., "Building the Foundation for Agentic AI," Bain & Company, September 23, 2025, <https://www.bain.com/insights/building-the-foundation-for-agentic-ai-technology-report-2025/>
20. Uniphore's "The Uniphore Business AI Cloud," Uniphore, <https://www.uniphore.com/platform/business-ai-cloud/>
21. "PwC's AI Agent Survey," PwC, May 16, 2025, <https://www.pwc.com/us/en/tech-effect/ai-analytics/ai-agent-survey.html>
22. David Crawford et al., Bain & Company, "State of the Art of Agentic AI Transformation," September 23, 2025, <https://www.bain.com/insights/state-of-the-art-of-agentic-ai-transformation-technology-report-2025/>
23. Pascal Gautheron et al., Bain & Company, "Building the Foundation for Agentic AI", September 23, 2025, <https://www.bain.com/insights/building-the-foundation-for-agentic-ai-technology-report-2025/>

About MIT Technology Review Insights

MIT Technology Review Insights is the custom publishing division of *MIT Technology Review*, the world's longest-running technology magazine, backed by the world's foremost technology institution – producing live events and research on the leading technology and business challenges of the day. Insights conducts qualitative and quantitative research and analysis in the US and abroad and publishes a wide variety of content, including articles, reports, infographics, videos, and podcasts. This content was researched, designed, and written entirely by human writers, editors, analysts, and illustrators. This includes the writing of surveys and collection of data for surveys. AI tools that may have been used were limited to secondary production processes that passed thorough human review.

About Uniphore

Uniphore is the Business AI Company that unlocks the agentic enterprise with a complete, composable AI platform spanning agents, models, knowledge, and data. Its platform, the Business AI Cloud, bridges the AI divide between consumer AI and enterprise AI – combining the simplicity of consumer AI with the rigor, security and scalability required for the enterprise. Uniphore allows business users to effortlessly harness AI and deliver results immediately, while providing CIOs the foundation to deliver powerful AI applications that are embedded into workflows, trained on enterprise data.

Trusted by more than 2,000 businesses globally, recognized by Gartner, Forrester, and listed on the Deloitte Fast 500, Uniphore delivers on the promise of AI as a transformative force for business. Learn more at www.uniphore.com.



Illustrations assembled by Tim Huxford with elements from Shutterstock.

While every effort has been taken to verify the accuracy of this information, MIT Technology Review Insights cannot accept any responsibility or liability for reliance by any person in this report or any of the information, opinions, or conclusions set out in this report.

© Copyright MIT Technology Review Insights, 2026. All rights reserved.

To cite this report, please use: "Redefining enterprise intelligence with autonomous AI," MIT Technology Review Insights and Uniphore, August 2026.



MIT Technology Review Insights

www.technologyreview.com

insights@technologyreview.com