ripple treasury

Use of AI in the Ripple Treasury Ecosystem – Security & Compliance Overview

As part of our innovation strategy, we are integrating *Generative Artificial Intelligence* (GenAI) and *Agentic AI* to enhance client experience, streamline operational workflows, and deliver intelligent insights.

Introduction

This document outlines Ripple Treasury's approach to security, governance, and compliance in the context of GenAI and Agentic AI usage, with particular focus on protecting client data, managing AI risks, and aligning with applicable laws (i.e., EU AI Act) and internationally recognized standards, including the ISO/IEC 42001 standard for AI management systems.

Ripple Treasury – Information Security & Compliance Team

Table of Contents

1. Acronyms and Definitions

Term	Definition
AI	Artificial Intelligence
GenAI	Generative AI – models capable of generating human-like text, code, or other outputs
Agentic AI	An artificial intelligence system that can accomplish a specific goal with limited supervision.
Hallucination	In AI, particularly with large language models, this refers to the generation of false or nonsensical information presented as factual. It's when an AI produces content that is not grounded in its training data or real-world facts.
ISO/IEC 42001	International standard for AI Management Systems
LLM	Large Language Model – advanced language models trained on vast data sets
Output	The result or response generated by an AI system after processing input data or a query.
Prompt	The initial input or instruction given to an AI model to elicit a specific type of response or to guide its generation of content.

2. Agentic AI Capability

Ripple Treasury's Treasury Management

Ripple Treasury TMS includes agentic AI capabilities to further enhance the automation potential of its treasury ecosystem. Agentic AI systems go beyond passive assistance by initiating, planning, and executing multi-step tasks autonomously—always within defined parameters and controls set by the client.

In the Ripple Treasury ecosystem, agentic AI may be configured to:

- Analyze multiple data sources and perform scenario-based financial simulations
- Monitor key financial parameters (e.g., looking for anomalies or threshold breaches), and automatically trigger alerts containing contextual insights and recommended next steps for review or action by treasury staff
- Chain together multi-step queries to produce summary dashboards or reports

All agentic actions are governed by Ripple Treasury's AI risk framework and include built-in constraints, human-in-the-loop checkpoints, and audit logging. Agentic AI capabilities are deployed only within secured and isolated environments.

2.1. Human-in-the-Loop (HITL) Oversight

Generally speaking, Human-in-the-loop (HITL) is a concept used in multiple domains, and it can be defined as a model in which human interaction is required.

In the context of agentic AI, HITL ensures that critical decisions, escalations, or autonomous actions initiated by the AI are reviewed or approved by a human before execution.

To maintain transparency, accountability, and control, Ripple Treasury's agentic AI services incorporate Human-in-the-Loop (HITL) checkpoints.

HITL functions ensure that:

- Critical decision-support outputs are confirmed by authorized personnel
- Clients can leave feedback on AI-generated content as needed, such as corrections, clarifications, or further context

- Workflow configurations allow adjustable thresholds for HITL based on risk level, sensitivity, or context of use

By including HITL mechanisms and checkpoints, Ripple Treasury reinforces the principle that AI enhances – but does not replace – expert judgment, maintaining human accountability across business workflows.

3. Responsible Use of Client Data

Ripple Treasury guarantees that client data is never used to train Generative AI models. Each client's data is strictly isolated from others. Generative models are only used in inference mode, and all AI interactions occur within Ripple Treasury's secured Azure and AWS environments.

4. Security and Governance Controls

4.1. Data Security

Ripple Treasury applies robust security controls to protect client data at every stage — from ingestion to AI-enhanced processing — ensuring confidentiality, integrity, and availability of the data.

The security and governance controls implemented across the Ripple Treasury ecosystem include:

- Complete data isolation between tenants: Each client's data, including treasury records and AI indexes, is completely isolated at the application, storage, and AI layers. There is no co-processing or data sharing across tenants.
- End-to-end encryption:
  - Data in transit is encrypted using TLS 1.2+ protocols.
  - Data at rest is encrypted with AES-256 and FIPS 140-2 validated cryptographic modules.
- Strict access controls: Role-based access control (RBAC) is enforced across all components and Multi-factor authentication (MFA) is mandatory for all privileged access. Additionally, specific conditional access policies ensure that only authorized systems and personnel can access AI data and models.
- Data retention policies: Client data and AI indexes are retained only for the duration specified in the client's contract and data processing agreement.

- Auditability: Access to client data and AI components is logged and logs are collected and monitored as part of our internal procedures operated by our Security Operations Center.

4.2. Infrastructure and Network Security

Ripple Treasury enforces strict infrastructure and network security measures to protect client data and ensure isolation between components that handle sensitive financial information.

- **Network segmentation between components:** All key components — including APIs, serverless functions, and storage — are deployed in segmented network zones. Only explicitly authorized communication paths are allowed between layers, minimizing lateral movement risks.

- **Tenant data isolation enforced at all layers:** Every client's data is logically isolated at the application, storage, and AI indexing layers. There is no co-mingling of data between clients in any AI workflow.

- **No public exposure of raw client data:** Generative AI models are invoked in inference-only mode and operate within Ripple Treasury's private, secured infrastructure. Client data is never transmitted to public AI endpoints or external unauthorized third-party services.

- **Licensed AI models executed in secure cloud boundaries:** Ripple Treasury only uses enterprise-grade licensed models provided by Microsoft (Azure) and Amazon AWS, with strict execution boundaries to ensure client data does not leave Ripple Treasury's controlled environments.

- **End-to-end encryption:** All data in transit is protected using TLS 1.2 or higher, and all data at rest is encrypted using FIPS 140-2 validated cryptographic modules.

- **Automated security monitoring:** Intrusion detection, anomaly detection, and continuous monitoring are in place to track security events across the AI stack and core platform infrastructure.

- **Zero trust access principles:** Access to AI components follows zero trust principles, requiring strong identity validation, minimal privilege assignment, and continuous evaluation of access context.

4.3. AI Governance

Ripple Treasury has established a dedicated AI governance framework to ensure that all AI components are deployed responsibly, securely, and in full alignment with evolving global regulations. This framework includes proactive risk assessments and privacy impact evaluations, conducted before deployment, with a strong focus on privacy protection, data integrity, and regulatory compliance.

Clients in both the European Union and the United States can trust that Ripple Treasury's AI services are designed to meet applicable privacy laws — including the General Data Protection Regulation (GDPR) for EU clients, and the California Consumer Privacy Act (CCPA) along with other U.S. state privacy regulations. In addition, Ripple Treasury classifies its AI use cases according to the EU Artificial Intelligence Act's risk framework, with solutions categorized as limited-risk applications within financial services.

Importantly, AI generated outputs are fully traceable: every AI-generated response is linked back to the specific source documents within the client's own dataset, supporting transparency and auditability.

5. Compliance Alignment

Ripple Treasury's AI strategy aligns with:

- EU Artificial Intelligence Act

- ISO/IEC 42001 and ISO/IEC 27001 standards.

Governance structures, impact assessments, and transparent operational policies ensure responsible AI use and stakeholder confidence.

6. Transparency and Client Assurance

Ripple Treasury is committed to providing full transparency to clients regarding the use of AI technologies in its treasury management ecosystem. Our approach ensures that clients maintain visibility into how AI capabilities operate within the system.

Ripple Treasury is also committed to:

- **No use of client data for training:** Ripple Treasury guarantees that client data is never used to train GenAI models—data is only retrieved at inference time.
- **Client-specific data isolation:** Each client's data and index are completely isolated, ensuring that one client's data is never exposed or used in another client's context.
- **Flexibility:** AI-powered features are not enabled by default, as they're protected by a "feature flag". If the flag is off, no client data is passed to the AI at all; and even when enabled, the usage of AI services is explicit, clients must trigger it (i.e., by submitting a prompt).
- **Response traceability:** AI-generated outputs can be traced back to their source content, allowing clients to validate the origin and accuracy of AI-generated information when requested. Each action is recorded as a trace with rich metadata: session ID, durations, token usage, model used, filters applied, etc.
- **Model behavior transparency:** Known limitations and safeguards of the LLMs (e.g., hallucination risk) are documented and monitored to control and improve the efficiency of the AI service.
- **Data lifecycle controls:** Ripple Treasury has implemented data retention policies for their clients' data retained in indexes; these are also reindexing as needed.
- **Independent assessments:** Ripple Treasury may periodically engage third-party assessments to validate AI practices and demonstrate transparency and compliance to stakeholders.

7. Conclusion

Through licensed GenAI and agentic AI tools from strict data and network isolation, and compliance with the EU AI Act and the ISO/IEC 42001 standard, Ripple Treasury ensures AI is deployed responsibly and securely. Clients can trust that their data is safe, private, and never used beyond its intended scope.