REPORT

STIBO SYSTEMS



State of Trusted Data:
The AI Trust Tax

Introduction

Enterprise artificial intelligence (AI) is moving beyond experimentation to become embedded in the decisions, workflows and customer experiences that drive daily business performance. As organizations refine their data environments, AI is helping teams improve productivity, accelerate decision-making, automate processes, and deliver more personalized experiences. True, enterprise-grade AI requires organizations to significantly invest in models, token consumption, both physical and virtual infrastructure, applications, agents, and implementation.

However, this technology investment is merely the first bill in a series. Now past the early stages of initial enterprise AI deployment, many organizations are realizing the hidden costs of a second bill: the time and labor inseparable from workforce AI-readiness training, output validation, data reconciliation, information preparation, missing context collection, explainability checks, and human oversight before actions are taken from AI-generated recommendations.

Stibo Systems calls this the AI Trust Tax.

To demystify how prepared enterprises are for this next stage of AI, Stibo Systems conducted a survey of 400 U.S. enterprise data and AI decision-makers. The research demonstrates that while organizations are realizing meaningful value from AI and reporting greater data maturity, trust has not been able to scale at the same rate. Notable challenges remain, with many still spending considerable time validating outputs, preparing data, and managing fragmented systems. Nearly four in five (80%) spend at least six hours each week preparing, reconciling, or validating data for AI and analytics, while fewer than three in ten (29%) say they can always explain the data behind an AI output. AI is presented as intelligence, but it is evident that leaders need trustworthy intelligence. The cost of AI output is the AI Trust Tax, a tax for which the accepted currency is significant validation, reconciliation, and human rework.

The First Enterprise AI Bill

Four in five

leaders have scaled AI across functions or embedded it into core business processes.

AI has progressed decisively beyond initial experimentation. Four in five (80%) respondents say their organization has either scaled AI across multiple functions or business units (46%) or deeply embedded it into core business processes (34%). A mere 6% remain in limited-pilot or non-use stages, reflecting an inflection point in the enterprise AI conversation. The question is no longer whether organizations should invest in AI, but how to scale and make AI dependable enough to scale across the workflows where accuracy, speed, customer outcomes, and compliance matter most.

Organizations exhibit deployment of a broad portfolio of AI capabilities, using or piloting an average of nearly five (4.84) applications of AI. Generative AI is the most common application, but AI utilization extends well beyond productivity tools into customer experience, autonomous workflows, predictive analytics, risk management, and enterprise data management.



The top motivations for AI investment included improving operational efficiency (51%), followed by increasing employee productivity (40%), improving customer experience (35%), and expediting decision-making (33%). These priorities raise the stakes for the data behind AI. It is evident that leaders view AI as a means to increased efficiency, but if AI relies on fragmented, inaccurate, or poorly governed business information, the labor it is meant to eliminate is at risk for manifesting in other ways.

The Second Bill: **The AI Trust Paradox**

72%

of respondents report high[1] confidence in AI outputs, yet 65% often or always manually validate, check, or rework them.

The Two Bills of **Enterprise AI**:

The First Bill Building AI	The Second Bill Making AI Trustworthy & Usable
AI models	Data gathering & preparation
Token consumption	Output validation
Physical & virtual infrastructure	Data reconciliation
AI applications	Data cleaning & preparation
AI agents	Explainability checks
Implementation	Human oversight & rework

Despite high AI confidence and data connectivity, quality, and consistency, leadership still reports significant manual labor associated with its usage. Nearly three-quarters (72%) report their business users are confident[2] in AI outputs. An overwhelming majority (94%) of respondents rate their organizational data quality and consistency for AI usage as excellent (42%) or good (52%). Nearly nine in ten (87%) respondents describe their organizational data as fully (28%) or mostly (59%) unified across business domains, systems, and functions. However, human intervention has not left the conversation. Nearly two in three (65%) report business users often (43%) or always (22%) manually checking, validating, or reworking AI outputs prior to action, suggesting a possible gap in data traceability and lineage systems. Significant oversight is still required at this stage of enterprise AI maturity, underscoring this paradox of perceived confidence and trust in AI.

Leaders are cognizant of the fact that AI output is fundamentally dependent upon its input. Manual work still belies the productivity and efficiency gains leaders intend AI to deliver. Teams spend significant time preparing data for AI and analytics, with four in five (80%) spending at least six hours per week on validating, reconciling, cleaning, or searching for data. Nearly half (48%) spend six to ten hours per week, and almost one-third (32%) spend more than ten. Nearly three in ten (28%) leaders say that AI outputs require too much manual review.

Inadequate data quality and governance prove to be hindering AI progress. Despite the overwhelming majority of respondents declaring their organizational data to be excellent or good, two-thirds (66%) still report that poor data quality, fragmentation, or insufficient governance has delayed an initiative in the past 12 months, with 16% saying that these issues have caused a significant delay. The most common challenges of unreliable or unusable data input into AI systems are teams being unable to explain or validate AI-generated recommendations (32%), data silos inhibiting AI development (31%), and AI outputs being inaccurate or inconsistent (30%). On the other end of the AI pipeline, the most common business consequences of unreliable outputs include heightened risk or compliance exposure (14%), reduced business user confidence (14%), and increased operating costs (12%).

1. Very or extremely confident
2. Very or extremely confident

The Autonomy Ceiling

As organizations transition from AI that supports employees, to AI that can take action on their behalf, the bar for entrusted intelligence rises. Agentic AI has the potential to automate complex workflows, increase response time, and reduce employee time spent on repetitive tasks. However, greater autonomy also inherently increases the surface area for unreliable, incomplete, or difficult-to-explain intelligence and actions.



The research shows that organizations are cautiously advancing towards agentic AI adoption, with more than one-third (37%) permitting autonomous action only with the caveat of human-review checkpoints and one-quarter (26%) permitting it in select business-critical workflows. Just 12% would allow broad autonomous action across core business processes. In order to grant increased autonomy to AI, leaders want greater confidence in AI accuracy (42%), better data quality (36%), stronger explainability and traceability (31%), and a clearer picture of ROI from AI initiatives (31%).

These results suggest an autonomy ceiling. While organizations recognize the potential of agentic AI, they remain hesitant to grant broad authority until there is greater confidence in the accuracy, quality, explainability, and traceability of the information directing AI actions. Allowing more to be under the purview of agents without improving the foundation can simply increase the potential for uncertainty. Agents lacking access to complete, current, and connected information may seemingly yield efficiency but later create downstream errors. Employees will then be forced to spend time reviewing and correcting these actions, further adding to the AI Trust Tax, rather than reducing it.

Characteristics of **Trustworthy Intelligence**

Artificial intelligence ought to be just that – intelligent. However, that intelligence comes with a caveat. Its intelligence only becomes truly trustworthy when it begins with unified, governed data. It depends on data that is connected, accessible, explainable, and relevant to the business decision at hand. AI cannot reliably recommend action without the relationships and context that connect products to suppliers, customers to transactions, inventory to demand signals, and assets to locations. Most organizations report progress toward this goal. Nearly nine in ten (87%) say their data is mostly or fully unified across business domains, systems, and functions. Similarly, 72% say most critical data is unified and governed across systems or available for AI-enabled workflows.

Further, these encouraging perceptions do not mean the work is complete. About one in two (51%) say trusted data is fully available when teams need it, while 46% say it is only moderately available. This underscores the difference between having a data foundation and operationalizing it for daily use. When prospectively considering the upcoming two years, nearly three in five (57%) leaders anticipate their organization's AI system making unreliable or harmful decisions due to lack of complete, accurate business context being a moderate to extremely significant risk.[3]

Organizations are strongest when it comes to linking data from customers and transactions (60%), products and customers (58%), and inventory and demand signals (57%). However, connecting other mission-critical data points remains more difficult: fewer than half can reliably link products and locations (45%) or assets and locations (44%).

For AI, individual data points are not enough. Its value depends on a holistic understanding of the relationships that shape business decisions. A supply chain planning system, for example, must connect product, supplier, location, inventory, and demand data; a product-compliance system requires product information, regulatory requirements, market data, and supplier relationships. Without this broader context, AI recommendations may be incomplete, unreliable, or difficult to explain. Too many disconnected systems (12%) was among the most frequently cited barriers to creating trusted, governed, and AI-usable organizational data, alongside competing technology priorities (13%), and the workforce talent gap (11%).



3. Moderately significant risk or very significant risk or extremely significant risk

The Answer: **Trustworthy Intelligence**

Note | This section compares High AI Maturity organizations (n=317) with Low AI Maturity organizations (n=83). Given the smaller Low Maturity base, indicated findings are directional. Analysis prioritizes trends and substantial differences.

- **High AI Maturity** was defined as those that responded, *"Have scaled AI across multiple functions or business units"* OR *"AI is deeply embedded into core business processes."*
- **Low AI Maturity** was defined as those that responded, *"We are not currently using AI"* OR *"We are experimenting with AI in limited pilots"* OR *"We have deployed AI in select functions or use cases."*

Organizations further along with their AI trajectory are not simply deploying more AI tools. Business intelligence that organizations can reliably trust is not incidental, but instead created through deliberate data quality improvement, allied systems and domains, formalized accountability and guidelines, dedicated AI teams, and established shared governance and access across functions. High-maturity organizations differentiate themselves through foundational investment to yield dependable AI. Investing in foundations and support structures translates to AI outputs that are more reliable, usable, and scalable.

Across all respondents, leaders at organizations most frequently improved data quality management (54%), lineage and traceability (50%), and AI-ready data infrastructure (50%). More than two in five have improved access to trusted business data (45%) and established novel data governance policies (44%). However, only 36% have consolidated or integrated their data systems, and just one third (33%) have created or expanded data stewardship roles. In comparison to low-maturity peers, high-maturity organizations are more likely to have invested in data quality management (56% vs. 48%), lineage or traceability (52% vs. 45%), and enterprise data management (38% vs. 27%). High maturity enterprises are more than twice as likely to experience strong positive ROI (31%) compared with low maturity organizations (14%). This suggests that enterprises recognize the need for heightened data ecosystems and AI-ready infrastructure, but many still have whitespace to bolster the organizational practices and connected data environments to enable trustworthy intelligence at scale.

- **4x |** High-confidence organizations are over four times as likely to always explain the data behind an AI output (37% vs. 9%).
- **2.3x |** High-confidence organizations are more than twice as likely to say trusted data is fully available when needed (61% vs. 26%).
- **2.7x |** High-confidence organizations are nearly three times as likely to report significant measurable AI value (51% vs. 19%).
- **2.2x |** High-confidence organizations are more than twice as likely to have scaled at least half of AI pilots into production (62% vs. 28%).

A similar pattern emerges amongst organizations with high confidence[4] in AI outputs. When business users do not fully trust AI outputs, respondents most frequently reported increased manual validation or rework (20%) as being the largest consequence. High AI confidence appears to reflect operational maturity, with nearly two in five (39%) saying AI is deeply embedded in their core operations, compared with just 18% of moderately confident organizations. Moderately confident organizations face greater likelihood that their business users exhibit AI tool adoption reluctance in comparison with highly confident enterprises (33% vs. 23%).

The research shows that highly confident entities exhibit greater data access, with the overwhelming majority (93%) saying their business data is mostly (59%) or fully (34%) unified. In comparison, 70% of moderately confident organization respondents report mostly (59%) or fully (11%) unified data across business domains, systems, and functions. Three in five (61%) at highly confident organizations say trusted data is fully available when needed for the teams building, deploying, or using AI, compared with only one-quarter (26%) at moderately confident organizations. Trusted data is available and accessible when it is most crucial.

Confidence is built on context and explainability. Those with high confidence in their organization's AI outputs are over twice as likely to say their company can connect relationships between key business entities extremely well when compared with those that are only moderately confident in AI outputs (39% vs. 18%). Those with just moderate confidence (38%) are more likely to say data silos inhibit AI development than their high confidence (28%) counterparts. Data traceability is also key. Highly confident organizations are over four times as likely to be able to track the source, history, and changes of data used in AI business systems than the moderately confident (41% vs. 10%).

Organizations with high AI confidence users exhibit structural differences from those with just moderate confidence. In the past 12 months, they are more likely to have invested in master data management or enterprise data management (41% vs. 25%), improved data lineage or traceability (54% vs. 41%), and increased investment in AI-ready data infrastructure (54% vs. 42%).

The governance model implemented by high-confidence organizations is also more developed when compared with moderately confident organizations. High-confidence organizations are significantly more likely to have a dedicated centralized data or AI governance council (25% vs. 16%), master data management or an equivalent enterprise data foundation (53% vs. 42%), and policies governing AI data use (61% vs. 52%). In addition, they are more likely to report that both business and IT teams share ownership, governance, and metrics for AI-ready data (28% vs 16%).

It is not enough to simply deploy AI more broadly. Hallmarks of AI maturity include greater operationalization, shared IT-business function AI responsibility, as well as dedicated AI teams, formalized policies, data management and traceability investment, and access to trusted business data. While these findings do not establish causation, they suggest that organizations that build a concrete foundation for trustworthy intelligence are more strategically positioned to utilize AI confidently, scale successful use cases, and ultimately, realize stronger business outcomes.

4. Very confident or extremely confident

AI Outcomes & Business Impact

AI produces tangible results, with the majority (83%) of all respondents realizing significant (42%) or some (41%) measurable business value rooted in AI, and 81% reporting strong (28%) or moderate (53%) positive ROI. A mere 5% report negative ROI from AI investments realized to date. Entities are measuring AI through fiscal and operational metrics, with organizations most commonly using productivity gains (47%), cost savings (42%), AI output accuracy (42%), revenue impact (38%), customer satisfaction / experience (38%), and decision-making speed (37%). Only three in ten (30%) utilize the number of use cases scaled to production from pilot.

AI value is evident, but the path from pilot to full-scale production is not immediate, with over one-third (34%) reporting that 25-49% of their AI pilots have scaled. Nearly half (48%) require four to six months to move from successful pilot into production, and one in five (20%) require seven months or longer. Sheer launched pilot numbers do not generate meaningful value, but instead value depends on enterprises' ability to operationalize successful use cases.

When business users exhibit confidence in AI outputs, there are clear positive impacts to the business. Compared with only moderately confident organizations, high-confidence entities are more than twice as likely to report significant measurable business value from AI initiatives (51% vs. 19%) and strong ROI (33% vs. 13%). They are also more likely to report transformational improvement in productivity (33% vs. 14%), decision-making speed (31% vs. 8%), and operational efficiency (34% vs. 6%). Additionally, high-confidence organizations are more likely to report that half or more of their AI pilots have moved into production (62% vs. 28%), suggesting greater scalability. While AI investment is costly, these associations suggest it is worthwhile.



Conclusion



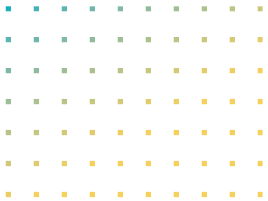
Enterprise AI has reached a new stage of maturity: it is widely deployed, delivering measurable value, and increasingly embedded in the workflows that shape operational efficiency, customer experience, and business decision-making. Yet scaling AI successfully isn't just about expanding use cases, it's about redefining what success looks like. AI succeeds when it produces trustworthy intelligence: outputs grounded in connected, high-quality, governed, and explainable data that give teams the context and confidence needed to act on outputs without excessive manual validation or risk.

The research shows that high-maturity organizations already treat this as the standard: they invest in the foundations that make AI dependable, including data quality, lineage, integration, governance, and access to trusted information. As businesses explore more autonomous and agentic AI capabilities, closing the gap between adoption and trust becomes the defining challenge. The organizations best positioned to realize sustained AI value won't simply be the ones with the most fragmented data transformed into something usable, they will be the ones for whom trustworthy intelligence is the measure of success itself: intelligence that is accurate, contextual, transparent, and ready to support action at scale.

REPORT

Call To Action: Three Priorities For Turning AI Into Action

AI usage necessitates due diligence, but the next step in its evolution is the infrastructure, systems, and accuracy improvement to minimize the manual labor inherent today. The goal is not to eliminate human oversight entirely, but to ensure that AI systems are properly monitored and securely given the tools necessary to yield more accurate outputs. Human oversight is imperative and irreplaceable to decisions affecting customers, compliance, safety, finances, and other high-stakes outcomes, but there is a difference between this sort of supervision and the rework incited by incomplete, disjointed, or difficult-to-explain information. When there is a paucity of trust in information completeness or data traceability and reconciliation, AI may simply shift or even proliferate labor rather than eliminate it.



1. **Invest in the conditions and processes for trustworthy data.** Prioritize data quality, traceability, lineage, integration, and governance to ensure AI recommendations are built on complete, current, and trustworthy business information. Invest in the human capital and training necessary for effective AI usage and implementation.

2. **Build context and traceability into AI, not just data access.** Link relationships across products, customers, suppliers, transactions, assets, locations, inventory, and demand so that AI can be more holistic in the analysis it bases recommendations upon. Reduce the manual validation burden by integrating trusted data practices into AI workflows.

3. **Measure the hidden work clouding AI processes.** Track the time and resources spent validating outputs, preparing data, reconciling records, and seeking context. These parameters will reveal where the AI Trust Tax is serving as a stop gap on productivity, speed, and scale.

About the study

Stibo Systems commissioned TEAM LEWIS to conduct an online survey of 400 U.S. full-time business leaders of manager level and above at organizations with 500 or more employees. All respondents influence their organization's enterprise data and AI strategy, with 82% serving as primary decision-makers. Respondents represent organizations across manufacturing, financial services, retail and quick-service restaurants, consumer packaged goods (CPG), life sciences, distribution, insurance, and other industries.

More than half (51%) of respondents work in IT, including data architecture, data engineering, governance, AI product, AIOps and related functions. The remaining respondents work across finance, operations, business administration, product development, sales, marketing, and customer-facing functions. Data was collected from 8/04/26-8/10/26.

About Stibo Systems | Stibo Systems is a leading enabler of trustworthy data through AI-powered master data management. Built on a robust and flexible platform, our SaaS solutions empower enterprises around the globe to deliver superior customer and product experiences. Our trusted data foundation enhances operational efficiency, drives growth and transformation, supports sustainability initiatives and bolsters AI success. Headquartered in Aarhus, Denmark, Stibo Systems is a privately held subsidiary of Stibo Software Group, which guarantees the long-term perspective of the business through foundational ownership. More at stibosystems.com.