


Global AI Confessions Report: CEO Edition 2026



Executive summary

In 2025, CEOs feared falling behind in AI. In 2026, they fear something far more dangerous: being held accountable for it. AI is no longer an innovation story — it is a performance mandate.

AI is embedded in decisions, tied to revenue expectations, scrutinized by boards, and increasingly linked to CEO tenure. The pressure hasn’t eased, it has only intensified. In fact...




The mandate is clear. The consequences are no longer theoretical.

The AI accountability era has arrived — and it’s personal

Yet beneath that urgency, something is shifting. Confidence is not keeping pace with rising expectations. CEOs are being asked to deliver outcomes from systems they don’t fully control, can’t always explain, and increasingly question at scale.

Many still believe in AI’s potential to reshape markets — but belief is softening, and certainty is eroding.



AI is operating, but not deciding

This year’s confessions expose a stark contradiction at the heart of enterprise AI: It is everywhere, but control is nowhere. CEOs claim control of AI, but don’t fully command it. CEOs are using AI constantly — to analyze performance, guide strategy, and influence decisions — yet they are not ready to let it act independently.

Human oversight remains deeply embedded, not as a preference, but as a safeguard. The result is a new operating reality: AI is a daily operator, but not yet a decision-maker. Leaders trust AI enough to use it, but not enough to rely on it — and that gap is becoming operationally expensive.

Executive summary

At the same time, the narrative around enterprise AI investment is flipping. The fear is no longer about falling behind — it's investing in the wrong things. As AI spending accelerates, CEOs are becoming more cautious, more measured, and more focused on proving ROI. AI is no longer a blank check; it's a budget line item under scrutiny, with steadily rising ROI requirements.

But even as organizations push for efficiency, they remain structurally exposed to failure — driven by vendor concentration, opaque cost structures, and systems scaling faster than they can be governed. The risk is no longer just wasted investment, it is enterprise-wide instability at scale.

AI ownership is claimed, but not exercised

Perhaps most critically, AI leadership itself is fragmented. CEOs overwhelmingly claim ownership of AI strategy, but far fewer are involved in most decisions, and even fewer control the systems end to end.

This creates a widening gap between accountability and execution — one that is becoming untenable as AI moves from experimentation to production. And the stakes are rising fast:


The risk is no longer abstract. It has implications not just on companies, but also on the people accountable for AI decisions and performance. In short, it's personal.

This is the defining tension of enterprise AI in 2026: AI is driving growth, but it is also becoming the largest source of enterprise risk. Despite the hyper-fast AI innovation cycles we are witnessing, this is no longer a primary constraint — control is.

The companies that succeed in this next phase will not be the ones that adopt AI the fastest, but the ones that can govern it, measure it, and prove its impact under scrutiny.

For CEOs, the objective is clear, but the path is not. And, increasingly, their careers may depend on getting it right.

Key finding



77% of CEOs believe it is likely that a CEO will be ousted in 2026 due to a failed AI strategy or an AI-driven crisis.

CONFESSION #1

My job depends
on AI, but I'm less
certain than ever
I'll get it right





CONFESSION #1

My job depends on AI, but I'm less certain than ever I'll get it right

AI is no longer optional — and bad decisions are no longer forgiven.

For CEOs, the mandate is clear and uncompromising: Deliver measurable business outcomes from AI, or face the consequences. This is not a future risk. It is already shaping how leadership is evaluated.



But that confidence is colliding with a much harsher reality. The stakes are already well established: A clear majority of CEOs (80%) say their role is on the line if AI fails to deliver by the end of 2026, and three-quarters believe CEO turnover tied to AI missteps is increasingly likely.

The expectation is no longer progress — it is proof. And, increasingly, that proof must withstand scrutiny from boards, investors, and the market in real time.

CONFESSION #1

My job depends on AI, but I'm less certain than ever I'll get it right

Boards, in particular, are no longer passive observers.




The pressure isn't increasing, it's hardening into expectation.

Yet underneath this rising accountability, confidence is becoming more complicated than it appears.

On the surface, CEOs project control. **Eighty percent say they would trust their company's AI governance frameworks — even if their job were on the line.**

But that confidence is showing cracks. Fewer leaders now describe themselves as “extremely confident” in their ability to deploy AI agents in production and at scale, dropping from **41% in 2025 to just 31% in 2026**, while more shift into a middle ground of “very confident.” Confidence isn't collapsing, but it is clearly recalibrating.

And the exposure is not limited to execution alone.



CEOs are being asked to tie their tenure to a technology they believe could reshape industries, but also destabilize them.

This is the contradiction defining AI leadership in 2026.

CEOs are all-in on AI. They are prioritizing it, funding it, and tying it directly to business outcomes. But they are doing so in an environment where the systems are scaling faster than certainty and where the margin has all but disappeared.

AI has become a leadership referendum. And for many CEOs, the risk is no longer whether AI will fail, it is whether they will be held responsible if it does.

CONFESSION #2

I rely on AI for
decisions, but I
still don't trust it
to act alone




I rely on AI for decisions, but I still don't trust it to act alone

AI is no longer on the sidelines of decision-making. It is in the room and, increasingly, at the table.

For CEOs, AI has become a daily operating partner. It is shaping how they analyze performance, assess competitors, and guide strategic direction.

CEOs report AI influencing or informing, at some level, up to an average of 40 business decisions per year — underscoring just how embedded AI has become in executive decision-making. From company performance analysis to competitive intelligence and operational automation, AI is now embedded in the executive workflow.

And CEOs are not hiding it.




...signaling that AI's role in decision-making is no longer experimental: It is expected, visible, and increasingly normalized at the highest levels.

But influence is not the same as autonomy.

Despite AI's growing role, CEOs are still unwilling to let it act independently. **More than half (51%) keep humans in the loop when it comes to making business-critical decisions: Over one-third (34%) operate under a model where AI recommends, but humans approve.** The implication is clear: AI may inform decisions, but it does not own them — and CEOs are not ready for it to.

This has created a new operating model defined by constant verification. CEOs are not just using AI; they are interrogating it. **Most (80%) report at least occasionally asking for justification or explanation behind AI-driven recommendations,** reinforcing that trust in AI is conditional, not absolute.




Every output is reviewed, every recommendation is questioned, and every decision still carries a human checkpoint — a clear signal that enterprises cannot scale AI until they can trust it to act, and that trust depends on effective governance.

And yet, the ambition for AI continues to expand.

I rely on AI for decisions, but I still don't trust it to act alone

CEOs believe AI could fundamentally reshape the executive layer itself.




AI is no longer just a tool, it's being evaluated as a potential peer in leadership.

Even here, the foundation is starting to crack.

As AI moves closer to the executive layer, belief in its ability to outperform human leadership is weakening.

Key finding



AI is being positioned as a decision-maker, but confidence in its decision quality is slipping. Without stronger trust and governance, that gap will stall progress at the top of the organization.

At the same time, expectations for deployment remain aggressive. **Eighty-three percent of CEOs say they are confident their organization will deploy AI agents in full production in 2026**, reinforcing that companies are moving forward, even as underlying trust remains fragile.

This is the paradox CEOs now operate within. AI is embedded in decisions, but not empowered to make them. It is trusted enough to influence strategy, but not enough to execute it without oversight. Leaders are scaling AI across the enterprise while deliberately building the guardrails needed to ensure it performs reliably, responsibly, and at scale.

AI is no longer a passive tool, but it is not yet an autonomous operator. Until that gap closes, every decision AI touches carries a tradeoff: speed sacrificed for certainty, scaled autonomy constrained by manual oversight, and a growing tension between what AI can do and what leaders are willing to let it do.

I'm not afraid
of investing in
AI, I'm afraid of
investing in the
wrong vendors



I'm not afraid of investing in AI, I'm afraid of investing in the wrong vendors


For the past two years, the dominant fear around AI was simple: Move fast, or fall behind.

In 2026, that fear has flipped. The greater risk is no longer underinvesting, it's betting big on the wrong players.

Nearly two-thirds (65%) say they worry more about over-investing in AI amid intense vendor competition and no clear market leader, compared to just **35% who worry more about under-investing now out of precaution and falling behind competitors**.

The concern isn't spending itself — it's committing too early, locking into the wrong vendors, and carrying those decisions forward at scale. The pressure to spend remains, but confidence in where to place those bets is eroding.

AI is no longer a blank check. It's a budget line item.



Meanwhile, only 5% say it is not reinforcing that AI is now directly tied to valuation and growth expectations.

That shift is being driven by rising expectations around ROI. CEOs are under increasing pressure to prove that AI investments translate into measurable business outcomes and not just experimentation or efficiency gains. And those expectations are becoming more concrete.




But while expectations are sharpening, the underlying economics of AI remain difficult to control.

I'm not afraid of investing in AI, I'm afraid of investing in the wrong vendors

That lack of control is increasingly structural.

Key finding



76%

Over three quarters (76%) believe their organization is overly exposed to operational or strategic risk due to reliance on too few AI vendors.

At the same time, AI ecosystems are becoming more complex, with layered dependencies across models, platforms, and infrastructure providers.

In other words, as AI scales, control is becoming harder to maintain — even as risk continues to grow.

That tension is compounded by the pace of investment. Organizations are scaling AI faster than they can standardize it, locking themselves into vendor relationships and cost structures that are difficult to unwind. Pricing models remain opaque. Consumption patterns are unpredictable. And as capabilities evolve, so does the cost of staying competitive.

Even the perception of resilience may be overstated. While **89% of CEOs say they could keep core functions running if key AI systems were disrupted**, the growing concentration of vendors and platforms suggests otherwise. The deeper AI becomes embedded in business operations, the harder it becomes to absorb failure.


This is the new financial reality of enterprise AI. CEOs are not pulling back on investment, but they are becoming more cautious, more selective, and more focused on proof. The question is no longer “Are we investing enough?” It’s now “Are we investing in the right things and can we prove it?”

I say I own AI
strategy — but I'm
not the one making
the decisions that
will determine its
success



I say I own AI strategy — but I’m not the one making the decisions that will determine its success


On paper, AI leadership is clear: The CEO is in charge. In reality, it’s anything but.



The message is unambiguous: AI is a CEO-level priority, and leadership believes ownership sits at the top.

But ownership is not the same as control, and that gap is where risk begins.

Despite claiming responsibility, many CEOs are not deeply involved in the decisions that ultimately shape AI outcomes.




That gap creates a critical disconnect: CEOs are accountable for AI success, but not consistently embedded in the decision-making processes that determine it.

As AI scales, that disconnect becomes harder to manage. Decision-making is increasingly distributed across CIOs, data leaders, technical teams, and external vendors, with each influencing architecture, model selection, deployment, and governance.

CEOs may set direction, but execution is fragmented across layers of the organization. And when those layers misalign, accountability still rolls up to the top.

I say I own AI strategy — but I’m not the one making the decisions that will determine its success

That tension is already surfacing.



Vendor selection is no longer a technical choice, it’s a strategic bet with financial, operational, and reputational consequences. And, increasingly, CEOs are stepping in after the fact, not before the decision is made.

At the same time, the surface area of AI decision-making is expanding rapidly.

Nearly all CEOs (94%) say low- or no-code tools are critical to scaling AI creation across the workforce — enabling more teams to build, deploy, and experiment with AI. But that democratization comes with a tradeoff: More builders means more decisions, more variability, and more risk.

As AI creation spreads beyond centralized teams, control becomes harder to maintain, and oversight becomes more complex.

And the expectations are rising just as quickly.



This is the new reality of AI ownership. CEOs are expected to lead AI strategy, defend its outcomes, and absorb its risks. But the systems, decisions, and dependencies that determine success are increasingly distributed across the organization.

AI leadership is not centralized, it’s fragmented. And that creates a fundamental challenge: CEOs are accountable for outcomes they do not fully control.

CONFESSION #5

AI is scaling faster
than my ability to
control it




AI is scaling faster than my ability to control it

AI is no longer just a source of growth. It's becoming a source of risk.

As organizations move from experimentation to production, AI is scaling rapidly across workflows, systems, and decision-making processes. But while adoption is accelerating, control is not keeping pace. The result is a widening gap between what AI can do and what organizations can safely manage.

CEOs see this risk clearly.


More than half (57%) believe that insufficient AI explainability could trigger a crisis that erodes customer trust or damages their brand. And the concern is not isolated.



As AI systems take on more responsibility, the consequences of failure are no longer technical — they are reputational, legal, and financial.

Yet, despite these risks, organizations are continuing to push AI forward and are now running into real constraints.

Regulation is already beginning to slow momentum.



Fewer projects are being canceled outright (down to 23% in 2026 versus 32% from the prior year), but more are being paused, reassessed, or modified — signaling that governance challenges are no longer theoretical. They are actively shaping how and when AI is deployed.

At the same time, external pressures are mounting. **Eighty percent of CEOs agree that emerging regulations, including the EU AI Act, will slow AI adoption within their organization**, reinforcing that the path to scale is becoming more complex, not less.


AI is scaling faster than my ability to control it

But some of the greatest risks aren't coming from regulation, they are coming from within.

Shadow AI remains pervasive. **Ninety-six percent of CEOs believe employees are using generative AI tools without approval**, and **42% estimate that more than half of their workforce is doing so**.


This is on par with our previous research: 91% of data leaders suspect employees are using generative AI tools without notice or permission and over half of CIOs (54%) report they discovered employees using unsanctioned AI tools, apps, or platforms to do work tasks or projects.

AI creation is no longer centralized or fully visible. It is happening across the organization, often outside of formal governance structures, increasing the risk of data exposure, inconsistent outputs, and unmanaged dependencies. This is why governance is emerging as the defining factor of AI success.



Organizations know how to build AI. The challenge is managing it responsibly at scale.

And CEOs understand the stakes extend beyond performance.



AI is no longer just a business initiative. It is becoming the metric by which leadership will be remembered.

Because as AI expands, so does its blast radius. Decisions made by AI systems are increasingly visible to regulators, to customers, and to the market. Failures are no longer contained. They are amplified. And as AI becomes more deeply embedded in business operations, the cost of those failures grows.

This is the new reality of enterprise AI. AI is no longer just a competitive advantage. It is a risk multiplier — and the companies that succeed in this next phase will not be the ones that scale AI the fastest, but the ones that can control it, govern it, and defend it when it matters most.

Regional editions



REGIONAL EDITION

 U.S.

U.K.

France

Germany

UAE

APAC

Japan


AI accountability hits hardest in the U.S.

In the U.S., AI isn't just under scrutiny — it's under pressure from every direction at once. U.S. CEOs are navigating a uniquely intense environment where boards are more involved, expectations are higher, and the margin for error is smaller.

The result is a leadership reality defined less by experimentation and more by performance, where AI outcomes are expected to be measurable, defensible, and immediate.


Board pressure is higher and rising faster than anywhere else

The U.S. stands out most clearly in how quickly accountability is intensifying.




AI is the defining business priority

While AI is a priority globally, the U.S. market shows a sharper level of intensity.



That difference matters. It signals that U.S. CEOs are not just prioritizing AI, but are elevating it above competing business initiatives.

Key finding




That visibility extends to the board: **Nearly all (97%) of U.S. CEOs say they would feel comfortable telling their board that AI influenced a strategic decision — among the highest globally —** reinforcing that AI is not just embedded in decisions, but openly acknowledged at the highest levels.

Confidence is high but not necessarily accurate

U.S. CEOs project some of the highest levels of confidence globally when it comes to their AI strategy and roadmap, **second only to France in the share who describe themselves as “extremely confident”** (28% for the U.S. vs. 24% overall), but that confidence may be overstated.



But that confidence exists alongside elevated pressure and risk.



Risk awareness is elevated, especially around failure and crisis

Concern about AI-driven crises is particularly pronounced.



U.S.

 U.K.

France

Germany

UAE

APAC

Japan


Regulation and restraint are shaping AI in the U.K.

In the U.K., AI momentum is strong, but it is increasingly shaped by caution. CEOs are navigating rising regulatory complexity, growing delays, and a more measured approach to deployment.

The pressure to deliver remains but, compared to the U.S., the U.K. reflects a leadership posture that is slightly more constrained, more deliberate, and more influenced by governance realities.

AI is a top priority, but not without hesitation

The U.K. remains one of the most AI-prioritized markets globally.




However, that prioritization is paired with a slightly less concentrated sense of urgency at the very top end than in the U.S., where more CEOs specifically rank AI as the single top priority.

U.K. CEOs are also increasingly involved in AI-related decisions, with levels of engagement holding steady year over year — reinforcing that accountability continues to sit at the CEO level, even as execution becomes more complex.

At the same time, **U.K. CEOs are the most likely globally to worry about over-investing in AI (77%, vs. 65% overall) rather than under-investing,** reinforcing a growing concern that while AI is a top priority, not all investments will deliver equal value.

Regulation is not theoretical, it’s actively shaping deployment


The U.K. stands out most clearly in how regulation is influencing AI timelines.



This signals a clear shift from experimentation to constraint, reflecting how regulatory pressure is already slowing AI deployment.

Confidence is high, even as constraints increase

U.K. CEOs report some of the highest levels of confidence globally when it comes to their AI strategy and roadmap...



But that confidence exists alongside growing operational and regulatory constraints. As delays increase and regulatory pressure intensifies, U.K. leaders are navigating a more complex path to execution, creating a dynamic where conviction remains high, even as the environment becomes more restrictive.

U.S.

U.K.

 France

Germany

UAE

APAC

Japan

In France, AI progress is measured against regulation


In France, AI ambition is real — but it is increasingly shaped by regulation, scrutiny, and a need to defend decisions. CEOs are not pulling back from AI, but they are approaching it with more caution and more discipline.

Compared to other regions, France reflects a leadership posture where scaling AI is not just about capability, it's about proving it can withstand regulatory and operational pressure.

Regulation is slowing deployment more than anywhere else

France stands out as the market most impacted by regulatory uncertainty.


Key finding



French CEOs are more likely than those in any other region to report that AI initiatives have been delayed due to regulation (61%, vs. 49% overall), making it the clearest example of how compliance requirements are already shaping AI timelines.

Yet at the same time, concern about regulation slowing AI usage is declining.

Key finding




Just 26% of French CEOs now say EU regulations could slow AI adoption, down from 39% in 2025, suggesting that while regulatory friction is real, leaders are becoming more accustomed to operating within it.

This dynamic is closely tied to the influence of the EU AI Act, which is pushing organizations to slow deployment, reassess risk, and ensure explainability and governance before scaling.

Confidence is present, but less absolute

Key finding



French CEOs report strong confidence in their AI strategy and roadmap, including **the highest share globally who describe themselves as “extremely confident” (32%, vs. 24% overall)** — a signal of strong belief in their strategic direction.

However, that confidence is more uneven beneath the surface. Compared to markets like the U.S. and U.K., where no CEOs report a lack of confidence in their AI roadmap, France shows more variability, with some leaders expressing lower levels of certainty.

This creates a more polarized confidence profile: high conviction at the top end, but less uniform belief across the broader leadership base.

U.S.

U.K.

France

 Germany

UAE

APAC

Japan


In Germany, AI is defined by risk, not speed

In Germany, AI adoption is not lagging, but it is being approached with greater caution and scrutiny than in most other regions. CEOs are balancing pressure to deliver results with a heightened sensitivity to risk, cost, and long-term exposure.

Compared to global peers, German leaders are less focused on accelerating AI at all costs and more focused on ensuring that what they build is sustainable, governed, and defensible. The result is a distinct posture: AI is advancing, but under tighter control.

Investment anxiety is higher — and tilting toward under-investment


Germany stands out as one of the few markets breaking from the global investment narrative. While 65% of CEOs globally are more concerned about over-investing in AI, German CEOs are nearly split...



This puts Germany in a unique position. Rather than fearing wasted spend, German leaders are more evenly weighing the risk of moving too slowly against the risk of overcommitting. The result is a more cautious, deliberate investment posture, one that reflects uncertainty not just about cost, but about where and how to invest effectively.

Regulation is already impacting execution and shaping behavior


Regulatory pressure is already affecting how AI is deployed in Germany.



And leaders expect that pressure to intensify. Eighty-six percent say EU regulations could slow AI usage in their organization, tied with Singapore for the highest globally, signaling that regulatory friction is not just being felt today, but is expected to shape how AI scales going forward.

CEO involvement is rising, but with less intensity than other regions

While German CEOs are increasingly involved in AI decisions, the pace of that involvement is slowing compared to other markets.



At the same time, **the share of German CEOs who say they are frequently involved in AI decisions (51%-75% of them) has dropped sharply, from 43% in 2025 to just 29% this year,** signaling a move away from hands-on engagement.

This suggests a subtle but important shift. While accountability remains high, German CEOs are stepping back from day-to-day involvement as execution becomes more complex — or distributing responsibility more across the organization. AI leadership is not disappearing, but it is becoming less centralized.

REGIONAL EDITION

U.S.

U.K.

France

Germany

 UAE

APAC

Japan

In the UAE, AI is a career-defining bet

In the UAE, AI is not just a strategic priority — it is deeply personal. Compared to global peers, CEOs in the region are more likely to see AI as a risk to their legacy, not just an opportunity to enhance it. While ambition remains high, confidence in the upside is less certain, and awareness of downside risk is significantly more pronounced.

The result is a leadership environment where AI is being pushed forward aggressively, but with a sharper sense of what failure could mean.


AI is more likely to define — or damage — a CEO’s legacy

Nowhere is the personal impact of AI more pronounced than in the UAE.



At the same time, only 50% say AI is shaping their legacy (vs. 57% globally) and just 13% say it is securing it (vs. 24% globally). This duality is critical. In the UAE, AI is a reputational bet, where the upside is career-defining, but the downside is equally visible.


AI influence is more distributed, with a stronger role for the CDO



That creates a structural tension. While AI outcomes are closely tied to CEO legacy, strategic influence in the UAE is more distributed than in other regions (55% believe the CEO has the greatest influence on the company’s AI strategy, vs. 70% overall). In practice, that means CEOs are carrying the accountability for AI, while more of the decision-making power sits across the leadership team.

Confidence in explainability is lower than anywhere else

UAE CEOs are also less confident than their global peers in their ability to defend AI decisions.



This gap is significant. As AI becomes more embedded in decision-making, the ability to explain and justify outcomes is becoming a core requirement, not just a technical capability. In the UAE, that capability is less established, reinforcing the broader tension: AI is advancing quickly, but the systems to defend it are still catching up.

U.S.

U.K.

France

Germany

UAE

 APAC

Japan


In **APAC**, momentum
is high and so is
executive exposure

Across APAC (based on CEOs polled in Japan, Singapore, and South Korea), AI is advancing with strong conviction and urgency. APAC CEOs are more likely than their global peers to believe in AI's potential and to tie outcomes directly to leadership performance.

But that momentum comes with growing pressure as leaders balance aggressive adoption with rising concerns about over-investment, vendor exposure, and long-term differentiation.

Confidence in AI’s capabilities is higher than the global average

APAC CEOs show stronger belief in AI’s potential to operate at the highest levels of the enterprise than EMEA and the U.S.



This elevated confidence reinforces a more aggressive posture toward AI adoption where leaders are more willing to trust AI in core decision-making processes.

Accountability is higher and more directly tied to outcomes

That confidence is paired with heightened personal exposure.



In both instances, the APAC numbers are higher than both EMEA and the U.S.


AI is scaling, but under tighter human oversight

Despite strong belief in AI, APAC CEOs are not fully relinquishing control. **More than a quarter (26%) say AI never makes business-critical decisions autonomously (vs. 17% globally), while only 12% say AI often acts autonomously with human monitoring after the fact (vs. 19% globally).**

This reflects a more controlled approach to scaling where AI is expanding, but human oversight remains deeply embedded.

Investment risk and vendor dependence are growing concerns

As AI adoption accelerates, concerns about making the wrong investments are rising.



This highlights a growing tension: Moving fast is necessary, but doing so without creating structural risk is becoming increasingly difficult.

REGIONAL EDITION

U.S.

U.K.

France

Germany

UAE

APAC

 Japan


In **Japan**, AI progress is deliberate and not driven by urgency

In Japan, AI adoption is advancing but with more skepticism about execution than in any other region. CEOs are not resisting AI, but they are less confident that their strategies are translating into reality and less pressured by investors to accelerate.

The result is a more deliberate posture: AI is moving forward, but with greater caution about whether it is delivering as expected.

AI is advancing, but not without human control


Japan stands out for maintaining stronger human oversight as AI scales.



This reflects a more conservative operating model. AI is being deployed, but not fully trusted to act independently, reinforcing a preference for human-in-the-loop decision-making.

Investment posture is cautious — with similar structural concerns

Japanese CEOs show a slightly more cautious approach to AI investment.



At the same time, structural risks remain aligned with broader trends. Eighty-five percent say they are overly exposed to relying on too few AI vendors (vs. 76% globally), reinforcing that even at a more deliberate pace, the challenges of scaling AI remain.


Confidence in AI execution is the lowest globally

Japan stands out for having the weakest confidence in AI execution. Only 60% of Japanese CEOs say their AI roadmap accurately reflects reality (vs. 81% globally), while 14% say they are not confident (vs. just 3% globally) — the highest level of skepticism across regions.

This gap is significant. It signals that for many Japanese leaders, the challenge is not defining AI strategy, it’s trusting that it is actually working as intended.

AI informs decisions, but is less likely to drive action

AI is influencing decision-making in Japan, but less often translating into action.



This suggests a more cautious operating model. AI is being used, but outputs are more likely to be scrutinized, challenged, or deprioritized before action is taken.

Survey Methodology

The research was conducted online by The Harris Poll on behalf of Dataiku from February 2-March 2, 2026. The survey was conducted among 900 CEOs in the United States, United Kingdom, France, Germany, UAE, Japan, South Korea, and Singapore.

For reporting purposes, EMEA is the aggregate of the U.K., France, Germany, and UAE; APAC is the aggregate of Japan, South Korea, and Singapore. A total of 900 online interviews were conducted (U.S. = 200, U.K. = 100, France = 100, Germany = 100, UAE = 100, Japan = 100, South Korea = 100, Singapore = 100).

EMEA (n = 400) and APAC (n = 300) regions are reported and significance tested in aggregate. Respondents to this study work for large companies with an annual revenue of or more than \$500 million or regional equivalents and hold the title of CEO.



Dataiku: The Platform for AI Success

Dataiku is the Platform for AI Success, the enterprise orchestration layer for building, deploying, and governing AI. In a single environment, teams design and operate AI agents, analytics, and machine learning with the transparency, collaboration, and control enterprises require.

For more than a decade, Dataiku has helped organizations turn data, analytics, and AI into measurable business value. As the market has evolved from analytics to machine learning to generative AI and agents, Dataiku has enabled customers to stay ahead — operationalizing new technologies while maintaining stability, governance, and trust.

Sitting above data platforms, cloud infrastructure, and AI services, Dataiku connects the full enterprise AI stack, enabling organizations to run AI across multi-vendor environments with centralized governance and oversight, at scale.

Today, many of the world's leading companies rely on Dataiku to deploy AI applications and agents in some of the most complex and highly regulated environments — turning AI into an enduring source of performance, competitive advantage, and long-term business value.

