

The **AI Governance Benchmark** for Leaders and Laggards

Why 75% of economic gains from AI are captured by just 20% of governed enterprises.

Table Of Contents

1. Executive Summary
2. The AI Governance Gap
3. The Dimensions of AI Governance
4. The AI Governance Maturity Curve
5. Leaders vs. Laggards: How Governance Differs
6. The Governance Gap in the Age of Agentic AI
7. Templatising the AI Governance Benchmark: A Scorecard
8. Regional and Sectoral Differences in AI Governance
9. Five Moves Toward Mature AI Governance
10. Conclusion
11. References

Executive Summary

From The Modern Data 101 Community

This report explores the shift passive policy documents to real-time agentic exception handling, the ultimate 2026 benchmark for measuring enterprise AI maturity.

As 2026 is nearly coming to an end, we would require AI governance maturity to be treated the way a lender treats a credit score, as a leading indicator of how a company will behave under stress. A credit score does not measure whether a borrower wants to repay a loan. It measures whether the borrower has a track record of doing so when it was inconvenient.

Every AI leaders-versus-laggards study of the past year, run by a different firm, on a different sample, asking different questions, keeps landing on the same dividing line. It is governance, and not spend, not model choice, and not even how early a company started.

A survey of 371 senior leaders found 60% of enterprises are already scaling AI across multiple departments, yet only 4% have governance mature enough to keep pace. That 56-point gap between deployment and control is the real subject of this benchmark.

PwC's 2026 AI Performance Study, based on 1,217 senior executives across 25 sectors, found that three-quarters of AI's economic gains are being captured by just 20% of companies. Those companies are not simply spending more. They are 1.7 times as likely to have a documented Responsible AI framework and 1.5 times as likely to have a cross-functional AI governance board.

We treat this as the central finding of 2026: governance has stopped being a compliance cost and started acting as a return multiplier. This benchmark works through what leaders do differently, what laggards are still missing, and why the gap is about to widen further as agentic AI scales faster than the controls built to supervise it.

AI governance works the same way, measuring whether anyone would actually notice, and act, the moment an AI system started behaving badly. Almost every company can produce the first kind of intention. Very few can produce the second kind of proof.

This brief maps that distinction across eight dimensions, builds a scorecard any organisation can use to locate itself on it, and argues that the gap is about to become a great deal more expensive to ignore.

The AI Governance Gap

Most enterprise AI governance fails for the same reason most New Year's resolutions fail. Building a model inventory, running red-team tests, or funding continuous monitoring shows up as next quarter's expense. The payoff, avoiding an incident that might not have happened anyway, is invisible and arrives on no one's scorecard. Against that incentive structure, governance loses to a feature launch in almost every budget conversation it enters, right up until the day an incident makes the cost of not governing suddenly visible too.

That incentive gap is large enough now to show up as a hard performance split rather than a soft cultural one. Sort companies by how seriously they treat this tradeoff and roughly 5% qualify as genuinely AI-driven, another 35% are still catching up, and 60% show almost no measurable return on their AI spend at all. The top group is not winning by a small margin either: it is outgrowing the rest by 1.7 times on revenue, 3.6 times on three-year shareholder return, and 1.6 times on EBIT margin.

Our key note

What strikes us most is how consistently the top slice of the market lands in the same narrow band no matter who is counting. One survey calls it 5% future-built firms. Another calls it 4% with governance mature enough to match their deployment. A third calls it the 20% of companies capturing most of AI's economic value.

These are not three separate phenomena. They are one population, viewed through three different lenses, and the consistency of the number is itself evidence that something structural, not incidental, separates that group from the rest of the market.

Roughly 5% of companies qualify as genuinely AI-driven, another 35% are still catching up

The Dimensions of AI Governance

Most corporate governance playbooks were written for systems that behave the same way twice. An AI model drifts as data changes, it responds differently to inputs no one tested it against, and an agentic system can take an action today that it would not have taken last week using the identical configuration.

Governing that kind of system with a one-time audit checklist is like inspecting a bridge once and assuming the traffic load never changes. The discipline AI actually needs looks closer to clinical trial monitoring: not a single approval gate, but an ongoing measurement practice that keeps checking the same thing long after launch.

Stripped of vendor branding, that ongoing practice breaks into four separable jobs, and NIST's AI Risk Management Framework gives them the clearest names available: Govern, which builds organisation-wide accountability and culture;

1. Map, which documents the context and risk of each specific AI system;
2. Measure, which tests systems against defined trustworthiness criteria; and
3. Manage, which turns those measurements into prioritisation, treatment, and incident response.
4. Govern sits above the other three, which is exactly why it is the one dimension almost every enterprise already has some version of. It is also, not coincidentally, the only one of the four that can be satisfied with a document instead of a system.

Where the effort actually goes

The uneven investment across these four functions is visible in every governance survey we reviewed.

That ordering is not random. It runs from cheapest to build down to most expensive, which is the same ordering an incentive-starved budget process would produce on its own.

Roughly half of organisations have usage policies and employee training in place, but only 34% maintain a model inventory, 21% run bias and fairness testing, and 17% conduct AI red teaming ~ [Source](#)

Why the expensive dimensions are the ones that count

A policy prevents nothing by itself. What actually catches a problem is something running continuously, which is why we weight Measure and Manage more heavily than Govern or Map when we assess an organization's real exposure. A parallel enterprise framework built specifically for agentic systems converges on the same three pressure points from a different angle: data protection, access control, and compliance, none of which is satisfied by writing something down once.

4

The AI Governance Maturity Curve

Ask an organisation to rate its own AI governance, and you will get a generous answer, because self-assessment measures intention. Ask whether a specific control actually fired the last time it was needed and you will get a much smaller number, because that measures proof. This is not a contradiction in the data so much as a description of two different instruments, and the size of the gap between what they report is itself the most useful number in this benchmark.

Four separate firms built four separate maturity curves in 2026, and every one of them puts the top tier in the single digits once the question moves from stated intent to operational proof.

The self-reported number

Asked to place themselves on a four-stage maturity model, 61% of organisations say they have reached the strategic or embedded stage, a number that reflects how an organisation sees its own intent and published policy more than what happens when that policy is tested.

72% of S&P 500 companies declared at least 1 material AI risk, a number that jumped from just 12% in 2023.

When it comes to AI governance, or even governance in general, it differs vastly from region to region, based on varying regulatory pressure. Even across the US, despite it being a major AI innovation centre, a substantial share of companies were developing or using AI without a publicly documented organisational AI policy.

The operationally verified number

Asked instead whether governance is mature enough to match the actual scale of AI running in production, only 4% of organisations qualify, and a separate benchmark finds only 34% describe their program as strategic and continuously improving, with no single function owning more than a quarter of governance responsibility.

Only 4% of organisations qualify, and a separate benchmark finds only 34% describe their program as strategic

“Organizations are increasingly focused on creating AI-ready data,” said Anurag Raj Senior Principal Advisor at Gartner.

“However, AI-ready data also requires AI-ready stakeholders who understand the value of trusted data, participate in data governance-related policy management activities, and overall maintain a culture of accountability and trust.”

Reading the curve correctly

We read this less as four competing maturity models and more as one curve with four checkpoints: Ad hoc, where a policy exists but nothing enforces it; Documented, where usage rules and training exist on paper;

Operational, where an inventory and periodic testing exist in practice; and Enforced, where monitoring and response run continuously without waiting for a human to notice first.

Most enterprise programs climb the first two rungs quickly, because they cost little and satisfy an auditor. The climb stalls at Operational, because that rung is the first one that requires a standing team rather than a one-time project, and standing teams compete for headcount against every other priority in the business.

Leaders vs. Laggards: How Governance Differs

The single behaviour that predicts almost everything else in this section is whether an organisation named one person to own AI risk before it named a single use case. Ownership forces every other decision to have somewhere to land.

Without it, governance becomes a shared responsibility in the worst sense of the phrase: everyone's job, which in practice means no one's. The market has been racing to fix exactly this: 76% of organisations now have a Chief AI Officer, up from just 26% a year earlier, and 60% of Fortune 100 companies are expected to name a dedicated head of AI governance in 2026, a role Sony, Bank of America, and UBS have already filled.

Capability follows ownership, and not the reverse

Once ownership exists, capability tends to scale quickly behind it, and the customer-care function shows the pattern cleanly: 67% of leaders have pushed foundational AI use cases to scale across the organisation, against just 16% of laggards, and 40% of leaders report meaningfully improved customer experience scores in the past year, against 12% of laggards. The capability gap is downstream of the ownership gap, not a separate problem requiring its own fix.

IBM's internal AI-driven cost analysis drove \$600 million in IT savings since 2022.

Trust in autonomy is earned, never assumed

Laggards are, primarily, not behind in the tooling aspects, but lack in not simply handing over control where they have no way to verify: 37% say they are not at all comfortable letting AI handle a customer interaction end to end, against just 4% of leaders.

That reluctance looks like caution from the outside, but it is better read as an accurate response to an unproven system, since 64% of CEOs at governed organisations now say they are comfortable making major strategic decisions on AI-generated input, a comfort level earned through evidence rather than granted on faith.

Governance and capability multiply, they do not add

The clearest evidence that these two levers reinforce each other rather than substitute for each other comes from a study of 755 mid-market leaders, where companies combining strong governance with real capability building reported an 81% measurable ROI on AI, against roughly 40% for either lever alone, and just 15% with neither. Spending on one without the other returns about half of what spending on both returns together, which is a strong argument against the common practice of funding them as competing budget lines.

6

The Governance Gap in the Age of Agentic AI

Every governance model discussed so far assumes a checkpoint exists somewhere before an AI system's output reaches the world: a human reads the draft, approves the recommendation, signs off on the decision. Agentic AI removes that checkpoint by design.

An agent that can call an API, update a record, or trigger a workflow has already acted by the time anyone could have reviewed it, which means governance has to shift from pre-approval to real-time exception handling, the operating model fraud detection systems use, not the one compliance sign-off uses. Very few enterprise governance programs have made that shift yet, which is why this is where the gap in this benchmark widens fastest rather than merely persisting.

The commitment is real, but the oversight to it costs hefty

Agentic AI is not a speculative bet for the companies pulling ahead. Agents already account for 17% of total AI value generated in 2025, projected to reach 29% by 2028, and future-built companies allocate 15% of their AI budget to agents, against 12% for scalers and almost none for laggards.

What Makes the Best Data Platform for AI in 2026 report reached from the platform side: agents don't create clean task boundaries and governed data access on their own, they inherit whatever architecture already exists. An enterprise with fragmented systems and undocumented processes will get agents that amplify that fragmentation faster than it can be cleaned up.

Set against that commitment, close to three-quarters of organisations plan to deploy agentic AI within two years, while only 21% currently report a mature governance model for it, the widest spread between intent and readiness anywhere in this benchmark.

We named this shortfall agent debt in prior research: the distance between how many autonomous decisions a system is authorised to make and how many of those decisions anyone can trace, own, or reverse. This benchmark's numbers are that same debt, confirmed a second time by an entirely different survey.

Shadow AI is the individual-level version of the same debt

Agent sprawl is what this debt looks like at the system level. Shadow AI, employees quietly routing work through public assistants and unsanctioned tools, is what it looks like one layer down, at the level of individual habit. Both failures share a root cause: an organisation cannot govern usage it cannot see, and most governance teams are working from an incomplete picture of which tools are already touching sensitive data. A policy document catches neither failure, because both happen in the gap between what a policy assumes and what employees and agents are actually doing.



Templatizing the AI Governance Benchmark: A Scorecard

Everything in Sections 3 and 4 is more useful as a tool than as an argument, so we built one. Cross the four governance dimensions against the four maturity stages, and the result is a scorecard an organisation can fill in honestly using its own operational evidence, not a vendor's proprietary rubric or its own self-flattering intentions. Score each cell by asking a blunt question: if this broke tomorrow, would anyone know before a customer did?

Dimension	Ad Hoc	Documented	Operational	Enforced
Govern	No named owner; AI risk is everyone's job and no one's	Policy exists; steering committee meets periodically	Named accountable role (CAIO or equivalent); board visibility	Owner has authority to halt deployment; board reports
Map	No inventory; nobody can list every AI system in use	Spreadsheet inventory, updated irregularly	Living model/agent registry tied to risk tier	Registry auto-updates from deployment pipeline
Measure	No testing beyond initial launch	Bias/accuracy checked once pre-launch	Scheduled re-testing; some red teaming	Continuous evaluation; red teaming built into CI/CD
Manage	Incidents handled ad hoc, case by case	Incident process documented, rarely rehearsed	Defined escalation path; incidents logged centrally	Real-time monitoring; automated containment for agents

Caption: A governance maturity scorecard mapping four dimensions (Govern, Map, Measure, Manage) against four stages, from Ad Hoc to Enforced. Source: Modern Data 101.

What an honest AI Governance score usually looks like

In our experience reviewing programs against this grid, the typical enterprise scores Documented on Govern and Map, and Ad hoc or Documented on Measure and Manage.

That pattern is not an accident. It is what an organisation looks like when it has satisfied an auditor but not yet built a standing capability. The fastest path off that plateau is not a clean sweep to Enforced across every row. It is moving Measure and Manage up a single column each, since those are the two dimensions that separate a scaling company from a genuinely future-built one in every survey cited in this benchmark.

87% of banks can improve their AI governance according to Deloitte's Trustworthy AI Governance Index ~ [Source](#)

A Banking Example

Banking sector is data-intensive, rule-bound, and sits under more supervisory scrutiny than almost any other industry. Deloitte's 2026 survey of 135 senior leaders at global and systemically important banks builds the same kind of index this benchmark does, across five pillars that map closely onto Govern, Map, Measure, and Manage, and finds the identical shape: 87% of global banks have clear scope to materially strengthen AI governance, with only 13% reaching an optimised level and 10% still effectively ad hoc. A fifth credible survey, built independently of the other four in this benchmark, lands on the same narrow band at the top.

The weakest pillar is the same one

Banking confirms exactly where the gap sits rather than introducing a new one. Organisational structure is the weakest pillar overall, with more than a quarter of banks ranked in the lowest category, and only 37% of banks report that senior leaders explicitly own or visibly exercise accountability for AI risk. That is a Govern-dimension failure wearing a banking-specific label, the same failure Section 3 argues is cheapest to fix and most commonly left unfixed.

Agentic AI lags here too

The pattern from Section 6 repeats sector by sector rather than breaking anywhere. 61% of banks have risk monitoring across the lifecycle for traditional AI and 59% for generative AI, but this drops to 44% for agentic AI, and 72% of banks have less than half of their AI use cases recorded in a central register, meaning most banks cannot currently produce a complete answer to what agentic systems they are even running.

The strongest evidence in this entire benchmark that governance pays

Every other return figure in this benchmark is a correlation between survey tiers. Deloitte's banking analysis is the one piece of econometric modelling in this benchmark, controlling for AI usage, bank size, headquarters location, and years in operation, and it still finds that a 10-point increase in governance index score is associated with a 10 percentage-point increase in revenue growth, roughly doubling the average bank's growth rate over the period studied. Deloitte is appropriately cautious about reading this as causal, and so are we: the sample skews toward smaller banks and the modelling cannot rule out that better-resourced banks simply find it easier to fund governance in the first place. But a controlled regression pointing in the same direction as four uncontrolled surveys is a stronger signal than any one of them alone, and it is the closest this benchmark comes to a dollar figure for what the gap in Section 2 actually costs.

Regional and Sectoral Differences in AI Governance

Regulation is a reasonable proxy for how expensive a mistake is allowed to be, and it tends to arrive well ahead of an organization's technical readiness to meet it, because lawmakers respond to liability exposure, not to engineering maturity. 78% of enterprises report themselves unprepared for their obligations under the EU AI Act. Readiness is not catching up to regulation. Regulation is simply outrunning it, which is the normal order of events rather than an exception to it.

Geography changes the incentive, and not the technology

Compare firms operating under different regulatory pressure and the behavioural gap shows up directly: only 19% of Irish firms say responsible AI is a business priority right across the organisation, against 28% in the US, with governance practices trailing by a similar margin. We do not read this as evidence that American firms are more sophisticated. We read it as evidence that litigation and disclosure risk arrived in the US market earlier, and governance investment tracks exposure, not technical capability.

Sector exposure tracks the cost of being wrong

The same logic explains why some industries govern harder than others without being more advanced. Financial services and insurance sit among the highest-pressure segments in this benchmark, and 42% of that group has already experienced a successful AI-enabled social engineering incident, a cost that is immediate and quantifiable. Technology and professional services, by contrast, are procuring agentic AI fastest while carrying the widest gap between deployment intent and oversight maturity, because for now the cost of a miss there is still mostly reputational rather than financial.

Five Moves Toward Mature AI Governance

None of the five moves below requires new technology. Every organisation scoring well on this benchmark already had access to the same models, the same vendors, and largely the same budget as the ones scoring poorly. What differed was sequencing and follow-through.

1. Name a single accountable owner before naming a single use case. Every other behaviour in this benchmark traces back to whether that ownership existed first.
2. Score honestly against all four NIST dimensions, not just Govern. Most organisations over-invest in policy because it is cheap and under-invest in Measure and Manage because they are not, which is precisely backwards given where this benchmark finds leaders actually separating from everyone else.
3. Move enforcement into the infrastructure layer rather than leaving it in a policy binder. Governance that lives inside the systems AI actually runs on catches what a document reviewed once a year never will.
4. Build agentic AI governance as its own workstream rather than an extension of existing policy, since the checkpoint that older governance models depend on no longer exists once a system can act on its own.
5. Size the investment to actual exposure, not to a generic industry average. A financial services firm and an early-stage technology company are not running the same risk clock, and funding both identically wastes budget in one direction while leaving real exposure open in the other.

Conclusion: From AI Policies to Operational Governance

Every credible survey we reviewed for this benchmark draws the line in roughly the same place: somewhere between 4% and 21% of enterprises have built AI governance that can actually prove itself under pressure, not merely describe itself on paper. That narrow band is not a data quirk. It is closer to a load-bearing wall in the modern enterprise, the kind of threshold that separates companies compounding an advantage from companies accumulating a liability they have not yet had to pay.

We expect that threshold to move in a jump rather than a gradual climb, the same way agentic AI adoption itself is moving in a jump, because boards rarely fund standing capability until a specific, visible failure makes the cost of not having it impossible to ignore. Every organisation reading this benchmark still has a choice about which side of that jump it lands on. The choice gets harder to make, and far more expensive, once it is made for you.

References

- [1] AI Leaders Outpace Laggards with Double the Revenue Growth and 40% More Cost Savings. Boston Consulting Group, press release, September 2025. <https://www.bcg.com/press/30september2025-ai-leaders-outpace-laggards-revenue-growth-cost-savings>
- [2] Three-Quarters of AI's Economic Gains Are Being Captured by Just 20% of Companies. PwC 2026 AI Performance Study, 1,217 executives across 25 sectors, April 2026. <https://www.pwchk.com/en/press-room/press-releases/pr-150426.html>
- [3] PwC's 2025 Responsible AI Survey: From Policy to Practice. PwC US, 2026. <https://www.pwc.com/us/en/tech-effect/ai-analytics/responsible-ai-survey.html>
- [4] PwC's 2026 Global CEO Survey. PwC, 2026. <https://www.pwc.com/gx/en/news-room/press-releases/2026/pwc-2026-global-ceo-survey.html>
- [5] Business and IT Leaders Report AI Agents Are Scaling Faster Than Their Guardrails. Deloitte Insights, survey of 3,235 leaders across 24 countries, April 2026. <https://www.deloitte.com/us/en/insights/topics/emerging-technologies/ai-agents-scaling-faster.html>
- [6] Building Trust: How Customer Care Leaders Pull Ahead with AI. McKinsey & Company, 2026. <https://www.mckinsey.com/capabilities/operations/our-insights/building-trust-how-customer-care-leaders-pull-ahead-with-ai>
- [7] State of AI Trust in 2026: Shifting to the Agentic Era. McKinsey & Company, 2026. <https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/tech-forward/state-of-ai-trust-in-2026-shifting-to-the-agentic-era>
- [8] IBM Study: CEOs Are Reshaping the C-Suite for the AI Era. IBM Newsroom, May 2026. <https://newsroom.ibm.com/2026-05-04-ibm-study-ceos-are-reshaping-c-suite-roles-for-the-ai-era>
- [9] Gartner: Global AI Regulations Fuel Billion-Dollar Market for AI Governance Platforms. Gartner, Inc., February 2026. <https://www.gartner.com/en/newsroom/press-releases/2026-02-17-gartner-global-ai-regulations-fuel-billion-dollar-market-for-ai-governance-platforms>
- [10] The State of AI Governance Report 2026. Credo AI, survey of 371 senior leaders, 2026. <https://www.credo.ai/downloadsopen/the-state-of-ai-governance>
- [11] AI Governance Stats for 2026: Adoption, Risk, and the Oversight Gap Defining the Year. Optro, May 2026. <https://optro.ai/blog/ai-governance-stats>

[12] AI Leaders vs Laggards: What the Data Shows. aibl, survey of 755 UK mid-market leaders, August 2026.

<https://aiblmedia.com/posts/ai-leaders-vs-laggards/>

[13] The Agentic Enterprise: How AI Agents Are Actually Being Used Inside Enterprises Today. Modern Data 101

Community Whitepaper, Vol. 03, September 2026. <https://www.moderndata101.com/modern-data-report-2026>

[14] AI Risk Management Framework. National Institute of Standards and Technology (NIST), 2023.

<https://www.nist.gov/itl/ai-risk-management-framework>


[15] Vision Compliance Releases 2026 EU AI Act Readiness Report. Vision Compliance, 2026.

<https://www.wisfarmer.com/press-release/story/49888/vision-compliance-releases-2026-eu-ai-act-readiness-report-finds-78-of-enterprises-unprepared-for-obligations/>

[16] EU AI Act Readiness, April 2026. Sprinkling Act, 2026. <https://sprinklingact.com/reports/eu-ai-act-readiness-april-2026>

[17] Responsible AI Practices Are Still in Their Infancy in Ireland. PwC Ireland, 2026. [https://www.pwc.ie/media-](https://www.pwc.ie/media-centre/press-releases/2026/responsible-ai-survey.html)

[centre/press-releases/2026/responsible-ai-survey.html](https://www.pwc.ie/media-centre/press-releases/2026/responsible-ai-survey.html)



540+ Respondents, 64 Countries, 29 Industries

[Access the report now](#)



Modern Data 101 Community



Subscribe 
moderndata101.com