



Building Defensible AI Decisions

How leading organizations test, evaluate, and stand behind AI systems —>

Only 23% of IT leaders are very confident in their organization's ability to manage security and governance when rolling out GenAI tools.

Source: Gartner survey of 360 IT leaders (May–June 2025), reported Oct 6, 2025

This confidence gap sits at the center of today's AI adoption challenge, leaving executive teams, risk and compliance leaders, and internal audit without the evidence they need to trust AI-influenced decisions in high-stakes environments.

Executive Summary

AI has entered a new phase of adoption. What was once confined to pilots is now embedded in operational workflows that power mission-critical decisions across enterprises, government, and regulated industries.

Leaders are no longer evaluated on whether they are adopting AI, but on whether they can stand behind the decisions those systems influence.


This requires more than incremental performance or productivity gains. **It requires the ability to explain how AI systems behave, show how risk is managed, and demonstrate that AI-influenced decisions are defensible over time.**

Despite rapid progress in model capabilities, many organizations remain underprepared for this level of scrutiny. AI is often evaluated using generic benchmarks, vendor-provided metrics, or one-time testing that does not account for how systems evolve once deployed. As a result, leaders are accountable for outcomes without clear, consistent evidence that their AI systems are behaving as intended.

In this eBook, we examine how that gap emerged and why it persists. Drawing on emerging AI policy frameworks and real-world deployment experience, we outline what trustworthy AI requires in practice, why traditional evaluation and governance fall short, and a practical blueprint leaders can use to make AI decisions defensible.

CONTENTS

- 03. — Executive Summary
- 04. — The Current State of AI Adoption
- 05. — The Need for an AI Evaluation Ecosystem
 - The AI Accountability Gap
 - Three Core Challenges in AI Model Evaluation
- 06. — Why Policy Frameworks Matter Beyond Government
- 07. — Where Traditional AI Evaluation Tools Fall Short
 - What Trustworthy AI Requires in Practice
- 08. — AI Trust Readiness Self-Assessment
- 09. — Where Organizations Struggle to Operationalize Trust
 - Moving from Assumptions to Evidence
- 10. — Seekr Guard: Making AI Trust Operational
 - Higher Standard for AI Decisions
- 11. — Example Use Case: Governing AI in Financial Risk Management
- 12. — Conclusion
 - About Seekr




The Current State Of AI Adoption

AI adoption has not followed a single, deliberate roadmap. In most organizations, it has unfolded incrementally through pilots, proofs of concept, and isolated tools that gradually became operational dependencies. Over time, AI systems have been integrated into workflows that matter, often without a corresponding evolution in oversight or governance.

Today, organizations manage complex AI ecosystems: multiple foundation models across teams, fine-tuning on proprietary data, AI agents embedded in workflows, and user interfaces that hide how models behave.


This complexity introduces a new level of risk. AI systems are not static assets. Their behavior shifts as prompts change, data distributions evolve, models are updated, or adversarial pressures emerge. Yet many governance practices still assume stability, relying on point-in-time evaluations that quickly become outdated.

The result is a growing disconnect between how AI systems are actually used and how they are governed in theory, which is where confidence begins to erode. While adoption accelerates, confidence often lags behind.



88% of organizations are now using generative AI regularly in at least one business function, a 10% increase from 2024, underscoring how fast adoption is outpacing governance.

Source: McKinsey & Company, "The State of AI in 2025: Agents, Innovation, and Transformation"



The Need for an AI Evaluation Ecosystem

America's AI Action Plan calls for a robust evaluation ecosystem to identify and mitigate risks in frontier AI systems, particularly in national security and other high-impact environments. For senior leaders, this is a signal that evaluation and testing are no longer optional; they are part of the expected control surface for AI.

Traditional, generic benchmarks often overlook mission-critical needs, which makes AI deployments ineffective and unreliable in practice. Many models carry embedded political and cultural biases and can be manipulated to generate deceptive or incomplete content, creating new vectors for operational, reputational, and regulatory risk.

The AI Accountability Gap

As AI systems move into high-stakes environments, accountability only intensifies. Leaders remain responsible for outcomes even when the systems influencing those outcomes are difficult to interpret or explain.

For many leaders, the real question is not whether AI works, but whether they can stand behind it when it is part of a critical decision.

In practice, this dynamic produces a familiar pattern. AI systems are deployed and deliver value, yet their influence remains limited. Teams use them cautiously, leaders hesitate to expand their role, and manual checkpoints re-emerge at critical moments, not because the technology fails, but because leaders are unsure they can explain and defend its behavior when it matters most.

Three Core Challenges in AI Model Evaluation

In practice, most organizations encounter the same three evaluation challenges, regardless of industry.

- 01. Choosing the right model for each mission.** Model selection is often driven by vendor claims and public benchmarks instead of rigorous tests on your own data and use cases.
For leaders, this means model choices are hard to defend when outcomes are challenged.
- 02. Knowing whether tuning made things better or worse.** Fine-tuning and prompt changes alter behavior, but teams rarely know exactly where performance improved, where it regressed, or how risk shifted.
Without clear before-and-after evidence, every change to an AI system introduces uncertainty.
- 03. Proving ongoing reliability, compliance, and governance.** Blind spots in drift, bias, and adversarial behavior leave leaders exposed without an audit-ready trail for boards, regulators, and governance teams.
Boards and regulators expect an audit trail, not assurances that "the model tested well once."



Why Policy Frameworks Matter Beyond Government

Much of the public discussion around AI governance has been framed in the context of government policy. Documents like America's AI Action Plan are often perceived as relevant primarily to public-sector institutions. In reality, the principles they articulate reflect broader shifts in expectations that now apply equally to enterprises operating in regulated or high-impact domains.

At their core, these frameworks emphasize that AI systems should be accurate in context, explainable to humans, secure by design, and accountable throughout their lifecycle.

For enterprises, the implication is straightforward: AI must be governed in ways that withstand scrutiny from regulators, boards, customers, partners, and internal stakeholders. The practical question is no longer "Do we align with these principles?" but "Can we show evidence that we do?"

The practical question is no longer "Do we align with these principles?" but "Can we show evidence that we do?"



Where Traditional AI Evaluation Tools Fall Short

Most AI evaluation methods were not designed to support this level of accountability. Public benchmarks and standardized tests provide useful signals about general model capabilities, but they rarely reflect the constraints, edge cases, or risk tolerances of real-world environments. Performance in the abstract does not translate reliably to performance in practice. A model that tops public leaderboards may still fail on your specific document formats, languages, or risk thresholds.

Vendor-provided metrics introduce additional challenges. Without independent validation or transparency into testing methodology, these metrics are difficult to audit or reproduce. When outcomes are questioned, they offer little defensible evidence. When outcomes are questioned and decisions are escalated, leaders are left defending opaque metrics they did not define and cannot reproduce.

Point-in-time testing presents another limitation. AI systems change continuously. Fine-tuning, prompt adjustments, data updates, and model replacements all affect behavior. Static evaluation creates a false sense of safety and leaves organizations exposed to silent degradation. Together, these approaches answer a limited question: whether a model performs well in general.

They do not answer the question leaders ultimately face: whether a system can be trusted in their environment, over time.


Today's Status Quo Creates Three Failure Modes:

- **Operational blind spots:**
Context-free, untested models mislead users and jeopardize integrity and security.
- **Inadequate benchmarks:**
Generic benchmarks and leaderboards miss business-critical requirements, making AI harder to validate, govern, and trust.
- **Embedded bias:**
Political and cultural biases shape outputs in ways that can distort facts and exacerbate inequities.

What Trustworthy AI Requires in Practice

In regulated and mission-critical environments, organizations build trust in AI only when a small set of operational requirements is consistently met.

01. **Evaluation must be grounded in context.** Models should be tested using real organizational data, workflows, and constraints. Context is the benchmark against which trust is established.
02. **Explainability must be built into the system.** If outputs cannot be traced, audited, and challenged, meaningful oversight becomes impossible. Explainability enables accountability by allowing humans to understand and contest model behavior.
03. **Risk must be made explicit.** Bias, hallucinations, model drift, and adversarial manipulation cannot be managed if they are invisible. Trustworthy systems surface these risks so they can be measured and addressed.
04. **Oversight must be continuous.** As AI systems evolve, evaluation and monitoring must evolve with them, with clear triggers for review and re-approval.



AI Trust Readiness Self-Assessment

The following self-assessment is intended to help leaders evaluate how prepared their organization is to deploy AI systems in high-stakes environments.

For each statement, rate yourself on a simple scale:

1 = rarely true, 2 = sometimes true, 3 = consistently true.

Evaluation and Context

Are we testing AI systems the way we actually use them?

- We evaluate AI systems using our own data, workflows, and scenarios, rather than relying primarily on generic benchmarks or vendor claims.
- Model performance is assessed against use-case-specific accuracy and risk thresholds that reflect real business and regulatory requirements.

Explainability and Oversight

Can we explain and control how systems behave?

- We can explain how AI systems produce outputs in a way that supports human review and accountability.
- When model behavior changes, we can trace those changes back to specific tuning, data updates, or configuration decisions.

Risk Identification and Monitoring

Are we actively looking for the ways our AI could fail?

- We actively test for failure modes relevant to our environment, including bias, hallucinations, and adversarial behavior.
- We monitor for model drift as data, usage patterns, or external conditions change.

Governance and Documentation

Can we show how AI decisions are governed and documented?

- AI evaluation results are documented in a consistent, repeatable manner.
- We maintain an audit-ready record of how AI systems were selected, tested, approved, and updated over time.

Operational Maturity

Is AI governance built into how we already run the business?

- AI governance is embedded into existing risk, compliance, or review workflows rather than treated as a one-time approval step.
- We have a clear process for reassessing AI systems as they evolve.

How to interpret your responses

Organizations that can confidently affirm most of these statements typically have the foundation needed to scale AI adoption in high-stakes use cases. Gaps often indicate where trust is being assumed rather than supported by evidence.

Focus first on any category where you scored mostly 1s; these are areas where a single incident could significantly undermine confidence if left unaddressed.

Where Organizations Struggle to Operationalize Trust

While many organizations agree with these principles, few have the infrastructure to apply them consistently.

Evaluation practices are often fragmented across teams.

Generating meaningful test data can be time-consuming and labor-intensive.

Comparing models across tasks or scenarios is difficult without standardized criteria.

Changes introduced through tuning or updates are rarely assessed systematically.

As a result, governance tends to be reactive. Attention focuses on incidents rather than prevention, and confidence is rebuilt only after failures occur.

This reactive posture undermines both trust and momentum—and it is exactly where a structured evaluation and governance capability can shift organizations from firefighting to prevention.

“Before deploying an AI model into a production environment, you need to know three things: the provenance of the data, the lineage, and the intent behind it.”

–Pat Condo, Seekr CEO

Moving From Assumptions to Evidence

Closing the AI trust gap requires a shift from assumption to evidence. Leaders must be able to show how a system behaves in their environment, and how that proof holds up as the system changes.

This fundamental shift is what enables scalable, defensible AI adoption. It transforms trust from an abstract aspiration into an operational capability.

SeekrGuard: Making AI Trust Operational

SeekrGuard is designed to turn AI trust from an aspiration into an operating discipline. It provides a structured way to evaluate, compare, and document AI model behavior using your organization's context as the benchmark, so leaders can move from assumptions to evidence.

By enabling consistent evaluation and documentation, SeekrGuard integrates governance into the AI lifecycle itself. Oversight becomes part of how AI systems are selected, tuned, and deployed, rather than an after-the-fact exercise.

SeekrGuard Core Capabilities

- **Model Scorecards:** Rank options across tasks, datasets, and scenarios for faster model selection.
- **User-Defined Risk Profiles:** Weigh what matters most by mission or business line and standardize governance criteria.
- **Configurable Evaluators:** Test edge cases, safety constraints, and adversarial prompts with consistent scoring.
- **Automatic Test Data Generation:** Turn internal documents and knowledge sources into test criteria and datasets, without heavy upfront work.
- **Flexible Testing:** Evaluate open-weight and proprietary models in approved environments.

A Higher Standard for AI Decisions

With the right evaluation infrastructure in place, organizations gain the ability to make AI decisions grounded in evidence rather than assumption. Model selection becomes more defensible. Tradeoffs between performance and risk become clearer. Emerging issues can be identified earlier, before they escalate into incidents.

Most importantly, leaders gain the confidence to stand behind the systems they deploy. In practice, this means leaders can answer three questions at any moment: which models we trust, where we are exposed, and whether recent changes made the system safer or riskier.

Governing AI in Financial Risk Management

Financial institutions increasingly rely on AI to support high-stakes decisions across credit risk, fraud detection, portfolio analysis, and regulatory stress testing. These models influence outcomes that affect customers, capital allocation, and compliance obligations. As AI becomes embedded in core risk workflows, leaders must be able to explain how models behave, how changes affect performance, and how risk is managed over time.

Target

Ensure AI models used in financial risk, credit decisioning, and compliance workflows are accurate, explainable, and defensible over time.

CHALLENGES

- AI models are often selected using generic benchmarks or vendor claims that do not reflect firm-specific data, risk thresholds, or regulatory requirements.
- Blind spots in bias, drift, and adversarial behavior make it difficult to demonstrate ongoing compliance to regulators, internal audit, and senior leadership.
- Model behavior can change significantly as a result of tuning, prompt updates, or data shifts, without clear visibility into where performance improves or degrades.

SOLUTION

With SeekrGuard, financial institutions can compare candidate AI models using their own data, scenarios, and risk criteria as the benchmark. Risk teams assess performance across relevant stress conditions, understand how tuning affects accuracy and bias, and apply consistent evaluation standards across the model lifecycle.

As models evolve, SeekrGuard supports repeatable testing against the same criteria, creating a documented record of performance, risk, and change over time. This integrates evaluation and governance directly into existing AI workflows rather than treating oversight as a one-time review.

- Compare credit risk models against historical downturn and regulatory stress scenarios.
- Produce audit-ready evidence for internal governance reviews and regulatory examinations.
- Detect model drift as market conditions or customer behavior shifts.

IMPACT

- Clear visibility into model behavior across real financial risk scenarios
- Stronger audit and compliance posture grounded in evidence
- More defensible model selection and deployment decisions
- Increased confidence in AI systems used for high-stakes financial decisions
- Earlier detection of drift, bias, and emerging risk

Conclusion

As AI systems take on greater responsibility, expectations around trust and accountability will continue to rise. Organizations that succeed will not be those that deploy AI the fastest, but those that can demonstrate understanding, control, and oversight as systems evolve.

Moving from policy principles to operational practice is the defining challenge of high-stakes AI adoption. Organizations that build evidence-based evaluation and governance into their AI lifecycle will not only deploy AI more effectively; they will be prepared to defend AI-influenced decisions with confidence when it matters most.

About Seekr

Seekr is a leader in explainable and trustworthy artificial intelligence designed to power mission-critical decisions in enterprises, government, and regulated industries. We provide secure, auditable AI solutions tailored to sectors where transparency, accuracy, and compliance are paramount.

As a SOC 2 compliant, privately held company, Seekr is committed to delivering innovative AI technology that meets the highest standards of security and regulatory compliance. Trusted by leading organizations, we partner with government agencies, defense organizations, and enterprises in highly regulated industries to enable AI solutions that drive real-world results with unmatched transparency and control.

