Origami Risk

The AI Maturity Roadmap for Risk and Insurance Leaders

Why moving deliberately isn't falling behind

AI is dominating the conversation. From boardrooms to industry panels, leaders are being told that now is the time to integrate AI or get left behind. The pressure is real, but so is the complexity.

For organizations managing underwriting, claims, compliance, or enterprise risk, adoption isn't as simple as flipping a switch. These are environments where precision matters, mistakes are costly, and change must be controlled. The gap between the promise of AI and practical implementation is easy to understand and not a sign of resistance, but of discipline.

> "There's a general acceptance that AI is pivotal. But even people who are very knowledgeable about AI largely still haven't changed their day-to-day habits."
>
> JAIME HENRY, VICE PRESIDENT, PRODUCT MANAGEMENT, ORIGAMI RISK

This disconnect isn't just a reflection of risk awareness; it's also a challenge of vision. Many organizations are still working through what a compelling, future-ready AI strategy looks like and how to align execution across teams once that vision is in place. AI is reshaping not only what's possible but also how it gets done, leaving many unsure of the right next move.

The real question becomes not whether to adopt AI but how to do it to deliver meaningful, long-term impact.

In this guide, we introduce our **TRUST framework** for AI adoption, built specifically for the complex, high-stakes environments of insurance, compliance, and enterprise risk.

This framework is designed to help leaders move with intent, deploying in ways that are:

- **Tactical:** Focused on real problems, not hypothetical use cases
- **Responsible:** Governed, explainable, and aligned with organizational values
- **User-controlled:** Configurable and embedded within existing workflows
- **Secure:** Designed for privacy, compliance, and risk-aware environments
- **Transparent:** Built on visibility, accuracy, and trust

These principles are not just conceptual—they are designed to be operational, helping to elevate areas where AI can add value, how to implement AI responsibly, and what it takes to move from isolated pilots to enterprise-scale impact.

You can read this guide front to back, or use it as a reference:

- If you're just beginning to explore AI, start with our TRUST framework to understand the foundational elements that should guide your approach.
- If you're further along, skip ahead to the AI maturity curve and use cases to assess readiness and identify where you can scale.
- If you're looking for strategic alignment, focus on the path forward section for insights on embedding AI into your broader risk and resilience strategy.

Wherever you are in your journey, the goal is the same: to make decisions about AI that are thoughtful, transparent, and built to last.

Part I: Origami Risk's TRUST framework

The core premise behind this TRUST framework is a practical model for responsible AI adoption. It is built around five principles that align innovation with the demands of risk-aware environments, each addressing a core requirement for AI to be both usable and sustainable. Together, they form a maturity model that guides not just where AI is deployed but also how it's governed, adopted, and scaled.

Tactical: Start small, solve real problems

AI should be grounded in business value. The most effective AI initiatives begin with specific use cases tied to measurable operational outcomes. This means identifying business pain points (areas where manual effort, delays, or inconsistencies create friction) and applying AI in ways that complement, rather than replace, human judgment.

WHAT TO DO:

- **Look for bottlenecks.** Where are claims getting stuck? What reports take too long to produce? What processes require repetitive tasks or handoffs between departments?
- **Prioritize use cases that save time and reduce friction.** Examples include summarizing long documents, drafting communications, flagging risks, or ingesting data from forms to eliminate rekeying.
- **Pick one use case.** Start with a narrow, low-risk pilot and measure the impact. Focus on repeatable tasks in well-understood workflows where AI can drive small, meaningful efficiencies.

WHAT THIS LOOKS LIKE:

- Drafting claims summaries automatically so adjusters can focus on decision-making.
- Turning long audit reports into a list of recommended actions.
- Pre-filling underwriting data from prior submissions or internal sources.

To ensure early pilots are generating real value, track simple, operational KPIs such as time saved p er task or user, reduction in manual entry or rework, or improvement in turnaround time.

Responsible: Governed, ethical, and human-centered

The Responsible part of the framework is about the organizational layer of governance: how decisions are made, who is accountable for them, and how AI is evaluated and improved over time. It focuses on the human, legal, and operational oversight that ensures AI use aligns with internal policies, ethical standards, and regulatory expectations.

Yet most organizations aren't there yet. According to Harness the winds of change, a 2025 report by Genpact, only 35% of insurers have the right KPIs in place to measure AI effectively. Without the ability to track accuracy, bias, or business impact, even well-intentioned initiatives can stall or quietly go off course. The risk isn't just that AI underdelivers but that no one notices until it's too late.

WHAT TO DO:

- **Assign clear ownership.** Someone must be accountable for how AI is implemented, monitored, and improved.
- **Keep humans in the loop,** especially for decisions with regulatory, financial, or reputational consequences.
- **Document the decisions AI makes,** how it's governed, and how you'll measure success over time.

WHAT IT LOOKS LIKE:

- A compliance officer reviews all AI-generated audit outputs before they're finalized.
- A risk manager logs which AI tools are used, what data they rely on, and how their performance is monitored.
- Quarterly reviews of KPIs, such as accuracy, rework rate, and user override frequency, are tied to AI adoption.

User-controlled: Configurable, *not* coercive

AI adoption hinges on whether your people trust the tool and feel trusted themselves.

AI gains traction when users feel in control. If a solution forces people to change how they work without context or choice, adoption stalls. But when AI is configurable and designed to support judgment, it becomes a trusted part of the workflow.

Importantly, this doesn't mean AI can't replace certain manual tasks or reshape how a workflow functions. In many cases, it should. But what users want to retain is control, visibility, and confidence that the system is working with them and not around them.

WHAT TO DO:

- **Let users choose where and how AI is utilized.** Allow for opt-in adoption and customizable configurations.
- **Make outputs editable and explainable.** Ensure users can review, adjust, and understand the results.
- **Match permissions to roles.** Not everyone needs full access. Ensure AI functionality aligns with responsibilities.

WHAT IT LOOKS LIKE:

- A claims adjuster generates a summary draft and reviews it before sending it.
- A risk manager pulls an AI-generated trend report but still chooses what goes into the final version.
- An underwriter uses prefill suggestions but validates before approval.

Adoption improves when AI enhances the user's role rather than complicates or obfuscates it.

Secure: Built for risk-aware environments

Secure focuses on the technical and data infrastructure layer of AI governance. It's about protecting sensitive information, ensuring compliance, and maintaining control over how AI systems operate behind the scenes, especially in high-stakes, regulated environments.

If your teams don't trust how an AI agent handles data, they won't engage with the output. This is especially true in regulated environments, where data privacy, auditability, and compliance aren't optional. Genpact's report reveals that 62% of business leaders cite data privacy as their top concern in adopting AI. That's for good reason: Most generative AI tools weren't built for enterprise risk use cases.

WHAT TO DO:

- **Use private, secure AI environments.** Avoid open tools that can use or leak your data.
- **Follow your existing compliance framework.** Make sure AI use aligns with standards like HIPAA and GDPR.
- **Control both inputs and outputs.** Define what data goes into AI tools and what can be generated, stored, or shared.

WHAT IT LOOKS LIKE:

- AI isn't allowed to send data outside your organization's secure environment.
- The system restricts sensitive claims data from being shared or used for training.
- Logs keep track of who used AI, what it generated, and where it went.

If the tool doesn't meet your standards for security and compliance, it's not ready.

Transparency: Trust through clarity

Transparency focuses on the explainability and visibility layer of AI adoption. It ensures users can understand, audit, and trust what AI is doing: why it made a recommendation, what data it used, and how its decisions can be validated or corrected.

Transparency is a prerequisite for scaling adoption. If users can't tell where an answer came from, how it was generated, or why something was flagged, they'll stop trusting it. In their report, State of AI Adoption in Insurance 2025, Roots found that 71% of insurers say accuracy is the most important factor in AI success. But accuracy is meaningless if it's not explainable.

WHAT TO DO:

- **Show your work.** Make sure AI-generated outputs come with context, not just conclusions.
- **Track performance over time.** Monitor where AI is helping and where it's not.
- **Make AI decisions auditable.** Ensure all AI actions and outcomes can be reviewed, traced, and explained.

WHAT IT LOOKS LIKE:

- A claims summary shows which case notes were used to generate it.
- A risk score links back to data sources or rules applied.
- Dashboards report on how AI is performing and how users are interacting with it.

Transparency is what earns trust in the system and in the decisions it supports.

Part II: Applying the framework

Every organization is at a different point in its AI journey. According to the Roots report, 45% of insurers are still exploring AI capabilities, 25% are testing, and 22% have a live solution in production. The remaining 8% have not yet started actively engaging with AI.

Wherever your organization falls, the important thing is to align your AI strategy with your current level of readiness and evolve it with intention. But readiness shouldn't become a reason to delay. With AI reshaping how work gets done across claims, underwriting, and compliance, organizations that postpone action risk falling behind peers already building capability and confidence.

Here's how to build TRUST across the AI maturity curve:

Exploration: Build awareness, not hype

The exploration stage is where most organizations begin and can get stuck. At this stage, the priority isn't building models or launching features. It's about building understanding. Teams are exploring what's possible, auditing their data, and identifying internal champions.

This is the time to lean into **Transparent** and **Tactical** thinking. Educate stakeholders on what AI can do, clarify its limitations, and start identifying where it could accelerate workflows, reduce friction, or enhance decision-making.

FOCUS ON:

- Auditing data quality and governance readiness.
- Introducing AI with clear internal messaging.
- Mapping potential use cases to business pain points.
- Setting realistic expectations about timeline and scale.
- Expanding AI features only where risk is clearly managed

Testing: Prove the value, protect the system

In the testing phase, organizations are piloting solutions, but success is still fragile. Without the right controls, even helpful tools can erode trust or create friction.

This is where **User-Controlled, Responsible,** and **Tactical** elements of TRUST become essential. Teams need tools that are opt-in, configurable, and easy to monitor. Human-in-the-loop design allows validation before decisions are made. Outputs should be explainable, and users should feel in control.

FOCUS ON:

- Tight pilot scope and clear success criteria.
- Role-based permissions and opt-in functionality.
- Gathering user feedback and adjusting in real time.
- Surfacing gaps in data, process, and oversight.

Production: Operationalize and govern

In production, AI becomes part of daily work. It supports claims, decisions, risk assessments, reporting, and communication. But with scale comes new responsibilities.

This is where Secure, Responsible, and Transparent pillars move from best practices to requirements. AI is now influencing real-world outcomes: financial, legal, and reputational. Outputs need to be traceable, access needs to be controlled, and usage needs to be reviewed.

FOCUS ON:

- Monitoring performance and usage trends.
- Auditing results and reviewing edge cases.
- Updating governance documentation and ongoing employee training.
- Expanding AI features only where risk is clearly managed.

Production: Operationalize and govern

In production, AI becomes part of daily work. It supports claims, decisions, risk assessments, reporting, and communication. But with scale comes new responsibilities.

This is where Secure, Responsible, and Transparent pillars move from best practices to requirements. AI is now influencing real-world outcomes: financial, legal, and reputational. Outputs need to be traceable, access needs to be controlled, and usage needs to be reviewed.

FOCUS ON:

- Monitoring performance and usage trends.
- Auditing results and reviewing edge cases.
- Updating governance documentation and ongoing employee training.
- Expanding AI features only where risk is clearly managed.

AI Workflow Readiness Checklist

Use this simple checklist to evaluate whether a use case is ready to move to the next phase — whether you're preparing to test, scaling to pilot, or expanding to production:

- ✓ What problem are we solving?
- ✓ Is our data clean, accessible, and compliant?
- ✓ Can we explain how the AI informs decision-making?
- ✓ Are stakeholders involved and informed?
- ✓ Do we have a plan for measuring impact?

If the answer to any of these is "no," pause and reassess. AI readiness is not a one-time check. It's a continuous filter for responsible progress.

Part III: The path forward

The purpose of AI isn't to replace decision-makers. It's to equip them with faster and actionable insights, automated workflows, and more time to focus on what matters. For organizations in risk-focused sectors, the goal isn't rapid adoption but sustainable innovation. That's why Origami Risk views the TRUST framework not as a checklist but as a mindset that informs how AI is introduced, scaled, and maintained.

So what does that look like in practice?

1. **Focus on foundations that last.** Success starts with structure: clean data, defined workflows, and aligned stakeholders. AI can only deliver if the inputs are reliable and the processes it supports are consistent. That means investing in the fundamentals: governance, access, and security. Then you can scale complexity.

   Want to go deeper on this topic? Read Get Your House in Order for AI to explore how data quality, system integration, and internal alignment lay the groundwork for effective AI adoption.
2. **Build alignment across teams.** AI doesn't live in a vacuum. It touches multiple functions and works best when those functions work together. Breaking silos between underwriting, claims, risk, compliance, and IT isn't just an organizational goal. It's a prerequisite for successful AI adoption.
3. **Move forward with intent, not urgency.** The most effective AI programs aren't always the flashiest. They're the ones that start small, learn quickly, and scale deliberately. Emerging capabilities like natural language search, automated summarization, and embedded decision support can add tremendous value, but only when deployed with clarity and control.
4. **Choose partners that support your vision.** As AI innovation accelerates, so does the need for infrastructure to support it. The right partner can help you adopt AI tools responsibly, integrate them seamlessly, and govern them at scale.

Platform-first, AI-ready.

Origami Risk helps risk and insurance teams build their foundation. Our configurable, secure platform enables the kind of data integrity, workflow control, and system connectivity that responsible AI adoption depends on. Whether you're integrating AI tools today or planning for them tomorrow, we help you get the fundamentals right, so you're not just using AI, you're using it well.

Let's talk about how Origami Risk can support your next step forward.

Connect with us.

About Origami Risk

Origami Risk empowers leaders in insurance, risk and safety with a purpose-built, cloud-native platform that optimizes workflows for better data, better insights, and better collaboration. Through highly configurable solutions integrated on a single platform, Origami Risk supports the management of the full lifecycle of risk, from prevention to recovery—helping the experts reduce harm and loss, and respond more rapidly and effectively when it happens. Grounded in continuous innovation and a foundational focus on client success, Origami Risk is trusted by leading organizations to enable greater resilience as they build for the future.

For more information, visit **origamirisk.com**