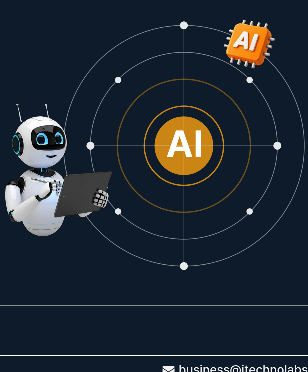
Who Actually Owns Your AI Program?

A 2026 survey of 900 CEOs says the board now holds them personally accountable for AI. It also says they are not in most of the decisions. This is what that gap means for the people who build the systems.



ISO 27001:2013 · ISO 9001:2015 · 500+ APPS DELIVERED · 300+ DEVELOPERS

EXECUTIVE SUMMARY

The accountability moved up. The decisions did not.

Most CEOs say they own their company's AI strategy. Far fewer sit in the decisions that shape it.

In a survey of 900 CEOs at companies with $500 million or more in revenue, **70% named themselves the primary driver of AI strategy** — well ahead of CIOs at 16%, chief data officers at 6% and IT departments at 5%.

Yet only 60% take part in more than half of AI-related decisions, and just 6% are involved in nearly all of them. Two-thirds have overruled a vendor or platform choice made further down the organisation in the past year.

That is the gap engineering leaders inherit. Strategy is claimed at the top. Architecture, model selection and vendor commitments are settled lower down. And the person who answers for the outcome arrives after the fact.

The pressure behind it is no longer abstract. **80% of CEOs say their role is at risk** if AI fails to deliver measurable business gains by the end of 2026, and 87% say they would stake their job on delivering those results.

When the person signing off on AI spend believes their tenure depends on the result, requests reach a delivery team differently. Timelines compress. The questions get sharper. What did the model decide, who approved it, and what happens if the vendor changes its pricing?

Those are engineering questions. This report reads the survey from that side of the table.

KEY FINDING

80% of CEOs say their **role will be at risk** if their company fails to deliver measurable business gains from AI by the end of 2026.

KEY FINDING

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

EXECUTIVE SUMMARY · CONTINUED

What the gap costs an engineering team

Read the survey as a delivery brief and five pressures come out of it. Each one lands on a CTO before it lands anywhere else.

01 **The proof deadline** p.04
Revenue is now the headline measure. Pilots that cannot reach a tracked number get cut.

02 **The control gap** p.06
Accountability sits at the top; decisions sit below it. Expect your choices to be revisited.

03 **The vendor trap** p.08
Over-investing now worries CEOs more than falling behind. Lock-in is the live risk.

04 **The shadow layer** p.10
Unapproved AI use is assumed, not suspected. A ban is not an answer; an alternative is.

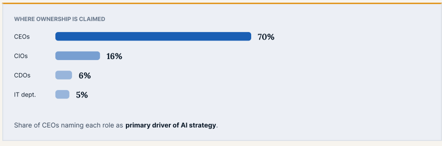
05 **The explainability test** p.12
Regulation is already moving timelines. Unexplainable output is now a brand risk.

KEY FINDING

62% say their board is **actively applying pressure** to deliver measurable AI-driven outcomes.

KEY FINDING

87% would **stake their job** on delivering measurable results from their AI program.



Share of CEOs naming each role as **primary driver of AI strategy**.

01

The proof deadline: AI now has to show up in a number the board already tracks.



Revenue growth has overtaken productivity as the most-cited measure of AI success. For delivery teams, that changes which projects survive review.

01 · THE PROOF DEADLINE

The expectation is no longer progress. It is proof.

Three-quarters of surveyed CEOs (73%) rank AI strategy and execution as a top or high priority, and nearly as many tie that priority directly to their own accountability. The survey's own phrasing is blunt: the margin has disappeared.

What changed is the measure. **Revenue growth is now the most-cited measure of AI success at 28%**, up from 16% a year earlier, and nearly level with increased productivity at 25%.

A pilot that saves hours is easier to build and harder to defend. If the saving never reaches a revenue line, a cost line or a cycle-time metric that finance already reports, it will not survive the next review.

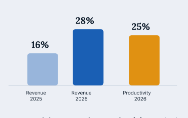
That is a scoping decision, and it belongs at the start of a project rather than at the end.

What this changes in practice

Write the success metric into the brief before the first sprint, and name the existing report it will show up in. If no such report exists, building the measurement is part of the work, not an afterthought.

Expect to be asked for attribution, not anecdote. A model that improved conversion needs a before-and-after a finance team will accept, which in most cases means instrumenting the baseline before the model ships.

KEY FINDING · MEASURE OF AI SUCCESS



Revenue growth has overtaken productivity as the leading measure of AI success, nearly doubling year on year.

KEY FINDING

35% say their job would be at **significant risk if the AI bubble were to burst** — a market bet layered on top of a performance bet.

ITECHNOLABS VIEW — Scope every AI project to a metric finance already reports — before the first sprint, not after the pilot.

02

The control gap: ownership is claimed broadly and exercised narrowly.



Seven in ten CEOs say they drive AI strategy. Six per cent are in nearly all the decisions. Two-thirds have overruled a choice made below them.

02 · THE CONTROL GAP

Accountable for the outcome, absent from the decision.

The most useful number in the survey is the distance between two figures. 70% of CEOs say they are the primary driver of AI strategy. Only 60% take part in more than half of AI-related decisions, and **just 6% are involved in nearly all of them**.

The consequence is already visible: 67% have challenged a vendor or platform choice made by their CIO or another team member in the past year.

From the CTO's chair, that is a pattern worth planning around. Decisions get made at the working level, then revisited at the top once the bill or the risk becomes legible. Without a record, every revisit restarts the argument from zero.

What this changes in practice

Keep a decision log. Not a document nobody opens — a short, dated record of what was chosen, which options were weighed, what it costs, and who signed. It converts a second-guess into a five-minute review.

Name an owner for each class of decision: model selection, vendor approval, data access, incident response. The survey shows ownership asserted in general and absent in particular. A one-page map closes most of that gap at almost no cost.

KEY FINDING · DECISION INVOLVEMENT

60% take part in **more than half** of AI-related decisions

6% are involved in **nearly all** of them

KEY FINDING

67% have **challenged a vendor or platform decision** made by their CIO or another team member in the past year.

KEY FINDING · GOVERNANCE FIRST



CEOs rank **governance ahead of people and orchestration** as what matters most for AI success.

ITECHNOLABS VIEW — A dated decision log turns a second-guess into a five-minute review.

03

The vendor trap: the fear is no longer moving too slowly. It is committing too early.



Two-thirds of CEOs now worry more about over-investing than under-investing, and three-quarters say they depend on too few AI vendors.

03 · THE VENDOR TRAP

Lock-in is the risk CEOs name first.

Two years ago the dominant worry was falling behind. In this survey the balance has flipped. **65% worry more about over-investing** amid heavy vendor competition and no clear market leader, against 35% who worry more about under-investing.

The exposure is specific. 76% say their organisation is overly exposed to operational or strategic risk because it relies on too few AI vendors.

Read plainly, the fear is about committing early to a stack that is hard to unwind — one provider's models, one provider's tooling, one provider's pricing power at renewal.

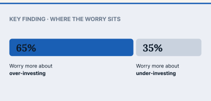
What this changes in practice

Put an abstraction layer between your application and any single model provider, so switching is a configuration change rather than a rewrite. Keep prompts, evaluation sets and fine-tuning data in your own repository, not in a vendor console.

Then prove it. Run your evaluation suite against a second provider before you need one. An exit path nobody has tested is a plan, not a capability.

This is our recommendation rather than a survey finding, but it answers the risk the survey names.

KEY FINDING · WHERE THE WORRY SITS



Heavy vendor competition and **no clear market leader** have reversed the dominant fear.

KEY FINDING

76% believe their organisation is **overly exposed to operational or strategic risk** because it relies on too few AI vendors.

KEY FINDING

89% say they could **keep core functions running** if their AI systems were disrupted — resilience is claimed, but rarely rehearsed.

ITECHNOLABS VIEW — An exit path nobody has tested is a plan, not a capability.

04

The shadow layer: unapproved AI use is no longer suspected. It is assumed.



Almost every CEO surveyed believes staff are using generative AI without approval, and two in five think most of the workforce is.

04 · THE SHADOW LAYER

A ban rarely holds. A better option usually does.

96% 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 not a suspicion at the margins; it is the assumed baseline.

The spread is partly encouraged from the top. 94% of CEOs call low-code and no-code tools critical to scaling AI across the workforce, which means building with AI is moving outside the teams that normally hold the controls.

The engineering consequence is twofold: company data leaving through channels nobody logs, and output quality that varies by whoever wrote the prompt.

What this changes in practice

Give shadow AI a sanctioned home. An approved internal assistant or retrieval tool, with SSO, access controls, retention rules and audit logging, competes with the unapproved tool on convenience instead of fighting it on policy.

Then watch the logs rather than the policy. Usage tells you which workflows people were trying to fix, which is a better roadmap input than a survey.

Set the retention and egress rules before you announce the tool, not after the first incident. Once a sanctioned assistant has traffic, changing what it keeps and where it sends data becomes a migration rather than a setting.

KEY FINDING · SHADOW AI

96% believe staff use generative AI **without approval**

42% estimate **more than half** the workforce is doing it

KEY FINDING

94% say **low-code and no-code tools are critical** to scaling AI across the workforce — so the spread is sanctioned at the top even where the tools are not.

CORROBORATING · EARLIER DATAIKU RESEARCH

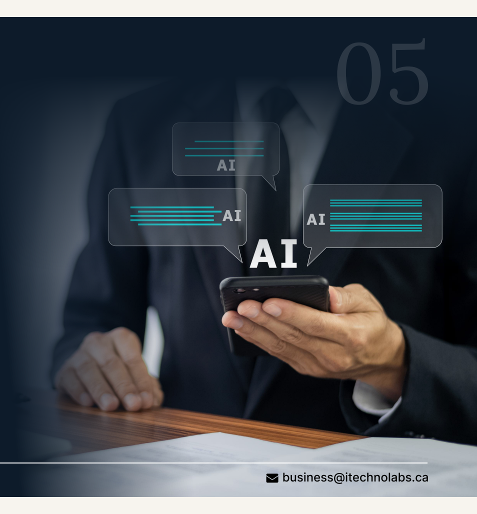
54% of CIOs report they have **already discovered staff using unsanctioned AI tools** for work tasks, and 91% of data leaders suspect it is happening.

From Dataiku's earlier research among data leaders and CIOs, cited in the 2026 CEO report — a different respondent base from the 900-CEO survey.

ITECHNOLABS VIEW — Usage logs from a sanctioned tool are a better roadmap input than any survey.

05

The explainability test: regulation is already moving delivery dates.



Half of CEOs have delayed AI initiatives over regulatory uncertainty, up sharply in a year. Explainability has become a brand risk, not a research topic.

05 · THE EXPLAINABILITY TEST

"Why did it decide that?" is now an engineering requirement.

51% of CEOs say they have delayed AI initiatives because of regulatory uncertainty, up sharply from 37% the year before. 80% agree that emerging regulations, including the EU AI Act, will slow AI adoption inside their organisation.

The reputational side is sharper still. 57% believe insufficient explainability could trigger a crisis that damages customer trust or the brand, and 79% express at least some concern that AI agents could expose the company to legal risk.

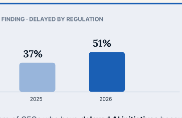
Fewer projects are being cancelled outright. More are being paused, re-scoped and re-architected — which is delivery work, arriving late.

What this changes in practice

Treat traceability as a functional requirement with a ticket, not a compliance task bolted on at the end. For every AI-assisted decision that touches a customer, you want the inputs, the model version, the output and the human approver retrievable months later.

Build that in from the first release. Retrofitting an audit trail onto a live system is the expensive version of the same work.

KEY FINDING · DELAYED BY REGULATION



Share of CEOs who have **delayed AI initiatives** because of regulatory uncertainty.

KEY FINDING

57% say weak explainability could **damage customer trust or the brand**.

KEY FINDING

79% have at least some concern that **AI agents create legal risk**.

ITECHNOLABS VIEW — Retrofitting an audit trail onto a live system is the expensive version of the same work.

THE OPERATING PICTURE

AI is operating. It is not deciding.

For all the talk of autonomy, the surveyed leaders are not handing over the keys. **51% keep humans in the loop for business-critical decisions**, and 34% run an explicit model where AI recommends and a person approves.

Interrogation is the norm, not the exception: 80% report at least occasionally asking for the reasoning behind an AI recommendation, and only 4% never ask it to justify one.

CEOs report AI influencing or informing up to an average of 40 business decisions a year. It is in the room, with a human checkpoint attached.

Ambition and certainty are moving in opposite directions. 83% are confident their organisation will deploy AI agents in full production during 2026, while the share describing themselves as extremely confident in deploying agents at scale fell from 41% to 31%.

What this changes in practice

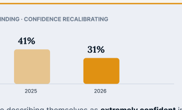
Design the checkpoint as a product surface, not a policy. The approver needs the recommendation, the reasoning and the confidence in one place, and the approval needs to be recorded with it. Build the review queue; do not assume email will do.

KEY FINDING · HUMAN OVERSIGHT

51% keep **humans in the loop** for business-critical decisions

34% run **AI recommends, human approves**

KEY FINDING · CONFIDENCE RECALIBRATING



Share describing themselves as **extremely confident** in deploying AI agents at scale.

KEY FINDING

~40 business decisions a year are **influenced or informed by AI**, on average, according to surveyed CEOs.

ITECHNOLABS VIEW — Design the checkpoint as a product surface. Build the review queue; do not assume email will do.

THE RESPONSE

Five things a CTO can do with this before the next board cycle

These are iTechnolabs recommendations, drawn from how the survey's risks translate into engineering work. They are not findings from the survey.

RECOMMENDATION · ITECHNOLABS

01 Write down who decides what.
Name an owner for model selection, vendor approval, data access and incident response. The survey shows ownership claimed broadly and exercised narrowly; a one-page decision map closes most of that gap.

02 Keep the human checkpoint, and log it.
Majority practice is recommend-then-approve for important decisions. Record the recommendation, the reasoning, the approver and the outcome so any decision can be reconstructed months later.

03 Design for vendor exit from day one.
Keep an abstraction layer between your application and any model provider, hold prompts and evaluation sets in your own repository, and run your evals against a second provider before you need one.

04 Give shadow AI a sanctioned home.
Offer an approved assistant or retrieval tool with SSO, access controls, retention rules and audit logging. Compete with the unapproved tool on convenience rather than fighting it with policy.

05 Tie every project to a number already tracked.
Revenue, cost or cycle time, named in the brief before the first sprint, and pointed at the report finance already reads. If a pilot cannot be expressed that way, rescope it before building.

THE THROUGH-LINE

Every item above turns a **board-level mandate** into something a team can build, explain and unwind. That is the work the survey describes without naming — and it is engineering work, not governance paperwork.

ITECHNOLABS VIEW — None of this is governance overhead. It is the delivery work that makes a mandate survivable.

HOW ITECHNOLABS SUPPORTS THIS WORK

We build the delivery layer the survey keeps pointing at.

iTechnolabs is an ISO 27001:2013 and ISO 9001:2015 certified engineering company with 300+ developers and 500+ apps delivered across Markham, Calgary, Ottawa and Sheridan, Wyoming.

For enterprise AI work, teams typically engage us for three things: LLM and RAG integration into platforms that already exist, AI agent development with the approval and audit surfaces built in, and dedicated engineering teams that extend an internal group without a hiring cycle.

All AI-assisted code on our own projects goes through the same review, security and performance checks as any other code before it ships.

300+ developers | **500+** apps delivered | **ISO** 27001 · 9001 certified

Ninety minutes, with your architects in the room. We map your current AI decisions, vendors, data paths and audit gaps against the five-point checklist, and you leave with the decision map and a written risk list whether or not you work with us.

Book Discovery Call

FREQUENTLY ASKED QUESTIONS

Questions this report gets asked

Who is responsible for AI strategy in most large companies?
Most CEOs say they are. 70% named themselves the primary driver, well ahead of CIOs at 16%, chief data officers at 6% and IT departments at 5%. In practice decisions are spread across data, IT and external vendors, so ownership on paper and control in practice often differ.

What is the biggest AI risk for CEOs in 2026?
Control rather than speed. CEOs ranked governance (39%) above people (34%) and orchestration (28%) as what matters most for AI success, and 76% said they depend on too few AI vendors.

Do CEOs trust AI to make decisions on its own?
Mostly not yet. 51% keep humans in the loop for business-critical decisions and 34% use an AI-recommends, human-approves model. 80% at least occasionally ask for the reasoning behind a recommendation, and only 4% never ask.

What is shadow AI and how common is it?
Shadow AI is employees using generative AI tools without company approval. 96% of surveyed CEOs believe it is happening and 42% estimate more than half their workforce does it. A sanctioned alternative with logging usually outperforms a ban.

How is regulation affecting enterprise AI projects?
51% of CEOs have delayed AI initiatives because of regulatory uncertainty, up from 37% a year earlier, and 80% expect rules such as the EU AI Act to slow adoption. Fewer projects are cancelled outright; more are paused and re-scoped.

How can a CTO reduce AI vendor lock-in?
Keep an abstraction layer between your application and any single model provider, store prompts and evaluation sets in your own repository, and validate a second provider in advance. This is an iTechnolabs recommendation, not a survey finding.

Are CEOs actually confident about AI agents?
Confidence is recalibrating rather than collapsing. 83% are confident their organisation will deploy AI agents in full production in 2026, but the share calling themselves extremely confident in deploying agents at scale fell from 41% in 2025 to 31% in 2026.

Does this survey apply to mid-market companies?
Only directionally. Respondents lead companies with $500 million or more in annual revenue. The pressures described here tend to appear in smaller organisations too, but the figures should not be quoted as representative of a 200 to 5,000 employee business.