LOOOM

WOVEN INTELLIGENCE. AUDIT-READY.

LOOOM AUDIT MANAGEMENT

LOOOM'S FOUR PILLARS OF AI GOVERNANCE

Executive Summary

Artificial intelligence is fundamentally reshaping the audit profession. Where AI once served as a search tool or document summariser, agentic AI systems can now plan, reason, take actions, and interact with external systems across multi-step workflows; all with minimal human intervention. For audit platforms like Looom, which integrate with large language models, this evolution opens extraordinary possibilities and equally significant responsibilities.

This white paper sets out how Looom approaches the governance of AI-assisted audit. It draws directly on the Model AI Governance Framework (MGF) for Agentic AI, published by Singapore's Infocomm Media Development Authority (IMDA) in January 2026, the most comprehensive and current guidance available for organisations deploying agentic AI. We map its four governance dimensions to the specific context of audit work, and translate each into concrete practices embedded in Looom's platform design.

The core argument is straightforward: audit findings carry real-world consequences. They inform financial decisions, regulatory compliance, and organisational strategy. An AI system that produces erroneous, biased, or unauthorised outputs in this context is not merely inconvenient, it can cause material harm to clients, regulators, and the public interest. Responsible deployment is therefore not optional; it is foundational to what Looom does.

Looom's Four Pillars of AI Governance

- Assess and bound risks upfront — defining the scope of what the AI can access and do
- Make humans meaningfully accountable — preserving auditor judgment at every critical step
- Implement technical controls and processes — across the full AI lifecycle
- Enable end-user responsibility — equipping auditors to oversee AI effectively

This document is intended for audit professionals, compliance officers, technology leads, and clients who want to understand how Looom governs its AI capabilities and what safeguards exist to protect the integrity of every audit engagement.

1. The Agentic AI Opportunity in Audit

1.1 What Agentic AI Means for Audit

Traditional AI tools used in audit - document classification, anomaly flagging, keyword search, operate as passive assistants. They respond to instructions and return outputs, but take no independent action. Agentic AI is categorically different. An agentic system can receive a high-level objective, plan a sequence of steps to achieve it, call external tools and data sources, adapt to new information mid-task, and complete multi-stage workflows autonomously.

For the audit profession, this means AI can now do far more than highlight relevant passages in a contract. It can cross-reference financial records with regulatory databases, identify discrepancies across multiple data sources, generate structured findings with supporting evidence, and flag items for human review, all within a single workflow. Looom's integration with large language models gives the platform access to these capabilities.

1.2 Why Audit Is a High-Stakes Domain

The MGF for Agentic AI identifies domain risk tolerance as a critical factor in assessing agentic AI risk. Audit sits at the high end of this spectrum for several reasons:

- **Irreversibility.** Audit findings are often irreversible in their consequences. Once a report is issued or a finding is communicated to a regulator, downstream effects cannot easily be recalled.
- **High-stakes outputs.** Incorrect findings can expose clients, auditors, and their firms to regulatory sanction, litigation, and reputational damage.
- **Sensitive data.** Data handled in audit engagements is typically sensitive, including personal financial information, trade secrets, and confidential communications.
- **Bias risk.** Bias in AI outputs; for example, in how the model weights certain types of transactions can produce systematically unfair or inaccurate conclusions across client portfolios.

These characteristics mean that the governance standards Looom applies to its AI capabilities must be commensurately rigorous. The MGF provides the framework for how we do this.



2. Pillar One: Assess and Bound Risks Upfront

The first pillar of the MGF for Agentic AI calls on organisations to assess risks before deployment and to limit the potential scope of impact through deliberate design choices. For Looom, this translates into two concrete practices: defining what the AI is permitted to do in an audit context, and enforcing those boundaries technically.

2.1 Scoping AI Involvement by Task Type

Not every audit task carries the same risk profile. Looom applies a structured approach to determining where AI involvement is appropriate, drawing on the MGF's risk factors:

Task Type	AI Involvement Approach
Document review and extraction	AI-assisted with human review of flagged items. Lower risk given read-only nature.
Anomaly and discrepancy detection	AI flags candidates; auditor confirms before any finding is recorded.
Cross-referencing data sources	AI executes lookups; outputs treated as draft pending human verification.
Final audit findings and reports	Human-authored. AI may draft; auditor must review, edit, and approve before issuance.
Client-facing communications	Human-composed only. AI may assist with drafting but cannot send independently.
Regulatory submissions	Strictly human-controlled. AI provides no direct output to regulatory systems.

2.2 Principle of Least Privilege

The MGF recommends that agents be given only the minimum tools and data access needed to complete their task. Looom enforces this through role-based access controls on the LLM integration. The model is granted read access to data relevant to an active engagement only. It is not granted write access to client systems, the ability to trigger external transactions, or access to data outside the current engagement scope.

This is not merely a technical preference; it is a design principle. As the MGF notes, an agent's action-space determines its potential for harm. Looom deliberately limits that action-space to what is necessary for audit analysis, reducing the blast radius of any malfunction or security incident.

2.3 Prompt Injection and Adversarial Input

Audit work involves processing large volumes of external documents; contracts, financial statements, correspondence that Looom feeds into the LLM for analysis. These documents may contain content designed to manipulate AI behaviour, a risk the MGF identifies as prompt injection. Looom implements input validation and sanitisation before external documents are passed to the model, and monitors outputs for signs of instruction-following inconsistent with the audit task.



3. Pillar Two: Make Humans Meaningfully Accountable

The MGF's second pillar addresses one of the most consequential risks in agentic AI deployment: the erosion of human oversight. As AI systems become more capable and reliable, the humans supervising them tend to scrutinise outputs less carefully. The MGF describes this as automation bias, and identifies it as a growing concern as agents become more sophisticated.

For an audit platform, this is particularly acute. The entire value of an audit rests on the professional judgment and accountability of the auditor. If that judgment is progressively replaced by unexamined AI output, the audit loses its integrity, regardless of whether the AI is technically correct.

3.1 Mandatory Human Checkpoints

Looom embeds structured approval checkpoints at points in the audit workflow where AI outputs translate into consequential decisions. These are enforced steps that cannot be bypassed. Consistent with the MGF's guidance, checkpoints are required before:

- Any AI-generated finding is recorded in the audit file
- Draft report sections containing AI analysis are finalised
- Any data flagged by AI as anomalous is escalated to the client
- Summary conclusions drawing on AI-assisted analysis are issued

At each checkpoint, the auditor is presented with the AI's reasoning, not just its conclusion. This design choice; showing the chain of reasoning rather than just the output is specifically intended to counteract automation bias by encouraging genuine evaluative engagement rather than a reflexive approval.

3.2 Auditing the Auditors' Oversight

The MGF recommends that organisations regularly audit the effectiveness of human oversight itself, recognising that over time human reviewers can become habituated to approving AI outputs without meaningful scrutiny. Looom implements this through oversight quality monitoring: where patterns of very rapid approvals are detected, the system flags these for quality review by a senior auditor or engagement manager.

3.3 Clear Responsibility Allocation

The MGF highlights that multi-stakeholder AI deployments can diffuse accountability in ways that leave harms unaddressed. Looom's governance model specifies clear responsibility at each level:

Role	AI Governance Responsibility
Engagement Auditor	Reviews and approves all AI-generated analysis before it enters the audit file. Accountable for every finding, regardless of AI involvement.
Engagement Manager	Reviews oversight quality metrics. Accountable for ensuring AI use on the engagement meets Looom's governance standards.
Platform Administrator	Configures access controls and tool permissions for the LLM integration. Responsible for technical boundary enforcement.
Looom Product Team	Maintains technical controls, monitors for model drift, and updates governance practices as AI capabilities evolve.

4. Pillar Three: Technical Controls and Processes

The MGF's third pillar covers the technical measures organisations should implement across the AI lifecycle during development, before deployment, and in ongoing operation. For Looom, this maps to the design of the LLM integration, pre-release testing, and production monitoring.

4.1 Technical Controls in the AI Integration

Looom has implemented the following controls in its LLM integration, consistent with the MGF's guidance on agentic system design:

- **Input validation.** Strict input formatting requirements are enforced before data is passed to the model, reducing the likelihood of malformed or adversarial inputs.
- **Constrained system prompts.** The model operates under a carefully constructed system prompt that defines its role, scope, and behavioural constraints for audit tasks. This prompt is version-controlled and reviewed when the underlying model is updated.
- **Comprehensive logging.** All model calls, inputs, and outputs are logged with timestamps and engagement identifiers, creating a complete audit trail of AI involvement in every engagement.
- **Data isolation.** The integration enforces per-engagement data isolation, ensuring that data from one client engagement cannot inadvertently appear in another's AI context.

4.2 Pre-Deployment Testing

The MGF recommends testing AI systems not just for output quality but for policy compliance, tool-use accuracy, and robustness to edge cases. Looom's testing protocol for its AI features includes:

- **Policy compliance testing.** Testing whether the model correctly follows Looom's audit-task scope constraints and does not attempt to perform actions outside its defined role.
- **Representative dataset testing.** Evaluating model performance on datasets representing the range of document types, financial structures, and anomaly patterns encountered in real audit engagements.
- **Edge case and adversarial testing.** Testing the model's handling of incomplete, ambiguous, or conflicting information common in real audit work to ensure it flags uncertainty appropriately.
- **Regression testing on model updates.** Retesting after any significant update to the underlying LLM to detect changes in behaviour that may affect audit output quality.



4.3 Continuous Monitoring in Production

Looom operates the following monitoring systems in production, consistent with the MGF's emphasis on continuous post-deployment monitoring:

- **Anomaly detection.** Automated alerts fire when the model produces outputs that fall outside expected patterns such as findings referencing data outside the engagement scope.
- **Regular output sampling.** A sample of AI-assisted audit outputs is independently reviewed each month to verify that model performance remains consistent with pre-deployment benchmarks.
- **User-reported issue tracking.** Any auditor who identifies an AI output they believe to be incorrect, biased, or inappropriate can submit it through a structured feedback channel reviewed by the Looom product team.
- **Graduated rollout on model updates.** Following any significant LLM update, Looom runs its full pre-deployment test suite in a staging environment before rolling the update to production engagements.



5. Pillar Four: Enable End-User Responsibility

The MGF's fourth pillar recognises that governance does not end with the platform it extends to the humans who use it. Organisations deploying AI systems are responsible for ensuring that end users have the information, training, and support they need to use AI responsibly and exercise effective oversight.

5.1 Transparency at the Point of Use

Consistent with the MGF's transparency requirements, Looom surfaces the following information to auditors at the point of AI use:

- Every AI-assisted output is clearly labelled, so auditors always know when they are reviewing AI-generated content versus human-authored content.
- The model's reasoning and the data sources it used are displayed alongside each output, enabling auditors to evaluate the basis of the AI's conclusions.
- The scope of the AI's data access for the current engagement is visible in the engagement settings, so auditors know exactly what the model has been able to see.
- Contact points for raising concerns about AI outputs are accessible directly within the platform.

5.2 Training and Capability Building

The MGF distinguishes between users who simply interact with AI outputs and users who integrate AI into their professional workflows. Auditors using Looom fall squarely in the second category and the MGF recommends layering education and training on top of basic transparency for this group.

Looom provides structured onboarding covering the AI's capabilities and limitations, common failure modes such as hallucination and overconfidence in ambiguous cases, best practices for prompting the system to get useful outputs, and the firm's policies on what AI can and cannot be used for in client-facing work.

5.3 Protecting Auditor Tradecraft

The MGF raises a concern that deserves particular attention in audit: as AI systems take over tasks that have traditionally served as training grounds for junior professionals, there is a risk of tradecraft erosion; the gradual loss of foundational skills. Looom is designed to augment auditor judgment, not replace it. AI outputs are presented as analysis for human evaluation, not conclusions for human approval.

Looom recommends that firms using the platform maintain deliberate programmes to ensure that less experienced auditors continue to develop core analytical skills, including working through some engagements with reduced AI involvement as part of their professional development.



6. Looking Forward: Governance as a Living Practice

The MGF for Agentic AI describes itself as a living document, acknowledging that agentic AI is evolving rapidly and that best practices will need to be continuously updated. Looom's governance approach is built on the same principle.

6.1 Monitoring AI Developments

The LLM that power Looom's analysis are updated regularly, and each update can bring changes in behaviour. Looom maintains a formal model update review process that assesses the governance implications of each update before it is deployed to production engagements. This includes re-running the full testing suite, reviewing any published documentation from the LLM on model changes, and assessing whether existing system prompts and controls remain fit for purpose.

6.2 Engaging with Regulatory Developments

AI governance frameworks for professional services, including audit are being developed by regulators across multiple jurisdictions. Looom actively monitors these developments and participates where possible in industry consultations. Our governance approach is designed to be readily adaptable as specific regulatory requirements for AI in audit emerge.

6.3 Commitment to Transparency with Clients

Clients of firms using Looom have a right to know that AI is being used in their audit, what it is being used for, and what safeguards are in place. Looom supports firms in meeting this disclosure obligation by providing clear documentation of AI involvement that can be shared with clients on request, and by maintaining the audit trail necessary to demonstrate human oversight of every AI-assisted finding.



Looom's Core Governance Commitments

- AI outputs are always presented as analysis for human evaluation, never as final conclusions
- Complete audit trails are maintained for all AI involvement in every engagement
- Human approval checkpoints are enforced at every consequential decision point
- Oversight quality is monitored to detect and address automation bias
- All AI features are tested against audit-specific scenarios before deployment
- Auditor training covers AI limitations, failure modes, and professional responsibilities
- Governance practices are reviewed and updated as AI capabilities and regulations evolve

7. Conclusion

Agentic AI presents audit with a genuine inflection point. The capabilities now available through platforms like Looom multi-step reasoning, cross-source data analysis, structured finding generation can meaningfully improve the quality, consistency, and efficiency of audit work. But these capabilities only create value if they are deployed responsibly.

The Model AI Governance Framework for Agentic AI provides a rigorous and practical foundation for responsible deployment. Its four pillars bounding risks upfront, ensuring meaningful human accountability, implementing technical controls, and enabling end-user responsibility; map directly onto the governance requirements of AI-assisted audit.

Looom has built its governance approach around these pillars. The platform is designed so that AI enhances auditor judgment rather than replacing it, that human accountability is preserved at every consequential step, and that clients and regulators can have confidence in the integrity of every audit engagement that uses Looom's capabilities.

As the technology continues to evolve, Looom's governance practices will evolve with it. This white paper is a statement of current practice and a commitment to the principles that will guide that evolution.

References

Infocomm Media Development Authority (IMDA). Model AI Governance Framework for Agentic AI, Version 1.0. Singapore, January 2026.

Infocomm Media Development Authority (IMDA). Model AI Governance Framework, 2nd Edition. Singapore, 2020.

Cyber Security Agency of Singapore (CSA). Draft Addendum on Securing Agentic AI.

World Economic Forum. AI Agents in Action: Foundations for Evaluation and Governance.

OpenAI. API Documentation and Model Cards. openai.com.