

JUNE 2026

Arrested Automation:

Why Agentic AI Stalls
at the Enterprise Level

Table of Contents

- 03** Executive summary
- 04** Key findings
- 05** The Agentic AI Maturity Index
- 06** The investment surge vs. the ROI deficit
- 09** Context fragmentation
- 11** Spotlight: The vertical divide in AI readiness
- 17** Spotlight: The regional landscape
- 20** Crossing the action bridge
- 22** Autonomous knowledge and the move from personal to organizational AI
- 24** Conclusion
- 25** Methodological notes

Executive summary

The first wave of enterprise AI brought incredible tools for personal productivity, sparking a wave of excitement and innovation across the business world.

Now, the market has begun its next major evolution: moving from individual assistants to complex, organizational AI that can safely manage operational decisions.

A global survey of 1,000 senior technology and data leaders across six markets reveals that while the enthusiasm to deploy AI is universal, our foundational data systems need rethinking to fully support this leap.

The primary opportunity lies in understanding the difference between two concepts: personal AI and organizational AI.

Personal AI includes tools such as chatbots that help individuals work faster, while organization AI works on behalf of the whole company, using shared knowledge, appropriate access levels, and well-designed governance.

The ROI that enterprises want to achieve doesn't happen until AI operates at the organizational level, but the hurdles are significant:



The ROI gap

Translating high initial investments into measurable operational outcomes requires connecting AI directly to core business systems.



Context fragmentation

Enterprise data needs to be enriched with business meaning, lineage, and governance so AI tools can understand the "big picture" at scale.



The action bridge

Getting AI to act requires solving for where output lives, if decisions are traceable, and whether the underlying data can be trusted.

The path forward looks bright and achievable, but it requires a shift toward **Autonomous Knowledge** — enterprise data that carries enough context, governance, and meaning for agents to act on it reliably and repeatedly, while learning and improving over time. This allows AI to graduate from simply offering helpful suggestions to performing intelligent actions that drive the business forward.

Key findings



90%

of leaders **expect to increase their agentic AI investments** in the next year

80%

of leaders **view the rise of agentic AI primarily as a competitive opportunity** to strengthen their position, rather than a threat (20%)

93%

of leaders **believe AI will eventually run core functions**, and nearly a third (32%) expect this shift to happen within the next two years

77%

of executives report that **20% or less of their data is sufficiently described and contextualized**

62%

of leaders **prioritize enterprise-wide ROI and business transformation** when evaluating AI investments, yet margin improvement as a success metric is only used by 30%

69%

of C-level executives **believe their organization is already running agentic AI**, but only 57% of VPs believe the same

60%

of leaders **experience some decision paralysis on technology decisions** due to a desire for long-term flexibility

40%

find that more than 40% of their **AI pilot projects pause before production** because their infrastructure is not ready for autonomy

45%

of organizations **apply a specific risk assessment for each AI use case**, demonstrating an advanced understanding of AI governance

The Agentic AI Maturity Index

The Agentic AI Maturity Index tracks how organizations are successfully growing their AI capabilities

69%

of C-suite executives say their organization is already operating with agentic AI, while only 57% of VPs say the same

1

Experimenting

Organizations here are exploring localized pilot projects. They are learning how AI helps individuals work better and are beginning to map out their broader data strategies.

2

Developing

These companies have successful models but no shared mechanism to assemble the right context for the right agent on behalf of the right process. Each team has its own definitions, its own data, its own rules. Navigating this personal-to-organizational transition is the defining challenge at this stage.

3


Building

These enterprises have structured local governance and automated basic workflows, but the data foundation underneath hasn't been built for organizational scale. The barrier here isn't frameworks – it's whether the underlying data can support agents operating across functions, not just within them.


4

Operationalizing


These organizations have made deliberate architectural investments in harmonizing significant portions of their data, established dynamic safety rules, and are starting to see AI confidently execute multi-step workflows. They are actively developing Autonomous Knowledge: data with enough context, lineage, and governance for agents to act on it reliably. Early results show measurable impact, but extending that foundation across the full enterprise is ongoing work.




of the market



of the market



of the market



of the market

The investment surge vs. the ROI deficit

90%

of senior technology leaders report that they expect to increase their agentic AI investments over the next 12 months.

Corporate investment in agentic AI is robust and optimistic, with 80% of senior leaders seeing agentic AI as a competitive opportunity to elevate their market position. But the enthusiasm and the spending on agentic are not translating into positive business outcomes easily or quickly. Despite aggressive investment, nearly two-thirds of tech leaders (63%) admit they haven't seen more than an emerging, minimal positive return on their AI investments to date.

This AI ROI gap is easier to understand when we separate personal wins from organizational ones. While helping employees write code or draft emails faster is helpful, individual productivity alone doesn't justify massive infrastructure investments.

“

Individual productivity gains — faster code, better drafts, quicker research — are real but they don't show up on the P&L in a way that justifies significant infrastructure investment. The ROI that executives expect requires agents operating at the organizational level: automating decisions, executing workflows, driving measurable business outcomes. Those returns require a completely different foundation. Most organizations are measuring enterprise AI ROI against personal AI infrastructure — and wondering why the numbers don't add up.

Louis Landry
Chief Technology Officer
Teradata


Enterprise ROI requires leaders to let AI step up and automate workflows, make operational decisions, and drive measurable business outcomes. The disconnect happens when companies try to achieve organizational goals using a foundation built only for personal AI.

Question:

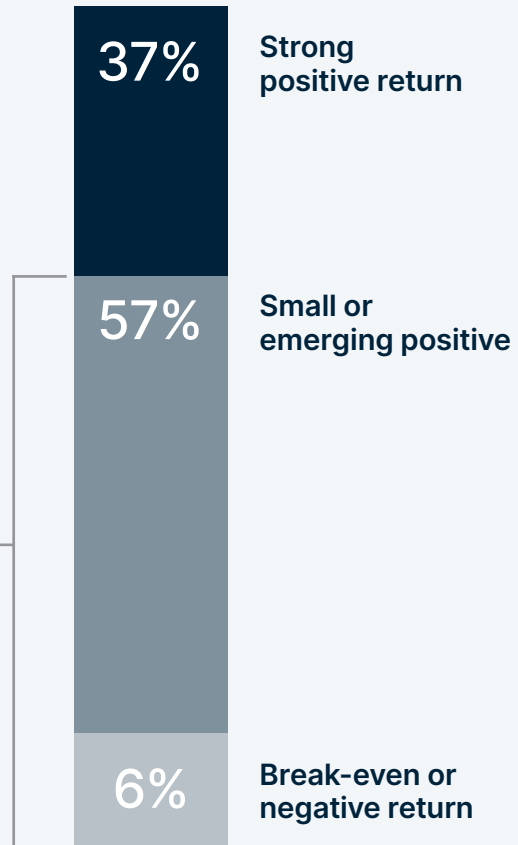
How would you characterize the return your organization has realized from its agentic AI investments to date?



Answer:



have yet to realize meaningful returns from their AI investments



The gap is also an absorption problem. When AI makes individuals faster, that freed capacity doesn't automatically move the bottom line for the organization. Someone in leadership has to decide where it goes — toward revenue, adding skilled headcount, or something else measurable. Until that reallocation happens deliberately at an organizational level, productivity gains create more activity without changing the numbers that corporate boards are watching.

The data reflects this: 62% of leaders prioritize enterprise-wide ROI over individual productivity gains when evaluating AI investments, yet margin improvement as a success metric is only measured by 30%.

Why is this happening? Organizations stuck in the Experimenting and Developing stages of the Maturity Index are burning through their IT budgets on disconnected pilot programs. They are purchasing expensive software licenses and investing heavily at the top of the tech stack, focusing on licensing large language models (LLMs) and building advanced user interfaces.

This spending frenzy is partially driven by a fear of missing out, and recent market volatility has also intensified high-stakes questions in the boardroom about AI investments.

As foundational models grow increasingly powerful, some enterprises still fear their core software services and institutional knowledge could be quickly replicated and commoditized by a simple AI wrapper. To survive and make themselves essential, companies must leverage the only unique, defensible asset they possess that the foundational model companies do not have: their proprietary, historical enterprise data. Unfortunately, that data is not prepared for AI ingestion, which often means that the millions being spent on the software layer aren't returning value.

Autonomous Knowledge is what enterprise data looks like when it carries enough context, lineage, and business meaning for AI agents to act on it reliably, not just access it.

The spending surge and the return gap aren't two separate problems; they are the exact same problem viewed from different angles. Companies are investing heavily in the model and software layer while the foundation underneath was never built to support what they're asking it to do. Until the foundation is fixed, the ROI gap will only widen.

“

In an enterprise there is no single truth. Sales, finance, and billing each count revenue differently, and each is right. The first problem is picking the definition the agent should use. The hard problem is assembling the rest of what the decision needs and knowing what to leave out.

Louis Landry
Chief Technology Officer
Teradata

The ROI gap threatens to widen as

77%

of leaders report that only 20% or less of their enterprise data and knowledge is ready for AI agents to act on reliably

Context fragmentation

The ROI gap's root cause is the data foundation, but more specifically, it's context fragmentation.

Organizations have no shortage of data, and in many cases, AI even has access to it. For an agent to act reliably however, the right data must be assembled with enough meaning and governance for the agent to understand what it's looking at. Enterprise data foundations were never built to do that because humans didn't need it.

Context fragmentation is what makes the leap from personal AI to organizational AI so difficult. Building an agent that works for one person is relatively straightforward. Building one that works across an organization requires significantly more effort because getting context right means knowing which data, assembled in what order, on behalf of which process, needs to reach an agent at the moment of a decision.


77% of executives report that 20% or less of their enterprise data is sufficiently described and contextualized – and 78% find it challenging to unify data and knowledge across business functions so agents can reason across the full enterprise.

When leaders rank the barriers to deploying agents, data lacking the necessary metadata, context, and relationships (43%) ranks nearly equal to data simply being fragmented across systems (42%). The problem isn't just where the data lives, but also whether it carries enough meaning to be useful and instructive.

Question:

How challenging is it for your organization to unify data and knowledge across business functions so that agentic AI agents can reason across the full enterprise context?

Answer:



report difficulty creating a connected data foundation for agentic AI



The costs of AI failure can be significant, evaporating millions of dollars and specialized labor hours without the expected operational outcomes. For the 68% of companies currently trapped in the Experimenting or Developing phases of the AI Maturity Index, context fragmentation is often the reason why. These organizations build highly successful, localized proof-of-concepts using clean, sampled data sets, but the moment engineering attempts to push that model into live enterprise production, the AI begins to struggle.

Currently, 40% of tech leaders find that more than 40% of their AI pilot projects fail simply because infrastructure systems were never built for this paradigm.

Infrastructure was built for reporting and analysis, not for agents that need to query, reason, and act across the full enterprise in real time. Adapting the data foundation to handle the workloads, governance requirements, and usage patterns that agentic systems create is a new requirement. Once that foundation is in place, organizations can begin to close the gap between what AI knows and what it can actually do.

This is the Autonomous Knowledge problem in concrete terms. Data without consistent definitions, traceable lineage, and embedded governance isn't agent-ready, regardless of how much of it exists or how accessible it is.

“

The goal of contextualizing your entire data estate is likely the wrong goalpost, and chasing it is part of why organizations stall. Data estates are always changing: new types, new sources, multimodal inputs. You'll never finish. Instead, identify the highest-value portion of your data — structured or unstructured — and focus on getting that portion fully described, governed, and agent-ready. If most of the data is unusable, the answer isn't to fix all of it at once. It's to be ruthlessly selective about which 20-50% you start with.

Josh Fecteau

Chief Data and AI Officer
& Chief Information Officer
Teradata

40%

of tech leaders find that more than 40% of their AI pilot projects fail simply because infrastructure systems were never built for this paradigm



SPOTLIGHT

The vertical divide in AI readiness

This data fragmentation challenge is universal, but it manifests differently depending on the specific regulatory, architectural, and operational hurdles of individual industries.

An analysis of the Agentic AI Maturity Index reveals striking disparities in how different verticals are navigating the transition to autonomy.

The Burden of Sensitivity

The healthcare sector* is grappling with its own deeply entrenched, and often deliberate, data fragmentation.

Highly sensitive and legally governed patient records, imaging data, and operational metrics are scattered across incompatible Electronic Health Record (EHR) systems, siloed hospital networks, and sovereignty boundaries that legally prevent data from moving at all.

A staggering 90% of healthcare leaders report that 20% or less of their data is sufficiently described and contextualized for agents to act on it reliably. As a result, healthcare organizations lag noticeably behind the general market in agentic AI maturity, appearing most frequently in the Experimenting stage of the Maturity Index (40%)

The inability to safely contextualize fragmented data makes it harder for healthcare providers and payers to move beyond basic, low-risk administrative AI use cases, like billing or scheduling. The ethical and legal risk of an AI agent acting autonomously on incomplete or fragmented medical history is simply too high.

An agent flagging drug interaction risks can't do its job if pharmacy records and patient charts sit in separate, incompatible systems.

Reflecting this caution, 43% of healthcare organizations cite governance, security, or access restrictions as a primary barrier limiting agent access to data, and 42% intentionally restrict their AI agents to analysis and recommendations where humans must execute the final action.



42%

intentionally restrict their AI agents to analysis and recommendations where humans must execute the final action



90%

of healthcare leaders report that 20% or less of their data is sufficiently described and contextualized for agents to act on it reliably

*Small base size; findings are directional



The Physical-Digital Divide

Manufacturing organizations are pushing hard to scale.

There are 40% currently in the Developing stage and around a quarter (24%) in the Building stage. The use cases are already well-defined as optimizing global supply chains, predicting equipment maintenance via IoT sensors, and automating quality control. Currently, 37% of manufacturing leaders prioritize supply chain, logistics, and procurement for autonomous AI deployment, but they haven't yet solved for fragmentation challenges.

Their data is split between modern cloud systems and decades-old operational technology that extends to the factory floor itself. A sensor predicting equipment failure on a production line has no way to automatically trigger a procurement order if those systems have never been connected.

Because these physical-to-digital data pipelines remain disconnected, 54% of manufacturing leaders cite AI model performance accuracy and reliability as a top barrier to deployment, alongside an equally severe shortage of AI and data specialists (54%).

37%

of manufacturing leaders prioritize supply chain, logistics, and procurement for autonomous AI deployment

54%

of manufacturing leaders cite AI model performance accuracy and reliability as a top barrier to deployment

87%

of leaders in manufacturing are the most likely to view agentic AI as a competitive opportunity

The Personalization Bottleneck

Retailers* sit in the middle of the pack, with 36% in the Developing stage and 32% in the Building stage.

They have massive volumes of consumer information, but it is often siloed across e-commerce platforms, brick-and-mortar point-of-sale systems, and third-party logistics providers. 46% of retail leaders prioritize customer engagement and service for autonomous deployment, but fragmentation is making that difficult to deliver. An agent trying to autonomously reorder inventory can't make a reliable decision if its demand

data stops at the e-commerce platform and doesn't include in-store sales. The same fragmentation undermines personalization because an agent launching a targeted campaign doesn't know the customer it's targeting returned their last three orders.



This is a critical vulnerability given that

48%

of retail organizations currently rely on customer satisfaction and experience metrics as their primary way to measure the ROI of their agentic AI

The Agentic AI Maturity Index by Industry

	Experimenting	Developing	Building	Operationalizing
Total	28%	40%	25%	7%
Financial Services	32%	33%	28%	7%
Manufacturing	28%	40%	24%	8%
Healthcare	40%	43%	15%	2%
Information Technology	23%	37%	27%	12%
Retail	28%	36%	32%	5%

*Small base size; findings are directional

The Compliance Anchor

Banks and wealth management firms are under pressure to deliver AI returns quickly.

With 65% prioritizing measurable enterprise-wide ROI over individual productivity gains, and 44% citing potential ROI as the leading factor when evaluating whether to expand an agent's scope.

But financial institutions face a particular version of the context fragmentation challenge. Their strict international compliance requirements, rigorous audit trails, and massive volumes of daily data make the stakes of getting something wrong unusually high. A fragmented data foundation could mean failing a regulatory audit, issuing a bad loan, or missing a fraud signal worth millions in real time.

A fraud detection agent monitoring thousands of transactions simultaneously is only as good as its ability to connect behavioral patterns across accounts, geographies, and transaction types. If that data lives in separate systems that can't communicate, the signal disappears.



65%

prioritize measurable enterprise-wide ROI over individual productivity gains



44%

cite potential ROI as the leading factor when evaluating whether to expand an agent's scope



50%

of financial services organizations report that governance, security, or access restrictions fundamentally limit their agents' access to enterprise data

The Industry Vanguard

Unsurprisingly, the IT* sector leads the pack, with the highest concentration of Operationalizing organizations (12%) and gaining ground in the Building stage (27%).

IT companies understand the requirements of modern data architecture and recognized early on that fixing the foundation was a prerequisite for AI deployment. They are heavily focused on backend stability, with 54% prioritizing IT, security, and infrastructure management for their autonomous deployments.



54%

prioritize IT, security, and infrastructure management for their autonomous deployments



Even in this sector

20%

of IT leaders cite infrastructure that simply cannot process data fast enough for agents operating at scale as the top ranked barrier to making data usable

*Small base size; findings are directional
Report prepared for Teradata

SPOTLIGHT

The regional landscape

The survey evaluated leaders across six distinct global markets: the United States, the United Kingdom, France, Germany, Japan, and Saudi Arabia.

This geographic diversity reveals that the challenges of agentic AI adoption are not isolated to Silicon Valley; they are a truly global phenomenon.



 **France and Germany**

In European markets like France and Germany, GDPR and the EU AI Act have pushed organizations to govern and document their data carefully by legal requirement. This effectively mandated a baseline of data sovereignty that other markets are still working toward, and the requirement has had an impact. Only 50% of French leaders and 54% of German leaders report decision paralysis, among the lowest of any surveyed market. French leaders are also less likely to report that data unification is challenging for their organization, just 52% compared to 84% in the US.

Compliance pressure has inadvertently pushed many European enterprises toward the kind of clean, contextualized data foundation that agentic AI requires. While this cautious approach may initially slow their pilot velocity compared to North American counterparts, it positions them well to see results.

 **United Kingdom**

The United Kingdom, however, stands in a different spot than its European counterparts. At 87%, UK leaders report the highest rate of difficulty unifying data across business functions – the most of any market surveyed and 35 points above France. Data fragmentation and infrastructure challenges are slowing deployment despite intent to move forward. Possibly contributing to this: post-Brexit, the UK is navigating data governance without the full structural backing of the EU framework.





66%

of US leaders report decision paralysis on durable infrastructure decisions

The United States

The United States leads markets on investment and ambition, but also on paralysis. 66% of US leaders report decision paralysis on durable infrastructure decisions, the highest of any surveyed market. The pattern is consistent with a market spending aggressively while the data foundation remains unaddressed. It has the highest proportion of leaders reporting pilot failures, 45%

Japan and Saudi Arabia

In Japan and Saudi Arabia, the push for AI autonomy is tied to national economic transformation initiatives. Both markets are heavily investing in building digital infrastructure from the ground up, which creates the opportunity to avoid tech debt.

At 52% and 51% respectively, decision paralysis for leaders in those countries is lower than for those in the US, but sovereignty is an active constraint in both markets. AI sovereignty, data sovereignty, and compute sovereignty shape which data can move, where it can be stored, and which infrastructure can be trusted. Organizations building their foundations now must consider those controls.

Crossing the action bridge

Even when organizations make progress on context fragmentation, they face activating that intelligence with autonomous action.

Today, AI is generating insights, such as summarizing 100-page legal documents, predicting quarterly sales trends, and drafting internal communications. But it isn't acting on them, and that gap between insight and execution is where much of the real-world value disappears.

This isn't a governance problem. Nearly half of organizations (45%) already apply a specific risk assessment for each individual AI use case. They understand the stakes, have frameworks in place, and know when humans should remain in the loop.

74% of leaders believe AI will be capable of running core business functions with minimal human oversight within four years. Almost none of them are letting it do that today.


Perhaps it's because 60% of leaders report decision paralysis on durable infrastructure decisions. The hesitation isn't really about technology selection. It's about confidence in what's being deployed.

Question:

How soon do you believe AI tools will be capable of running a core business function with minimal human oversight?



Answer:



anticipate AI will be capable of running core business functions within 4 years

Until organizations trust the data their agents are operating on, they won't let those agents act autonomously. In this way, it comes full circle to the context fragmentation issue, which is causing the automation hesitation and feeding the investment and ROI mismatch.

There's also a location problem. Most AI output currently lives outside the systems where consequential work actually happens, such as within ERPs, CRMs, and service management platforms. Even when the insight is good, a human has to manually carry it into the workflow. When intelligence is surfaced inside the tool where someone is already working, action follows naturally. When it lives in a separate dashboard, it competes for attention and usually loses.

Two additional concerns are keeping organizations from crossing the action bridge.



As long as AI output lives outside core systems, even good insights require a human to bridge them manually into a process, and that bridge is fragile and inconsistent.

Josh Fecteau

Chief Data and AI Officer
& Chief Information Officer
Teradata

The first is traceability.

As agents move from generating recommendations to executing decisions, the bar for accountability rises. An agent that approves a vendor payment, reroutes a supply chain, or denies a customer claim has to leave a trail that holds up to a regulator, a board, or an audit. That requires data with lineage built in, not bolted on after the fact. 51% of leaders cite accuracy and reliability of outputs as a significant deployment barrier, and that concern is directly connected to whether the underlying data can be traced back to the source.

The second concern is flexibility

30% of leaders cite vendor lock-in as an active worry, which is a rational response to a market where the leading model today may be obsolete in six months. The organizations navigating this best aren't trying to pick the right vendor. They're building their data into reusable, governed assets that any model can reason over. When enterprise data carries its own context, lineage, and business meaning, it stops being dependent on the tool consuming it. The foundation becomes portable, and the flexibility problem largely solves itself as it becomes Autonomous Knowledge.



Autonomous knowledge and the move from personal to organizational AI

The core problem running through the ROI gap, context fragmentation, and the action bridge is the same: most enterprise AI is currently built for individuals, and it needs to be rebuilt for organizations.

The meaningful target in making that happen is Autonomous Knowledge: a foundation that stays current, contextually rich, and governed over time without constant manual intervention.

Where to start depends on where you are. Organizations in the Experimenting and Developing stages should focus almost entirely on Phase 1 because the foundation has to come before anything else.

Those in the Developing stage are beginning to connect what they've built and will find Phase 1 unfinished work and Phase 2 starting to become relevant. Organizations in the Building stage likely have some context in place and are ready to tackle governance seriously.

“

The organizations seeing real, quantifiable value from agentic AI share a few traits: they've identified high-value, reusable context; they've built in autonomous learning so the system improves over time; and they've deployed at scale across domains, not just a single use case. Those patterns compound in a way that individual productivity tools don't.

Josh Fecteau

Chief Data and AI Officer
& Chief Information Officer
Teradata

Those approaching Operationalizing are the ones who can begin Phase 3 and start making the architectural bets that will determine their flexibility two or three years from now.

PHASE ONE



Audit and contextualize the data estate

Start by identifying the highest-value portion of your data estate and make that portion described, governed, connected, and genuinely agent-ready. The goal is to build shared context across people and processes. The data must carry consistent definitions, lineage, and business meaning across the organization, not just within a single team or function. This is what helps solve context fragmentation.

PHASE TWO



Embed dynamic governance at the source

Once that foundation is visible, embed governance directly into the data itself rather than the application layer. Policy enforcement applied at the software level is fragile. Governance that scales has to be built into the foundation, so that whatever model or agent is querying the system, the guardrails hold and decisions are traceable. This is what makes agent decisions explainable and the action bridge crossable.

PHASE THREE



Build for portability


Architectural decisions made today shouldn't become liabilities when the market shifts or the organization grows. A foundation that doesn't require rebuilding is what allows executives to make long-term bets with confidence, without disrupting the underlying data governance or enterprise context. This is what resolves the flexibility concerns that 30% of leaders report.

Conclusion

The argument this report makes is straightforward, even if the problem isn't.

Enterprise AI is stalling because the foundation underneath it was built to inform humans, not to ground autonomous systems operating at scale. What separates the organizations seeing returns isn't budget as much as it's the order of operations. The data foundation has to come before the software layer. Most enterprises have it backwards, and the ROI gap is the result. The shift from personal AI to organizational AI is the defining challenge of the next few years. Personal AI makes individuals faster, but organizational AI changes what a business can do.

The gap between those two things is a foundation of Autonomous Knowledge with enough context, lineage, and governance for agents to act on it reliably, across functions, on behalf of the whole enterprise.



The data foundation has to come before the software layer. Most enterprises have it backward, and the ROI gap is the result.



METHODOLOGICAL NOTES

The Teradata Survey was conducted by Wakefield Research among 1,000 senior technology and data leaders with a minimum seniority of vice president. Respondents were sourced from companies with a minimum of 500 employees across six key global markets: the United States (500), the United Kingdom (100), France (100), Germany (100), Japan (100), and Saudi Arabia (100). The survey was fielded between March 23rd and April 5th, 2026, utilizing an email invitation and an online survey instrument.

Results of any sample are subject to sampling variation. The magnitude of the variation is measurable and is affected by the number of interviews and the level of the percentages expressing the results. For the interviews conducted in this particular study, the chances are 95 in 100 that a survey result does not vary, plus or minus, by more than 3.1 percentage points in the total sample from the result that would be obtained if interviews had been conducted with all persons in the universe represented by the sample.

WAKEFIELDRESEARCH.COM

MARKET RESEARCH | PR POLLING | MARKET INTELLIGENCE | RESEARCH SUPPORT SERVICES