

Insights supported by



DATA AND AI IMPACT REPORT: **THE NEW ECONOMICS OF TRUST**



SEPTEMBER 2026

Endorsements from partners and customers

“With trust gaps narrowing across industries like life sciences and financial services, the research highlights the practices that drive real impact in regulated environments. It’s an essential perspective for industries where oversight is paramount, and strengthening trust and governance is key as AI becomes more autonomous.”

Mona Chadha

Director of AI and Strategic Partnerships,
Amazon Web Services

“Governance should not be seen as a prohibitive or restrictive step; governance is the structure that makes AI adoption trusted, safe and scalable. Organizations that implement AI governance well will differentiate themselves from their competition.”

Karen Zieker

Vice President, Global Head of Biostatistics and
Statistical Programming, IQVIA

“The organizations creating the greatest value from AI aren’t necessarily those with the most data, but those that use it to make better business decisions. These insights reinforce that turning information into actionable insight is where AI can have the greatest impact, helping organizations respond more quickly and make more informed choices.”

Andy Fletcher

Vice President of Technology and Digital
Products, Liverpool Football Club

“As AI becomes embedded in core business processes, this report shows that trust will be the defining factor that separates isolated pilots from enterprisewide transformation. We are finding that organizations that establish strong governance, data integrity, and accountability today will be best positioned to realize measurable and sustainable AI value tomorrow — accelerating efficiency, empowering innovation, and laying the foundation for an agentic future in which AI not only generates ideas and insights but also orchestrates decisions, actions, and outcomes at scale.”

Christian Lau

Chief Technology Officer, Los Angeles Football
Club and BMO Stadium

“This research clearly shows how strong governance accelerates speed to value as organizations move from GenAI to agentic systems. Its insights on trust gaps, copilots and early quantum adoption are especially timely, offering a sharp view of what it takes to scale advanced AI responsibly.”

Jeff Markowitz

General Manager, Software and Digital Platforms,
Microsoft

“Organizations are tightening their AI practices with stronger governance, clearer validation and more disciplined production use. These findings on trust shift away from unstructured experimentation and offer timely insight into how AI programs are maturing to bring much-needed clarity to the real conditions shaping AI impact today.”

Dez Tsai

Global Senior Director of AI and Data,
TD SYNNEX

Introduction - A letter from Bryan Harris

"It's not a technology problem. It's a process problem."

"It's not an accuracy issue. It's an adoption issue."

"It's not the model. It's the human."

Sound familiar? I hear these arguments all the time.

But here's the reality: Generative AI has an accuracy problem. In fact, it is well documented that state-of-the-art agents can have error rates that exceed 25% on complex tasks. And those errors compound across multistep workflows. An agent operating at an 85% accuracy rate per step may only successfully complete a workflow a fraction of the time.

For critical, high-stakes business decisions, these outcomes simply aren't good enough.

So why are organizations racing to deploy AI?

Because when it works, it's remarkable. A model generates a compelling answer. An agent completes an intricate task. A process that once took months now takes hours.

These moments create excitement, opportunity and urgency. They show us what's possible when AI is embedded into a business. But they also create a temptation to focus on AI's potential – without fully addressing its limitations.

That's where the trust gap emerges.

The first challenge is knowing when we should use an agent and when we shouldn't. Many organizations apply agents to a task that simply requires standard automation with precision and determinism. As a result, they create unnecessary complexity and cost to achieve the same goal.

So when should we consider agents?

Today, agents deliver the greatest value when applied to goal-oriented tasks with complex workflows. In these scenarios, human expertise is still essential to evaluate results and overcome errors introduced by nondeterministic behavior. This oversight can be achieved through human-in-the-loop review or by embedding domain expertise into agentic workflows.

Ultimately, trust is earned through the accuracy and repeatability of results. Whether organizations use an agentic workflow or standard automation, once the result is accurate, consistent and repeatable, the trust gap disappears.

This is the goal. And it requires strong leadership to invest in a persistent technology strategy while ensuring people remain at the center of oversight and governance. Organizations that do this successfully will overcome the trust barrier, achieve groundbreaking outcomes, and gain a competitive advantage in the market with AI.

Only then can we confidently say, "It's not a technology problem."



Bryan Harris

Executive Vice President and
Chief Technology Officer, SAS

Contents

Introduction - A letter from Bryan Harris	
Executive summary.....	1
Methodology	3
Autonomy raises the stakes for trust	4
Agentic AI ambitions outpace trust	5
Understanding the value of trust and trustworthiness	6
Five elements of trustworthy AI	7
The override reflex	9
Data maturity is still an AI bottleneck	10
AI maturity raises sovereignty concerns	12
How to succeed: Lessons from leaders	13
Operating as a trustworthy AI leader	14
Five behaviors that set AI leaders apart	15
How does AI trust affect ROI?	16
Three mistakes stalling AI value	17
How industry shapes leadership	18
Lessons from leaders in banking.....	19
Lessons from leaders in insurance	23
Lessons from leaders in life sciences	26
Lessons from leaders in the public sector.....	29
Small and midsize organizations	32
Key takeaways for all	34
From trust to advantage: The road ahead	35
Glossary of terms	36

Executive summary

The more AI acts on its own, the more trust it has to earn. Right now, trust isn't keeping up.

A CFO handing a procurement decision to artificial intelligence (AI) was science fiction a few years ago. Today, it's a Tuesday. Generative AI (GenAI) no longer just informs decisions. Agentic AI executes them, and leaders are standing at a turning point.

The outcomes are already real: assessing credit risk, pricing insurance, processing government benefits, matching patients to clinical trials, writing code, orchestrating workflows and making calls once reserved for human judgment. As autonomy climbs, so does exposure to regulatory scrutiny, operational risk and reputational damage.

That makes trust an operational concern, not a philosophical one. It doesn't come from the technology. It comes from how organizations design, govern and validate the systems they deploy. And the payoff is concrete: Organizations that build trustworthy practices into their AI are **15 times more likely to report strong ROI**.

The implication is simple. Trustworthiness now decides whether AI investments scale into production or stall out in overrides, hesitation and quiet underuse. Leaders are pulling ahead. Everyone else is falling behind, often without realizing it.

This second edition of the *Data and AI impact report*, drawn from **2,699 decision-makers across 28 countries**, captures a market at an inflection point. Trust is under pressure. Governance is racing to catch up. And a distinct group of trustworthy AI leaders has emerged, converting AI into measurable outcomes at a rate the rest of the market can't match.

This report maps what trustworthy AI leadership actually looks like: what the data shows, what leaders do differently and where organizations need to focus now to build what's next.

What is trustworthiness?

Trustworthiness is measured across five dimensions: Data Quality & Governance, Model Governance & Oversight, Explainability & Fairness, Responsible AI Policy, and Audit & Accountability.

Together, these dimensions form the Trustworthiness Index.

Organizations scoring 80 or higher are classified as trustworthiness leaders. While they represent a relatively small segment of the market, they are significantly more likely to achieve stronger AI outcomes and ROI.

Learn more: See "The five elements of trustworthy AI (page 7)" and "Lessons from leaders (page 13)."



KEY FINDINGS

Taken together, these findings point to a defining moment in AI adoption. Organizations are moving from systems that generate answers to systems that take action. As AI gains more autonomy, trustworthiness becomes the determining factor in whether organizations realize value, manage risk and scale AI successfully.

Trust drops the moment AI starts to act. (Page 4)

Agentic AI is trusted by 66% of respondents, trailing generative AI by 10 points as AI moves from generating insights to taking action.

Deployment is outpacing trust. (Page 4)

Roughly 89% of respondents say their AI agents now play a role in decision-making, from informing human-approved decisions to acting autonomously, yet trust in agentic AI remains 10 points lower than GenAI.

Explainability is the biggest barrier. (Page 4)

“Lack of explanation” is the top reason users override AI (34.5%), and Explainability & Fairness improved the least of the five trustworthiness dimensions.

Data maturity is holding AI back. (Page 10)

Only 17.5% of organizations run data infrastructure at the highest maturity level. Most can't yet support the lineage and validation that agentic AI requires.

Leadership is behavioral, not structural. (Page 15)

Leaders aren't defined by org charts. They're defined by five behaviors, and they're 44 times more likely to run governance and explainability reviews.

Trust drives outcomes. (Page 16)

Organizations with trustworthy AI practices are 15 times more likely to report strong or high ROI (62% vs. 4%).



AI can be reliable, but not necessarily trustworthy. If it's not transparent and you don't understand what it's doing or how it's doing it, you don't necessarily understand what's in the black box.

IT Director, Local Government (US)

Methodology

We set out to measure trust, trustworthiness and the business impact of AI. What we found was something larger: Organizations are entering a new phase of AI adoption, and that changes the trust dynamic. As systems evolve from generating outputs to taking action, the requirements for trust are changing as well. Many of the results in this report – from override rates and explainability challenges to governance investments and ROI performance – reflect organizations adapting to this new reality.

Sample composition



Focus industries



The trust gap equation




Figure 1: How the trust gap is calculated, methodology, sample size, and who we spoke to

Trust and trustworthiness are not the same thing:

- **Trust** is how much confidence people have in AI.
- **Trustworthiness (TW)** reflects the practices, controls and governance mechanisms that make AI worthy of that trust.

This report measures both. While trust can rise or fall based on perception, trustworthiness reflects the underlying capabilities organizations build to make AI reliable, transparent, accountable and safe.



Autonomy raises the stakes for trust

This transition from assistance to autonomy is the central story of this year's research. The findings that follow – from override behavior to governance investment – reflect how organizations are adapting to AI systems that are increasingly capable of acting on their own.

Generative AI is trusted by 76% of respondents globally. Agentic AI, which takes actions on its own, is trusted by 66% of respondents. That 10-point difference is not random.

The further AI moves from making suggestions to taking actions, the more it can do for a company, but the more it must do to be credible. Each step adds capability and raises the bar for trust.




Figure 2: The relationship between agentic AI, autonomy and human override

Agentic AI ambitions outpace trust

Unquestionably, autonomous AI is here. The question now is whether trustworthiness is advancing quickly enough to support it.

Organizations want agentic AI for automation (56%), competitive advantage (52%), innovation (51%), scaling digital labor (48%) and faster decision-making (48%). And the deployments are not theoretical.

It reflects a reasonable expectation that systems capable enough to deploy are capable enough to act. The problem is that trustworthiness, which ensures the technology can be trusted, has not kept pace with deployment ambition. The ambition for agentic is bigger than last year's, but the data shows the confidence underneath it has not kept pace.



Control is also proportional [to risk]. Human involvement scales with the level of risk, but not with skepticism about the technology. If we realize that the degree of human intervention is too high – that means that the system needs to be worked on.


Chief Information Security Officer, Global Bank (US)



Understanding the value of trust and trustworthiness

Perceived trust

Subjective user/organization trust in AI



Trustworthiness

Objective governance practices that prove AI reliability, fairness and oversight




Figure 3: Global score for perceived trust and trustworthiness

As any technology matures, the gap between what it can do and what people believe it can do tends to close. AI adoption is confirming that pattern, and this year the data is explicit about who is doing the work to close this gap.

The Trustworthiness Index, our measure of the governance, data and oversight practices that prove AI reliability, jumped from 49.7 to 61.3 out of 100.

Perceived trust slipped from 73 to 70.7, showing that while trustworthy capabilities are rising, people's overall trust in AI is dropping.

The gap between the two shrank by 13.9 points, from 23.3 to 9.4. Capability is finally catching up to meet expectations, and systems are earning the trust they previously captured on faith alone. That progress didn't happen on its own.

Organizations have been putting real work into the governance, controls and transparency that make AI worth trusting. And users have been giving it a fair chance by trying it on real tasks, pushing back when something feels off and giving it another shot when it improves.

These nuances aren't captured by the overall trustworthiness score alone. Progress varies across the individual elements of trustworthiness, so a closer look is needed to see how organizations are truly advancing.

Five elements of trustworthy AI

In this research, trustworthiness is measured by five key components: Data Quality & Governance, Model Governance & Oversight, Explainability & Fairness, Responsible AI Policy, and Audit & Accountability, each measured on a 0-100 scale.

These five elements become increasingly important as AI moves into the agentic age. Together, they provide the foundation organizations need to deploy increasingly autonomous AI with confidence.



Figure 4 - Five dimensions of trustworthiness



Responsible AI Policy and Audit & Accountability are the highest scoring and biggest movers year-over-year. There was a lot of policy-and-paperwork investment: Writing the rules, documenting the controls and assigning accountability are steps that require organizational commitment rather than deep technical infrastructure. Most organizations have prioritized them accordingly.

Data Quality & Governance and Model Governance & Oversight sit lower because they require more sustained investment in pipelines, monitoring and engineering practice, but still show a significant change.

The outlier is Explainability & Fairness, which shifted the least year-over-year. That looks insignificant until you set it against the override data: “Lack of explanation” is the single most common reason users override AI, named by 34.5% of respondents. The dimension that has improved the least is the one users react to the most.

The other four dimensions are largely policy and paperwork: Write the rules, document the controls, schedule the audits and manage the data. Explainability is a different kind of problem. It requires the model, the data lineage and the user interface to surface reasoning a human can actually follow, which means engineering work across the whole stack rather than a governance memo. That’s why it moves slowly. But the override data says organizations can’t keep leaving it for later.

The five elements provide a blueprint for trustworthy AI. The challenge is that many organizations are deploying increasingly autonomous systems before these foundations are fully in place. The consequences are visible throughout this research – in override rates, explainability challenges, data maturity gaps and uneven business outcomes.

Perhaps the most compelling effect appears in a clear behavioral measure of trust: how often people override AI decisions.



When we think of trust, it truly goes back to the source of the data that we are trying to navigate through. Because for all the AI systems that we are trying to build, they are only as good as the data that we are providing to these systems.

Chief Information Security Officer Global Bank (US)

The override reflex

The five elements of trustworthy AI provide a framework for understanding what organizations need to build. But how can we tell whether those investments are working? One of the clearest signals is human behavior. When people trust an AI system, they let it operate. When they don't, they intervene. Override rates offer a real-world measure of where trustworthiness is succeeding – and where gaps remain.

How often are respondents overriding AI decisions? Seventy-nine percent of respondents say they override AI recommendations in more than 10% of cases, and only 0.6% never override at all.

When do users override AI?

- **More than one-third (34.5%) say when the output is sound but its reasoning is hidden.** The most common reasons for override are a lack of explanation and missing context. The output may be correct, yet people step in because they cannot see how the system reached it. This is a trust and transparency gap, and the remedy lies in explainability and the data lineage that supports it.
- **More than one in four (25.6%) say it's when the output doesn't reflect context.** Respondents override because the system missed a situational detail, an edge case or information it wasn't built to account for. This is a contextual blind spot, and the remedy is richer inputs and tighter integration with the systems and data that carry that context.




Figure 5: Reasons for override by trustworthiness tier, comparing trustworthiness leaders with trustworthiness laggards

The practical implication points back to the data foundation. The largest and most addressable share of overrides comes from outputs people cannot interpret: A lack of explanation and missing context together account for the majority of reasons cited. An AI system can only explain a decision it can trace, which makes explainability a property of the data and governance beneath the model. The sections that follow examine where that foundation is strong, where it is fragile and what leaders do differently.



Data maturity is still an AI bottleneck

Most organizations deploying AI today are doing so using data infrastructure that cannot fully support the explainability, lineage or validation that agentic AI requires. This lag in maturity is not abstract, and it shows up in the override rates (see Page 8). Organizations are making steady progress, but until data maturity catches up, trustworthy autonomy will remain out of reach.

Data infrastructure maturity is measured along a five-stage ladder (Ad Hoc, Siloed, Standardized, Managed, Optimized), and the year-over-year shift is the clearest progress signal in the study.




Figure 6: Comparing the year-over-year change of organizations in terms of AI maturity

Standardized, the most common stage in 2025, dropped 16.1 percentage points as organizations moved up to higher maturity levels. In 2026, the Managed level rose 8.8 percentage points; Optimized rose 7.3 percentage points. More than half of respondents (52%) now see themselves at the Managed or Optimized level, but only 17.5% are fully Optimized.

Most organizations deploying AI today are doing so using data infrastructure that cannot fully support the explainability, lineage or validation that agentic AI requires.

The maturity payoff: Ad Hoc vs. Optimized data infrastructure

	Ad Hoc	Optimized	Multiple
Mandate data quality tools	5%	34%	6×
Mandate explainability controls	7%	40%	6×
Automated AI validation & monitoring	9%	45%	5×
Expect strong/high AI ROI	14%	52%	4×

Organizations with optimized data infrastructure are 4-6x more likely to mandate AI quality and explainability controls - and nearly 4x more likely to expect strong ROI.

Figure 7: Comparison table for organizations with Ad Hoc vs. Optimized data infrastructure

Figure 8 illustrates the investment trends among two groups of users. Compared with Ad Hoc peers, Optimized organizations are six times more likely to mandate data quality tools across every AI project, six times more likely to mandate explainability controls, five times more likely to run automated, continuous AI validation, and four times more likely to expect strong or high ROI.

Leaders are not marginally ahead. They are multiples ahead. The leaders are pulling away because each control they put in place unlocks the next level of data maturity. Validation catches problems before they ship, which builds confidence to deploy. This confidence justifies more investment in explainability, leading to lower override rates – and stronger ROI.



AI maturity raises sovereignty concerns

Data maturity levels also change how much sovereignty concerns weigh on an organization. Asked how much data sovereignty and residency requirements affect their ability to deploy AI, low trustworthiness organizations report little impact: 19.5% say “significantly,” while 4.1% say “critically.”

High trustworthiness counterparts answer differently: 48.3% respond “significantly” and 29.7% respond “critically.” This is roughly 2.5 times the impact at the significant level and seven times at the critical level.

The finding may appear backward at first: Surely, mature organizations should be the least affected. In reality, the opposite is true. With more maturity comes more awareness of potential risks.

More mature companies see exposures their low trustworthiness peers have not yet noticed: which jurisdictions their training data crosses, which extraterritorial regimes reach into their cloud contracts, which deployments would have to be paused if a residency rule was tightened. Among the most AI-mature enterprises, roughly 78% globally flag sovereignty as a material constraint, and the figure barely fluctuates across regions (North America 78.6%, Europe 74.4%, LATAM 80.4%, APAC 79.2%, META 81.6%, global average 78%). These figures reflect data residency and jurisdictional control over data (where it lives and which regimes reach it), not national AI-capability sovereignty, and the near-flat regional spread confirms the pattern tracks maturity, not geography.

Sovereignty awareness is mainly a maturity indicator, though it does come with local nuances. Organizations investing in trustworthiness tend to see the next problem first.

Sovereignty concern rises with trustworthiness.

How much do data sovereignty and residency requirements affect your organization's ability to deploy AI?




Figure 8: Sovereignty concern for High Trustworthiness and Low Trustworthiness organizations

How to succeed: Lessons from leaders

The behaviors that separate leaders from laggards are increasingly the same behaviors required to deploy agentic AI and traditional AI safely and at scale.



Figure 9 Characteristics separating high trustworthiness organizations from low trustworthiness organizations



OPERATING AS A TRUSTWORTHY AI LEADER

We define leaders as organizations scoring 80 or above on the Trustworthiness Index. They treat trustworthiness as foundational, and the results show a massive divide between the behaviors of leaders and laggards.

The behavioral divide (% of orgs)

● Trustworthiness laggards ● Trustworthiness leaders




Figure 10: Behavioral divide between leaders and laggards

1.85× Realized gain across 13 outcomes

High TW 27.6% avg vs. low TW 14.9%

1.86× Expected ROI per \$1 of AI spend

High TW \$2.35 vs. low TW \$1.26

15× Reach strong/high ROI (\$2+ per \$1)

High TW 62% vs. 4% of low TW

Figure 11: Key benefits trustworthiness leaders realize

FIVE BEHAVIORS THAT SET AI LEADERS APART

Our research shows five practices that explain why trustworthy AI leaders convert AI spend into returns while others don't.

1. OPTIMIZE YOUR DATA INFRASTRUCTURE

Leaders are 11 times more likely to run 'optimized' data infrastructure

Data infrastructure is built so that it can be audited and traced. When you can show the lineage of a decision, you can defend it. When you can't, you're managing risk through human override, which can be expensive and hard to scale.

2. VALIDATE CONTINUOUSLY

Leaders are 10 times more likely to use finance-validated ROI models for AI projects

AI leaders measure AI value at the workflow level using finance-validated models (i.e., the CFO's team validates the expected return, the target metrics are tracked against forecast and a named owner is accountable for delivering them), not through pilots or sentiment. This forces clarity on what actually works and what doesn't, so failed experiments don't hide as invisible costs. The same rigor should apply to risk and harm, not only ROI.

3. TRAIN ALL EMPLOYEES ON TRUSTWORTHY AI

Eighty-five percent of leaders train across the organization on their responsible AI policy

Rather than merely writing a policy, leaders actively train the entire organization on their responsible AI policy. When everyone knows the rules, governance is owned across the business rather than left to a specialist team, so problems are caught early instead of discovered by users in production. Accountability becomes structural.

4. GOVERN BY DEFAULT

Leaders are 44 times more likely to conduct governance and explainability reviews

Leaders create independent governance functions with real authority. This paradoxically accelerates deployment because teams know what they can and cannot do.

5. INVEST AHEAD OF THE CURVE

Eighty-five percent of leaders are increasing their spend on trustworthy AI by more than 10%

Leaders invest in trustworthiness itself, treating it as strategic rather than discretionary. The difference isn't simply that leaders spend more: 85% of leaders increase their trustworthy AI investment by more than 10%.

Each practice enables the next. Together, they compound.



How does AI trust affect ROI?

Organizations that invest in trustworthy AI see a markedly different return on their AI investments. Those with high trustworthiness scores are 15 times more likely to report strong or high ROI than organizations with low scores.

The difference is significant. Nearly six in 10 organizations scoring 80 or higher on the Trustworthiness Index achieve at least a two times greater return on their AI investments. Among organizations with low trustworthiness scores, fewer than one in 20 report the same outcome.

This relationship becomes even more important in the agentic era. Organizations cannot scale autonomous systems if employees, regulators and customers do not trust them enough to use them.

The findings suggest that trustworthiness is more than a risk-management consideration. It is a business advantage. Organizations that establish clear governance, maintain strong data foundations and make AI decisions understandable are better positioned to move from experimentation to operational use. AI becomes part of how work gets done rather than another pilot program competing for attention and resources.

This dynamic is particularly important in regulated industries, where questions of accountability, explainability and oversight shape how quickly new technologies can be adopted. Organizations that address those requirements upfront are often able to scale AI with greater confidence and fewer barriers.

By contrast, organizations that treat trustworthiness as an afterthought frequently struggle to realize value from otherwise capable technology. Models may perform well, but adoption lags. Use cases remain limited. Deployments move more slowly. The result is a widening gap between AI investment and AI impact.

The data points to a clear conclusion: Trustworthiness does not merely protect AI value. It helps create it.

With this in mind, 86% of organizations report they are increasing trustworthy AI investment in the year ahead, with near-universal acknowledgment that the agentic shift requires a different, stronger governance posture than [GenAI](#) did.

15×

**More likely: Strong/
high ROI**

*15 to 1. About 6 in 10 high-trust
orgs see \$2+ back for every \$1 of
AI - fewer than 1 in 20 low-trust
orgs do (62% vs. 4%).*

Figure 12: The ROI multiplier by trustworthy AI organizations

Three mistakes stalling AI value

AI leaders are defined as much by what they avoid as by what they adopt. Across our interviews and survey data, three pitfalls consistently separate leaders from laggards.

DEPLOYING AI PEOPLE DON'T TRUST

Organizations that cannot explain AI decisions end up paying for them twice. As noted earlier, 79% of organizations override AI recommendations in more than 10% of cases, and lack of explanation is the leading reason. Every override adds review, rework and delay. Organizations would be well advised to make transparency and explainability of their AI systems and models a key requirement.

CONFUSING PILOTS WITH PROGRESS

Roughly half of AI budgets go to sales and marketing pilots, yet many of the highest-ROI use cases remain underfunded. Leaders invest in workflows that change how work gets done, not demonstrations that showcase what AI can do.

TREATING GOVERNANCE AS A HURDLE INSTEAD OF AN ENABLER

Organizations that bolt governance on after development create friction, delays and resistance. Leaders treat governance as a system accelerator vs. system speed bump, making it easier to scale AI with confidence.

Across industries, the same pattern appears repeatedly. The organizations pulling ahead are not using fundamentally different technology. They are avoiding a handful of costly mistakes that prevent AI from delivering value.





How industry shapes leadership

In the next section, we'll examine four priority industries explored in this research: banking, insurance, life sciences and the public sector. Each will show how the behaviors of trustworthy AI leaders and laggards practically appear in real-world industry use cases.

When read together, these help to map how the universal playbook can adapt to each individual industry's regulatory, operational and reputational realities.

Finally, AI value is theoretical until someone applies it to a real-world problem. Seeing how leaders in each of these key industries around the globe are approaching AI helps bring the universal principles of trustworthy AI into sharp focus, and allows us to learn from a variety of perspectives regardless of what industry you sit in today.

LESSONS FROM LEADERS IN BANKING



Banking

TRUSTWORTHINESS

63.5 +12.8

Y1: 50.7

Key concern & priority

Model Governance & Oversight, at a 10-point gap, is the only remaining soft spot. Closing this gap is banking's next priority, and it grows more urgent as agentic autonomy makes approval and oversight harder to retrofit.

The deeper tension is operating tempo: Leaders and laggards target the same mission-critical use cases, but leaders let AI act autonomously within defined boundaries (41% deploy AI in investigations vs. 7% of low-trust banks).

6.6 times increase in ROI outcomes for trustworthy AI leaders

High trustworthiness banks have cascaded governance through the organization in a way laggards haven't started, and the scale of the difference is hard to overstate. Eighty-five percent of leader banks have rolled out a responsible AI policy that reaches every employee, vs. 29% of laggard banks.

Automated, continuous AI validation exists in more than two-thirds of the leader group, compared to effectively zero in the laggard group. These differences are fundamental: They determine whether the operating model exists at all.

The outcome follows directly: 66% of leader banks expect strong or high AI ROI. This represents a 6.6 times difference between leaders and laggards in a sector that already started from a higher base than most.



The biggest challenge with the agent is: If the agent makes a decision, how do we go back and trace it and explain why the agent took certain actions? That's one of the reasons why we haven't really embraced autonomous AI completely yet.

Chief Information Security Officer, Community Bank (US)



Split on automation investigations

Both leaders and laggards agree on where AI matters most: credit, risk, eligibility and financial crime detection. The divide is not in priorities, but in autonomy.

Leader banks are far more likely than laggards to view automated investigations and case management as mission-critical (41% vs. 7%), the widest gap of any use case in our study. These systems do more than surface insights for human review; they take action within defined guardrails.

The implication is clear: The advantage does not come from choosing different use cases. It comes from building enough trust, governance and accountability to allow AI to operate at speed and scale.

Concerns shift with maturity

Sixty-four percent of leader banks flag data privacy and security as a primary risk, compared to 46% of laggard banks.

Laggard banks, by contrast, are more often blocked by reliability worries (39%) and unclear AI ownership (22%), the problems you encounter before governance is built. Agentic concern is rising sharply on both sides and the divide is telling: 38% of leader banks say that managing autonomy and human oversight is a concern, vs. just 12% of laggard banks. The leaders are running the deployments where that question is live, but the laggards have not yet reached it.

Banking: The five dimensions of trustworthiness




Figure 13: Banking – Five dimensions of trustworthiness

Banking scores above the global average on every dimension of trustworthy AI, with year-over-year gains running from roughly eight to 14 points across the board. Trust gaps have narrowed on all five dimensions; however, capability still meaningfully trails perception in Model Governance & Oversight, at a 10-point gap.

We see striking uniformity: No dimension is neglected, no gap is left to widen while another closed. Banking is extending two decades of model risk management (MRM) into AI, faster than any sector that started from zero.



The one dimension still trailing, Model Governance & Oversight, is the layer that operationalizes MRM: the committees, reporting lines, documentation and approval workflows that decide how banks “manage the managing.” Closing that 10-point gap is banking’s next priority, and it grows more urgent as agentic autonomy makes approval and oversight harder to retrofit. This MRM-led pattern is most pronounced in markets like the US, where AI-specific rules remain nascent and supervisors fold AI into existing model risk expectations. In jurisdictions with dedicated AI regimes, such as the EU, AI governance is defined more explicitly beyond MRM alone.



We have an AI governance committee or council ... it has business stakeholders, technology stakeholders, security stakeholders, plus risk and compliance stakeholders. We have some legal folks as well to just make sure that they weigh in on if there’s anything we need to be aware of from a reputational risk perspective.”

Managing Director, Global Bank (US)

LESSONS FROM LEADERS IN INSURANCE



Insurance

TRUSTWORTHINESS

63.2 +17.6

Y1: 45.6

Key concern & priority

The priority is keeping capability and perceived trust locked together as deployments scale, because insurance is the only sector where the Responsible AI capability score has crossed zero, ahead of perceived trust.

The immediate concern for leaders is managing autonomy and human oversight at scale: 34% of high-trust insurers flag this, vs 21% of low-trust peers, because leaders are the ones running those deployments today.

Largest disparity between leaders and laggards

Insurance shows the cleanest trust-driven split of any focus industry, with a wider leader/laggard distance than banking on almost every measure. Seventy-eight percent of leaders have cascaded a responsible AI policy to every employee, vs. 3% of the laggards. Automated, continuous validation and independent governance functions exist in a majority of the leader cohort; in laggards, they are effectively absent.

The result lands at the framework's limit: 59% of leaders expect strong or high AI ROI, vs. 0% of laggards, making this the only focus industry where the laggard cohort produces no strong/high ROI signal at all. In a sector built on pricing risk, the organizations that priced governance risk early are the only ones converting AI spend into AI return.



100% explainability is key. The underwriting ops team now gets a monthly report. Instead of the ops team having to review every single document, now they only focus on about 10% of the effort to complete the task manually rather than the whole thing.

Head of Data & Governance, Global Insurance (UK)



Leaders concentrate on outcomes that create value

For leaders in insurance, AI sits at the heart of the decision intelligence stack: risk assessment, underwriting, regulatory compliance, pricing, capital management, claims and fraud. Organizationwide analytics and customer support are all in scope.

Asked what they most want AI to achieve, laggards name cost savings above all else, at 48.5%, significantly more than the 23.3% of leaders who prioritize it. Leaders instead concentrate on outcomes that create value rather than strip it out: operational efficiency (37.2% vs. 27.3%), driving new revenue growth (30.2% vs. 15.2%) and increasing profits (30.2% vs. 21.2%). The pattern echoes one of this research's most durable findings: Deploying AI chiefly to cut costs is among the least effective ways to use it. The least mature insurers still reach for cost reduction first; the leaders have moved on to growth, efficiency and profit.

Concerns mature with deployment

Leaders have moved past the basics and are worried about scale. Regulatory compliance, accountability and reputational risk are the concerns that only become real once AI is doing work that the regulator and the brand can be held to. Fifty-two percent of the laggards flag data privacy and security as a top concern, vs. 42% of leaders – showing that the laggards are still trying to remedy data infrastructure issues.

On agentic, 34% of leaders say managing autonomy and human oversight are a concern, compared to 21% of laggards. Trustworthy leaders in insurance are worried about autonomy because they are actually running those deployments today.

In a sector built on pricing risk, the organizations that priced governance risk early are the only ones converting AI spend into AI return.

Insurance: The five dimensions of trustworthiness




Figure 14: Insurance – Five dimensions of trustworthiness

Insurance’s five dimensions moved together this year, with year-over-year gains ranging from 11.3 to 16.2 points. Model Governance & Oversight is the biggest mover, up 16.2 points, as insurers built out the approval workflows and oversight structures needed for automated underwriting and claims decisions. Responsible AI Policy is the smallest mover, up 11.3 points, and it is

also the sector’s highest-scoring dimension at 69.6 out of 100, so its smaller gain reflects a dimension with less ground left to cover. Insurance has done the foundational work across governance and policy alike. The next task is sustaining that pace as deployment scales, so capability and confidence keep moving together.



LESSONS FROM LEADERS IN LIFE SCIENCES

In highly regulated environments, the path to greater AI autonomy runs directly through stronger governance and trustworthiness practices.



Life Sciences

TRUSTWORTHINESS

64.8 +13.7

Y1: 51.1

Key concern & priority

The priority is the data foundation. Life sciences has a structural advantage no other industries can replicate: the highest Audit & Accountability score, with that gap collapsing 20.9 points year-on-year to -0.6, the only inverted dimension in the entire data set.

62% of high-trust organizations flag data privacy as a concern vs 44% of low-trust peers. This is because they understand the full scope of what is in their data and its high-risk nature.

Leaders outpacing laggards at a remarkably high rate

The life sciences organizations surveyed span pharmaceuticals, medical device and equipment companies, and clinical research and development. The life sciences trustworthy AI leaders operate at a markedly higher level of maturity than their laggards in governance, continuous validation and responsible AI policy.

Ninety-one percent of leaders have cascaded a responsible AI policy to every employee, vs. 7% of laggards. Continuous validation and independent governance functions exist in the majority of leaders; in laggards, neither does.

Consequently, 55% expect strong or high AI ROI vs. 11% of laggards. The talent gap is equally notable: 22% of laggards report having no staff with AI ethics or risk expertise, limiting their ability to scale AI responsibly.

The leader group is not marginally ahead; it is in a different category.



We learned the hard way. It took us part of a full year before we could ascertain that our data is good enough for the system to learn and understand. The first year was generally learning, trying things out, making a lot of changes along our processes and our data.

Head of Data & Governance, Global Insurance (UK)

Governance is what converts AI into outcomes

For leaders in life sciences, AI sits where the regulator and the patient meet: regulatory submissions and evidence generation, manufacturing quality and supply chain, drug discovery and clinical trial design. Leaders are two to three times more likely to report these use cases as mission-critical to their business.

Thirty-three percent of laggards say AI is not yet mission-critical anywhere, compared to just 3% of leaders. Oversight splits sharply at the deployment line, with 56% of laggards requiring full human review. Until organizations apply the same audit rigor to AI governance that they use elsewhere, a highly cautious approach remains the safest, yet most costly way to manage autonomous AI systems.

Leaders have made that extension: They are far more likely to approve only exceptions or let AI act autonomously within defined boundaries. Leaders have earned the right to let the system run; the laggards have not.

The payoff shows up where life sciences competes: The governance maturity that lets leaders run AI with lighter human review is what makes it dependable enough to compress

regulatory submission and evidence-generation timelines, sharpen clinical trial design and efficiency, and tighten manufacturing quality. In this sector, governance is what converts adoption into faster, more reliable outcomes in the highest-stakes work.

The most demanding concerns of any industry

The concerns of the life sciences industry are the most demanding of any focus industry, fitting the consequences of getting AI wrong here. Leaders are more worried about data privacy than laggards (62% vs. 44%). Unexpectedly, the leader cohort is more worried because they understand the full scope of what's in their data and its high-risk nature.

Regulatory compliance, output reliability and ethical use all rise sharply with trust. The starkest comparison: 26% of leaders cite criminal or regulatory liability as a concern, compared to 7% of laggards. The organizations closest to clinical and manufacturing risk know exactly where personal and corporate liability lands if an AI decision goes wrong, and they govern accordingly.



In our initial journey, we realized the output wasn't trustworthy because our data wasn't ready. It took us almost a year to get our data good enough for the system to learn from.

Senior Director APAC Technology Services, Life Sciences (Singapore)



Life sciences: The five dimensions of trustworthiness




Figure 15: Life sciences – Five dimensions of trustworthiness

Life sciences posts the highest Audit & Accountability score of any focus industry. The sector's long-established systems for regulated quality and clinical trial audit trails extended into AI faster than perception has caught up. Responsible AI Policy follows the same logic.

The leaders are pulling further ahead and keep investing to keep it that way. The sector's structural advantage lies in already having a regulated, operational audit standard. The cost of extending it to AI is a fraction of building one from scratch.

LESSONS FROM LEADERS IN THE PUBLIC SECTOR

For public sector organizations, trust is not simply a deployment requirement; it is the condition that makes broader AI autonomy politically and operationally possible.



Public sector

TRUSTWORTHINESS

55.5 +6.5

Y1: 49.0

Key concern & priority

The priority is transparency and explainability. 41% of high-trust public sector organizations flag this as the key remaining gap compared to 26% of low-trust peers.

This is a sign that leaders have named limited transparency as the obstacle standing between where they are and where they intend to go.

Leaders are quickly scaling trustworthy AI investment

In the public sector, trustworthy AI leaders are operating fundamentally differently from laggards. Ninety-one percent of leaders have cascaded a responsible AI policy to every employee, compared to 5% of laggards. Eighty-four percent qualify as governance leaders, vs. just 1% of laggards. Continuous validation and independent governance functions exist across most of the leader cohort; in laggards, they are

almost entirely absent. Half of leaders expect strong or high ROI, in contrast to 4% of laggards.

On investment commitment, public sector leaders report the highest trustworthy investment-scaling rate of any industry surveyed. Forty-one percent are increasing trustworthy AI investment by more than 20% in the year ahead, which is as fast as the most ambitious organizations across any industry in the research sample.



Anything to do with AI, you can expect to wait a year before the product can go live. That's the problem we're facing. But understandably, it's also new. We have to make sure the solutions we're building are safe, because it can affect someone's life.

Senior Director of Technology, National Health Service (UK)



Citizen impact drives AI focus

For leaders in this sector, the AI use cases that matter most are analytical: data analysis, forecasting, policy evaluation, eligibility assessment, public safety and risk prediction. Leaders cite data analysis and forecasting as mission-critical at 75%, vs. 36% of laggards. The alignment is unsurprising, as these are the areas with the greatest citizen impact and the highest demands for explainability.

AI output validation divides on the familiar line. Fifty-nine percent of laggards require full human review: the most cautious posture of any industry-laggard pairing. This reflects both governance immaturity and the cost of an AI mistake in public service.

Leaders are more likely to let AI act autonomously within defined boundaries or under post hoc review only. Leaders have built enough confidence in their controls to give the system the ability to run; the laggards have not.

Leaders have moved past foundational concerns

The concern profile splits along a clear maturity line. Laggards remain focused on foundational risks: data privacy and security breaches, named at higher levels than in any other sector's laggard group. Leaders have moved beyond them.

Forty-one percent of leaders flag transparency and explainability compared to 26% of laggards – a sign that leaders have named limited transparency as the key remaining gap, and that closing it is how they intend to move further, faster. Ethical use, bias and managing AI autonomy all register prominently as well.

Reputational risk is another concern of leaders: 34% name it, against 19% of laggards. Like transparency, it is no longer a foundational risk but an external one, fitting for a sector where a single visible failure can affect the safety, security and livelihoods of whole populations, traveling further than in most private-sector contexts. The leaders are worried about the right things. The laggards are still worried about foundational risks.

Public sector: The five dimensions of trustworthiness




Figure 16: Public sector – Five dimensions of trustworthiness

The public sector scores below every other focus industry on the elements of trustworthy AI.

Banking, insurance and life sciences all score in the high 50s to low 70s across the five dimensions, with Responsible AI Policy the strongest dimension in each industry (69, 69.6 and 70.5 respectively). The public sector trails all three on every dimension, clustering instead

between 52 and 61 points: Responsible AI Policy still leads at 60.5, while Data Quality & Governance lags at 52, with Model Governance & Oversight, Explainability & Fairness, and Audit & Accountability in between. Model Governance & Oversight posted the sector’s biggest year-over-year gain (+3.8 points), while Responsible AI Policy moved the least (+0.7 points).



SMALL AND MIDSIZED ORGANIZATIONS

Size does not determine whether an organization builds trustworthy AI. The trustworthiness distribution for SMBs is functionally identical to that of enterprises: in the 60-to-80 band, where the bulk of the cohort sits, 41.4% of SMBs vs. 40.6% of enterprises; in the leader band (80+), 14.2% vs. 15.5%. Across every level of trustworthiness, the gap is at most 1.5 percentage points.

The small gaps that exist tend to favor SMBs. They match enterprises on the Trustworthiness Index (61.9 vs. 62) and score higher on realized impact (56 vs. 54.5) and perceived trust (71.5 vs. 70.9).




Figure 17: The five components of trustworthy AI – SMB

They also lead on the practices that build trust:

- More SMBs pair a responsible AI policy with training for every employee (33.8% vs. 31.6%).
- They invest in transparency and explainability (34.1% vs. 30.6%).
- They fund AI ethics hiring and training (32.5% vs. 29.4%).
- On explainability and dedicated ethics talent, SMBs match or exceed large firms.

SMBs also cite certain trust concerns less often. Just 46.7% cite data privacy and security as a top limit on trust, compared with 51.7% of enterprises, and fewer flag ethics or bias (24% vs. 30%). Their commercial priorities are more revenue- or commercial/outcome-focused: SMBs are significantly more likely to use AI to drive customer outcomes, prioritizing customer satisfaction (49.7% vs. 40.8%) and revenue growth (43.7% vs. 37.9%), and treating customer engagement and personalization as mission-critical (47.8% vs. 42.2%).

SMBs match enterprises on trustworthiness and outperform on both realized impact and perceived trust.

Their governance is lean and disciplined. More SMBs use exception-based human approval (34.5% vs. 31.4%), an efficient model that scales. The discipline shows up in results: Of the SMBs that quantified the effect, 39.7% realized an 11% to 20% improvement from AI in the past year, compared with 28.9% of enterprises.

Two gaps remain, and both track according to size. SMBs adopt the newest technology later (generative AI 65.9% vs. 70%, agentic AI 37.7% vs. 42.9%), and far fewer have stood up a dedicated AI governance or risk function (8.6% vs. 11.7%), relying instead on funded plans and named individuals. Scale still matters here.

Even with those two gaps, the section's finding holds: Resource abundance helps but does not determine the outcome. The same pattern appears beyond size.



Key takeaways for all

Organizations that invest in AI trustworthiness realize a 1.85 times gain across 13 business outcomes and are 15 times more likely to report strong or high ROI. Three structural implications follow.

Viewed collectively, the findings tell a consistent story. Organizations are entering a new phase of AI adoption where autonomy is increasing faster than trust. The leaders are distinguishing themselves not by deploying more AI, but by building the trust needed to use AI confidently in high-consequence decisions.

The compounding advantage

The advantage that data and trustworthy AI maturity deliver is already large, and it is compounding. The most consequential forward number in this study: 85% of leaders are growing trustworthy AI investment by more than 10% in the year ahead, vs. 26% of laggards. Leaders are not just doing it to protect their lead. They are very much funding their way further ahead.

The 15 times ROI multiple is where the two groups are today; the spend posture is where they will be 12 months from now. Leaders are reinvesting the trust dividend into increasing their advantage, and nothing in the data suggests laggards will catch up on their current trajectory. For every quarter without further investment, the cost of catching up gets larger.

The work to be done

The dimension that has improved least across this study, Explainability & Fairness, is also the single most common reason users override AI. That is not a coincidence: Most organizations did the visible governance work first (policy, audit, accountability frameworks) and left the harder engineering work like data lineage, model interpretability or interface design for later – and later has arrived. The human override is what it costs.

Importantly, data infrastructure is the foundation under all of it. A system can only explain a decision it ran on traceable data. Only 17.5% of organizations have reached a fully optimized data infrastructure. The rest are deploying AI on a foundation that does not support the validation that agentic AI requires. Organizations cannot resolve the agentic trust drop without first closing that data gap.

85% of leaders are growing trustworthy AI investment by more than 10% in the year ahead.

From trust to advantage: The road ahead

The path forward depends on where your organization is today. Some have already built strong foundations for trustworthy AI. Others are still developing the data, governance and explainability needed to scale. We've grouped our final recommendations into two categories: market leaders and organizations still building trustworthy AI capabilities.

For organizations still building

Close the five behavioral gaps. Progress beats perfection.

Prepare data infrastructure

Without traceable, high-quality data, the governance layers above it cannot function.

Validate continuously

Catching problems before deployment is structurally cheaper than managing overrides.

Codify policy and cascade it

The organization is required to codify policy and cascade it. You build this with a program, not a PDF.

Staff independent governance

You must staff an independent governance body with real authority to manage AI in your organization.

Fund explainability

Fund explainability as an engineering priority rather than a "nice to have" feature.

For trustworthy AI leaders

The next proving ground is agentic AI.

- ✓ Extend explainability, validation and audit disciplines into agentic systems that take actions rather than just making suggestions.
 - Treat the trust drop from generative
- ✓ to agentic AI as an opportunity if your governance is already built.
 - Make AI reasoning visible in the interface
- ✓ people interact with, not buried in documentation.
 - Treat the 85% growing their AI investment
- ✓ as a starting point, not your benchmark, as laggards aren't standing still.

For every organization

Trustworthiness makes AI deployments durable rather than slowing them down. Every dollar invested in data quality, model governance, explainability and audit buys the right to deploy more AI, more autonomously, with fewer overrides and lower overhead.

Leaders in this study have demonstrated this at scale, with measured returns. We know it works. The question is whether your organization moves before the gap becomes structurally difficult to close.

The cost of not building trust into AI is a here-and-now risk, visible in override rates, stalled pilots and ROI that accumulates while laggards wait for confidence to catch up with capability. Organizations investing now are closing the gap and building a lasting advantage.



Glossary of terms

Perceived trust

Subjective user/organization trust in AI.

Trustworthiness

Objective governance practices that prove AI reliability, fairness and oversight.

Trust gap

Perceived trust – Trustworthiness.

Larger gap = more risk of misplaced reliance.

Impact index

Aggregate of realized business outcomes from AI.

GenAI/agent AI

GenAI generates output; agent AI plans and acts autonomously across steps.

Low trustworthiness organizations = laggards

Organizations scoring 0-39 on the Trustworthiness Index.

High trustworthiness organizations = leaders

Organizations scoring 80-100 on the Trustworthiness Index.

Five components of trustworthiness

Data Quality & Governance
Model Governance & Oversight
Audit & Accountability
Explainability & Fairness
Responsible AI Policy

Agents

Agents perceive, reason and execute multistep work with limited human interaction points. Multiple agents can be orchestrated in an agent AI system.

Traditional (predictive) AI

AI that forecasts outcomes from data, the “predicts” layer beneath generative AI (creates) and agent AI (acts).

Override reflex

The tendency of users to override AI output; the clearest behavioral signal of how far people trust agent AI in daily work.

Trustworthiness Index

Composite 0–100 score of an organization’s governance practices across the five dimensions; used to classify leaders and laggards.

Data sovereignty

Data residency and jurisdictional control over data, where it lives and which regulatory regimes reach it.

Responsible AI

An organization’s policies and practices for using AI ethically, safely and in compliance with regulations.

Trustworthy AI

AI designed, governed and validated so it is reliable, fair, explainable and accountable.

Data maturity scale

1. **Ad hoc.** Data architecture is unstructured and lacks formal processes or governance. Data is inconsistently managed, and decisions are reactive.
2. **Siloed.** Basic frameworks and processes begin to emerge, but they remain fragmented with gaps in consistency and governance.
3. **Standardized.** Clear governance procedures, standards and operating models are established but compliance is incomplete.
4. **Managed.** Data architecture processes are integrated across the organization, with iterative updates reflecting business needs.
5. **Optimized.** Data architecture is fully optimized and continuously improved using metrics like KPIs.

To learn more, please visit
sas.com/ai-impact

To contact your local SAS office, please visit
sas.com/offices

