

# AI Governance Playbook



## A Simplified Guide to Jumpstart Your AI Governance Efforts

AI adoption is accelerating, but teams are getting stuck on governance, compliance, and real-world execution.

We've compiled this playbook to help you understand where risk is really emerging, what regulators actually care about, and how to evaluate and implement AI responsibly.

Read on for information and resources to learn:

- What AI governance actually requires (and what's overkill)
- Where legal and compliance risk is increasing
- How to evaluate AI vendors before you sign
- What responsible AI-driven consumer engagement looks like

# AI Governance 101 - Gaps We See in Financial Services

AI tools are being implemented by businesses at a breakneck pace, but effective governance is just as important as the efficiencies the technology offers. In fact, according to Gartner's Data Intelligence Monthly: Executive Insights on AI Governance (published February 11, 2026), organizations that deploy AI governance platforms are 3.4x more likely to achieve high effectiveness in their AI initiatives.

Here are AI governance gaps we've noticed in financial services:

- 1. AI governance isn't just a list of "nos"** - businesses need to develop a clear framework to develop a backbone of rules to follow.
- 2. Effective governance moves away from siloed AI oversight** and defines clear ownership with humans-in-the-loop who are responsible for monitoring AI performance.
- 3. Work to limit shadow AI** - employees shouldn't be using unvetted AI tools or personal accounts.

Ready for a full breakdown of these AI governance gaps?



[Watch the webinar recording](#)



[Read the blog](#)

Get Your AI Governance Starter Checklist on the Next Page>>

# AI Governance Starter Checklist



The debt recovery industry evolves almost as fast as technology. With more collection strategies than ever before using AI to optimize performance, these tools bring a new element of risk along with their efficiency.

In such a highly regulated industry, you need to make sure factors like data privacy, algorithmic biases, and consumer protections are all taken into account. **AI governance is about building a framework that helps keep your operations compliant. This checklist is designed to help you get started:**

## AI Governance Starter Checklist

### ☒ **Do you have an AI use policy?**

*An AI use policy defines how your business legally and ethically uses AI in your operations. It's human oversight over technology.*

### ☐ **Is there a governance committee or cross-functional AI oversight group?**

*A cross-functional oversight group is essential since it helps ensure AI being used in your business is compliant, safe, ethical and aligned with your goals.*

### ☐ **Does your oversight framework match the model's risk profile?**

*Oversight is not one-size-fits-all. While low-risk administrative AI requires minimal tracking, 'High-Risk' models—those impacting consumer financial outcomes or legal rights—require more rigorous oversight. This ensures your governance resources are focused where the regulatory stakes (like UDAAP or ECOA) are highest.*

### ☐ **Are vendors contractually obligated to follow your AI policies?**

*This practice helps limit the risks of AI hallucinations, bias, unauthorized use of your company data, legal liability, and more.*

☐ **Do you have a list of permitted AI tools?**

*Includes the business purpose of each approved tool and any prohibited uses of AI tools to help mitigate risk and uphold transparency.*

☐ **Are humans the final decision makers for AI systems?**

*Businesses cannot blame AI tools for a “bad outcome” or “compliance failure”, so it’s important to have humans at the helm.*

☐ **Has your business conducted AI risk assessments?**

*Identify the legal, ethical, and security threats for every AI use case in your business so there’s a clear roadmap to follow.*

## Power Your Collection Strategy with AI

If your business wants to optimize its debt collection strategy with AI, TrueML has solutions to help. Our third-party services and white-label debt collection software follow AI governance best practices set by legal experts.

Ready to learn more? Let’s get your AI governance and recovery strategy ready to deliver superior results.

“

Using Retain (powered by TrueML’s AI) has helped us optimize our communication channels by meeting customers where they are. **We see better interactions with digital** than just relying on phone calls, and have maintained performance levels **while reducing our operational costs.**

**Atlanticus** Kyle Sutter, SVP - Head of Collections

**Embrace AI smartly, safely, and effectively**

Learn more at [sales@trueml.co](mailto:sales@trueml.co)



### About TrueML

TrueML pursues innovative solutions to revolutionize the financial services industry through its subsidiaries’ products and services that improve client collection rates and customer experience. TrueML Products, maker of Retain, develops AI-powered products for companies and consumers seeking financial health. TrueAccord offers digital-first debt collection, leveraging data science and technology to deliver superior results.

<https://trueml.co/>



# AI Risk Isn't Where You Think It Is

By implementing AI tools into operations, there are plenty of opportunities, and more importantly, risks to consider. The risk of using this technology isn't always found in the most obvious spots.

With AI technology and tolls, businesses have to weigh efficiency against new risks. While it's impossible to eliminate risk entirely, your business should be working towards gaining a better understanding of the legal, evidentiary and structural shifts that accompany the technology. The risk present with AI tools just might not be where you think it is.

## → Attorney-Client Privilege Meets AI Risks

A court found non-attorney chats asking AI legal advice on a “publically-available, non-enterprise” LLM tool were not confidential nor privileged attorney-client communications.

## → The Risks of AI Scribes

A cross-functional oversight group is essential since it helps ensure AI being used in your business is compliant, safe, ethical and aligned with your goals.

## → Innovation Ownership with AI Tools

According to US law, every inventor listed on a patent has to be a person. So even though AI can help in the process, it can't “invent” in the legal sense. It's important to document the human work that goes into code and algorithms to help any IP be legally protectable.

### Want More Insight Into AI Risk?



[Watch the webinar recording](#)



[Read the blog](#)

### Don't Get Surprised by AI Risks

Your guide to defending against the next generation of AI risks is on the next page >>

# The Unseen Threat Defending Against Next- Generation AI Risks



By implementing AI tools into operations, there are plenty of opportunities, and more importantly, risks to consider. The risk of using this technology isn't always found in the most obvious spots.

To help you identify potential risks, TrueML's expert legal team crafted a list of questions as a launching pad to help you think about how you can mitigate risks and to set your business up for success.

## 1. Legal & Regulatory Landscape

The regulatory landscape around AI technology can change quickly. To stay ahead, businesses need to update risk assessments to evaluate the impact of AI on legal privilege and evidentiary discovery:

- **Privilege Waiver:** Are employees inputting privileged or NPI data into public/consumer LLMs? (*Reference: U.S. v. Heppner*)
  - **Mitigation Strategy:** Restrict sensitive work to enterprise-grade platforms with "no-training" contractual guarantees.
  - **Mitigation Strategy:** Educate employees that GenAI is not a substitute for legal counsel. Prohibit privileged information and/or attorney work product from being inputs into LLMs by non-attorneys.
- **Wiretap & Consent:** Does the organization use AI scribes/meeting assistants in "all-party consent" jurisdictions?
  - **Mitigation Strategy:** Implement mandatory verbal consent recording before transcription begins.
  - **Mitigation Strategy:** Educate employees about being thoughtful about what to record/scribe with AI, and to review all AI generated transcripts for accuracy.

- **Discovery Scope:** Is there a centralized list of AI tools used for business?
  - **Mitigation Strategy:** Establish an "AI Inventory" to track inputs/outputs for potential discovery requests.
  - **Mitigation Strategy:** Ensure your record retention policies account for AI-generated content (e.g., meeting minutes, outputs, etc.)

## 2. Agentic AI & Autonomy

Tools that leverage agentic AI often have some form of autonomy that needs to be monitored closely to reduce risk. It's important to assess risks associated with semi or fully autonomous systems that perceive, reason, and act independently:

- **Excessive Agency:** Does the AI agent have access to tools or data beyond what is essential for its task?
  - **Mitigation Strategy:** Enforce "Least Privilege" protocols and define specific "acceptable agency" boundaries.
- **Probabilistic Failure:** Are traditional deterministic tests (achieves reliable accuracy) being used for nondeterministic AI behavior?
  - **Mitigation Strategy:** Adopt an Agent Development Lifecycle (ADLC) that prioritizes "evaluation-first" over "code-first."



**Probabilistic Failure** (e.g., hallucination) is when an AI output is based on statistical probabilities instead of fixed rules.

- **Velocity Risk:** Can the organization intercept an autonomous attack advancing at "light speed"?
  - **Mitigation Strategy:** Design and test "Kill Switches" (soft and hard stops) at the orchestration layer.

### 3. Intellectual Property (IP) Inventorship

As AI tools push innovation forward, it is important to be on the lookout for new risks to intellectual property. Be sure to evaluate the protectability of AI-assisted innovations under current USPTO guidelines:

- **Inventorship Gap:** Are innovations primarily conceived by AI or a natural person? (*Reference: Thaler v. Vidal*)
  - **Mitigation Strategy:** Require contemporaneous logs detailing human formulation and refinement of AI outputs.
- **Vibe Coding:** Do vendors rely on probabilistic AI without human "conception" trails?
  - **Mitigation Strategy:** Audit vendor patent portfolios; mandate IP indemnification in all AI licensing contracts.

#### Proactive Defense Strategy: AI-to-AI Combat



As threats become "linguistic" (vibe attacks) rather than technical, traditional firewalls are insufficient. Action: consider including the deployment of defensive agents designed to identify and intercept offensive AI attacks in real-time.

“ Using Retain (powered by TrueML AI technology) has helped us fully automate email and SMS outreach efforts. **Engagement increased as did the number of scheduled payments**, reducing the number of eventual charge-offs on our books.”



**Stop Surprise Risk in Digital Debt Collection with Retain**

Learn more at [sales@trueml.co](mailto:sales@trueml.co)



#### About TrueML

TrueML pursues innovative solutions to revolutionize the financial services industry through its subsidiaries' products and services that improve client collection rates and customer experience. TrueML Products, maker of Retain, develops AI-powered products for companies and consumers seeking financial health. TrueAccord offers digital-first debt collection, leveraging data science and technology to deliver superior results.

<https://trueml.co/>



# 4 Pillars for Using AI to Communicate with Consumers

Most businesses aren't thinking about if they should use AI to communicate with consumers, but how they can put the technology in action. For highly regulated industries, it's important to put responsible engagement into practice to help limit risk. The legal experts at TrueAccord have identified four key pillars for ethical AI to consumer communications.

## 1. Building a Proper AI Data Integrity Foundation

Some of the best AI foundations are built to prevent hallucinations (instances where AI models make-up facts) and limit the amount of context gaps the technology has to fill in to generate an output.

## 2. Designing the Right Omnichannel Blueprint

After an AI is trained, your communication strategy needs to have a set playbook that the technology can follow. It's important to never let AI operate in a silo, and to have human talent in place to reduce risk.

## 3. Transparency in AI to Consumer Interactions

The next step is AI actually communicating with consumers, and the digital handshake should go smoothly. For starters, it's best to be transparent and make it clear the customer is talking to AI so they're not caught off guard.

## 4. AI Post-Performance Review

It's essential to have a performance review of AI technology interactions with consumers. Similar to cars, AI tools can experience performance drift and there needs to be humans-in-the-loop to recalibrate them so that outputs stay accurate.



**[Watch the webinar for a more detailed breakdown of each pillar](#)**

**Check Out the Next Page for 8 AI to Consumer Tips >>**

## 8 Tips Use AI to Communicate with Customers



Since AI has evolved from novelty to necessity for business communications, it's important to know best practices to help ensure this technology has a tangible impact. Here are 10 tips for businesses using AI to communicate with customers:

### **1 Keep Humans Accessible**

It's important to make the option to speak with a human agent clear. Chatbot and voice AI agent conversations shouldn't create endless loops for customers. And once a handoff to a human agent happens, the AI should provide all key details to set them up for success.

### **2 Maintain a Single Source of Truth (SSOT)**

AI is only as smart as its training data. Make sure to centralize your knowledge base and ensure all those resources have up-to-date information. This helps limit the possibility of hallucinations where incorrect information opens up liability risks.

### **3 Consider Using Sentiment Analysis for Prioritization**

Every customer message shouldn't be treated the same way. Natural Language Processing (NLP) technology can detect frustration, urgency, anger, tone or other keywords set in advance to automatically transfer complex conversations to human agents.

## 4 Transparency is Essential

It's best practice to be clear and explicitly state when a customer is speaking to AI, and in some states like Utah, it is a requirement if your company is licensed. Many consumers don't have issues talking to AI agents, but being up front about it helps build trust and keeps businesses ahead of compliance risks.

## 5 Implement Personalization at Scale

Customers expect a personalized experience when interacting with businesses. AI can use customer data like browser behavior, past interactions and real-time analysis to customize the conversation to their unique needs.

## 6 Embrace Proactive Support

AI has the ability to get in front of customer problems to create a more frictionless customer experience. For example, a debt collection AI tool could detect that a specific customer prefers to receive text messages for financial matters and send them an easy link to a self-service portal.

## 7 Start Using Agent Augmentation

AI can use the data collected from customers to empower your human agents. This technology can streamline essential yet time-consuming work like summarizing account information, which frees up more time for your agents to handle complex cases.

“ Using Retain (powered by TrueML AI technology) has helped us fully automate email and SMS outreach efforts. **Engagement increased as did the number of scheduled payments**, reducing the number of eventual charge-offs on our books.”



## 8 Keep Humans in the Loop

Like any other team member, AI should be regularly reviewed for performance and accuracy. Human managers should “coach” AI tools by correcting any missed opportunities or mistakes when communicating with customers.



Using Retain (powered by TrueML's AI) has helped us optimize our communication channels by meeting customers where they are. **We see better interactions with digital** than just relying on phone calls, and have maintained performance levels **while reducing our operational costs.**



*Kyle Sutter, SVP - Head of Collections*

### TrueML Uses AI to Prioritize the Customer Experience

Whether you're looking for a third-party collections partner or want to handle delinquencies more effectively in-house, TrueML offers AI technology to help you collect more from happier customers.

**Embrace AI smartly, safely, and effectively**

*Contact Us at [sales@trueml.co](mailto:sales@trueml.co)*



#### About TrueML

TrueML pursues innovative solutions to revolutionize the financial services industry through its subsidiaries' products and services that improve client collection rates and customer experience. TrueML Products, maker of Retain, develops AI-powered products for companies and consumers seeking financial health. TrueAccord offers digital-first debt collection, leveraging data science and technology to deliver superior results.

<https://trueml.co/>

# Essential Questions to Ask AI Vendors

AI technology has reshaped the way many industries do business, and there's a surge of AI tool implementation to help drive efficiencies. This trend has introduced a new challenge for businesses - figuring out the best way to evaluate complex AI technologies before a contract with a vendor is signed.

With the right questions, you can learn more about AI vendors and identify any potential risks before signing a contract. Ready for some key questions you should ask vendors?

## → What Exactly is the Underlying AI Tool?

It's important to have a detailed understanding of the AI tool. You need a clear understanding of the technology being purchased and how it will be implemented into your current system.

## → What Data is Being Used & How Is It Handled?

Data privacy is crucial when you're talking with an AI vendor. You need to know if your data is being used for training purposes or is being shared with any third-parties. It's considered best practice to ask for a full breakdown of how your data will be segregated and used.

## → How is the AI Maintained, Explained and Monitored?

AI tools aren't "set-it-and-forget-it" systems. Over time, AI technology can experience drifts or degrade and create unreliable outputs. Ask AI vendors how they're tracking the tool's performance to ensure it's running as intended.

### Learn How to Limit Vendor Risk with These Key Questions



[Watch the webinar recording](#)



[Read the blog](#)

**Don't Sign an AI Vendor Agreement Without Asking These Questions >>**

# AI Vendor Evaluation Worksheet



The right AI tools have the potential to optimize your business operations and strategy. With many AI vendors in the marketplace, it can be challenging to find a partner that follows industry best practices, understands your business, and doesn't raise compliance or operational red flags.

Use this unified master worksheet to vet potential AI vendors across your Legal, IT Security, and Business Operations teams. This ensures a 360-degree defense before your company puts pen to paper on any contract.~

## Theme 1: Corporate Governance, Risk & Framework Alignment

*This section establishes how the AI tool integrates into your high-level risk program, regulatory standards, and corporate accountability guidelines.*

- ☐ **AI Governance Policies:** Does the vendor provide documented corporate policies to formally evidence their internal AI governance?
- ☐ **Compliance Framework Mapping:** Are the vendor's internal AI lifecycle policies mapped directly to recognized industry standards (e.g., NIST AI Risk Management Framework, ISO/IEC 42001)?
- ☐ **Security Certifications:** Does the vendor hold up-to-date, verified security certifications (e.g., SOC 2 Type II, ISO 27001) covering their infrastructure?
- ☐ **Legal Documents & Model Provenance:** Has the vendor provided explicit license agreements, terms of service, and a technical Software Bill of Materials (SBOM) detailing underlying foundation model versions or software libraries?
- ☐ **Regulatory Adaptability:** Is the tool operationally supported and designed to dynamically adapt and update based on evolving compliance, regulatory, and legal requirements?



## Theme 2: Human Oversight & Agentic Workflows

*This section ensures humans remain at the helm, mapping explicit permissions, defining risk baselines, and tracking automated system dependencies.*

- ☐ **Human-in-the-Loop Safeguards:** Does the AI tool have a strict "humans at the helm" or "human-in-the-loop" workflow built into its daily operational architecture to prevent unsupervised risk?
- ☐ **Agent Blast-Radius Evaluation:** If an agentic workflow or automated model decision path were to completely fail, what is the calculated maximum amount of operational, financial, or regulatory damage it could cause?
- ☐ **Clearly Defined Agent Permissions:** What can the AI agent explicitly do, and what is the AI agent strictly prohibited from doing within your organization's infrastructure?
- ☐ **Administrative Feature Controls:** Which high-risk or generative features can be turned on natively by everyday end-users, and which features are safely restricted behind technical Admin approval?



**Red Flag Warning:** The AI provider does not provide clear AI permissions or how human oversight influences the outcome(s).

## Theme 3: Data Architecture, Privacy & Regional Security

*This section targets data inputs, ingestion rules, third-party sub-processors, and local data protection parameters.*

- ☐ **Data Element Verification:** Does the tool isolate and utilize pseudonymized metadata and transaction logs while strictly filtering out unvetted PII or sensitive account identifiers?
- ☐ **Pre-Ingestion Data Scrubbing:** Does your organization have the technical capacity or internal middleware to filter, tokenize, or clean data before it is ever sent to the vendor's servers?
- ☐ **Native Privacy Filters:** Does the vendor provide automated, built-in privacy masking filters to strip out personal data prior to model processing?
- ☐ **Third-Party Pipeline Boundaries:** If underlying third-party APIs are utilized by the vendor, do they maintain strict Zero Data Retention (ZDR) or isolated data contracts with those sub-processors?
- ☐ **System Integration Architecture:** Does the AI tool plug directly into your existing system of record via secure APIs, and has a clear data map been drawn defining how data is structured and partitioned?
- ☐ **Geographic Residency Tracking:** Does the vendor contractually guarantee that data processing, storage, and global traffic routing comply with local data residency laws?



**Red Flag Warning:** The AI provider does not make it clear how they treat PII or isolate your data sets.

## Theme 4: Accuracy, Explainability & Model Maintenance

*This section covers automated guardrails, performance tracking, bias testing, and mathematical transparency.*

- ☐ **Model Feature Transparency:** Does the tool provide deep mathematical transparency into the specific "features" or variables driving each individual decision to support strict auditability?
- ☐ **Plain-Language Explainability Audit Trails:** Does the system record and maintain plain-language reasoning logs for automated actions to allow legal and regulatory auditing?
- ☐ **Hard-Coded Regulatory Guardrails:** Is the AI built with an override-proof, hard-coded logic layer to enforce specific compliance guardrails, such as frequency caps or time-of-day restrictions?
- ☐ **Active Performance Monitoring & Progress:** How will your internal team systematically check in on the AI tool's work, and how will continuous positive or negative performance feedback be provided back to the AI vendor?
- ☐ **Post-Deployment Drift Tracking:** Does the vendor actively track and alert on statistical model drift or sudden drops in analytical accuracy on a continuous basis?
- ☐ **Bias Auditing:** (If high risk model) Are protected class variables cleanly segmented out, and does the vendor conduct routine bias testing to proactively mitigate disparate impact risk?



**Red Flag Warning:** The AI provider does not provide a means to review and audit the decisions or work done by the AI.

## Theme 5: Strategic Alignment, Implementation & ROI

*This section covers the commercial reality of procurement—buy vs. build decisions, baseline metrics, and vendor viability.*

- ☐ **Build vs. Buy Value Proposition:** Can this specific predictive capability be engineered internally by your current technical team? If not, what specialized, proprietary dataset or value layer is the vendor charging for?
- ☐ **Vendor Consolidation Check:** Does an existing platform currently inside your tech stack already offer this feature, and if so, what justifies adding a new vendor dependency?
- ☐ **Domain & Industry Expertise:** Is the vendor's engineering and product team backed by deep, native experience in your specific industry and regulatory vertical, or are they generalist software developers?
- ☐ **Transparent Delivery Timelines:** Does the vendor provide a multi-stage, predictable deployment blueprint with structured stages for data mapping, parallel safety testing, and production rollout?
- ☐ **Measurable Business KPIs:** Are the expected commercial improvements quantified upfront with strict, auditable baselines?
- ☐ **Sandbox & Staging Testing:** Can you test the AI tool inside a separate, air-gapped sandbox or fully contained system before full deployment?
- ☐ **Commercial Fee Transparency:** Are all variable processing expenses, server calls, and hidden platform costs contractually capped without any surprise downstream usage fees?



**Red Flag Warning:** The AI provider does not provide a means to review and audit the decisions or work done by the AI.

## Rate Your AI Vendor

How does your AI vendor measure up?

The more boxes you check, the more likely the vendor provides a stable, predictive solution rather than a high-risk tool.

Are there boxes they didn't check off that are critical for your business, strategy, or risk appetite?

You may want to vet other AI vendor options before signing.

“ When you **scale through technology**, you can achieve **better margins** while operating in a manner that is **more consumer-centric**.”



Thomas Overton  
Chief Product & Technology Officer  
TrueML

### Disclaimer

*This worksheet is not designed to be a fully exhaustive vetting process. The questions above highlight enterprise best practices and key trends with AI tools and vendors.*

**Embrace AI smartly, safely, and effectively. Let our experts help!**

**Contact Us at [sales@trueml.co](mailto:sales@trueml.co)**



### About TrueML

TrueML pursues innovative solutions to revolutionize the financial services industry through its subsidiaries' products and services that improve client collection rates and customer experience. TrueML Products, maker of Retain, develops AI-powered products for companies and consumers seeking financial health. TrueAccord offers digital-first debt collection, leveraging data science and technology to deliver superior results.



## Want More Insight on AI Governance? TrueML Can Help

This playbook and the accompanying webinar series has broken down the foundation of AI governance. If you have more questions or want to figure out the next step in your governance journey or vetting AI vendors, TrueML's expert legal team can help.

**Connect with our team today to get the conversation started.**



**TrueML**

### About TrueML

TrueML pursues innovative solutions to revolutionize the financial services industry through its subsidiaries' products and services that improve client collection rates and customer experience. TrueML Products, maker of Retain, develops AI-powered products for companies and consumers seeking financial health. TrueAccord offers digital-first debt collection, leveraging data science and technology to deliver superior results.

<https://trueml.co/>