ripple treasury

Knowledge Studio

Every treasury team has a liquidity policy. Almost none of it lives anywhere a system can act on.

Your investment policy is one document. Your counterparty limits are in another. Your intercompany rules are in a spreadsheet someone built three years ago. And the funding logic your team actually uses every morning, the adjustments, the trade-offs, the constraints your most experienced person carries in their head, lives nowhere at all.

So every day, the platform runs the workflow. And every day, your team supplies the judgment the platform cannot hold. When that person is out, the process slows. When they leave, the institutional knowledge leaves with them. And when an auditor asks how a decision was made, someone reconstructs the answer from memory.

GSmart Knowledge Studio makes your treasury policy the platform.
Your floors, ceilings, limits and rules become live machine-readable controls that every GSmart agent reads from, acts within and cites by clause. The policy that governed every decision is captured at the moment it happens, not rebuilt afterward.

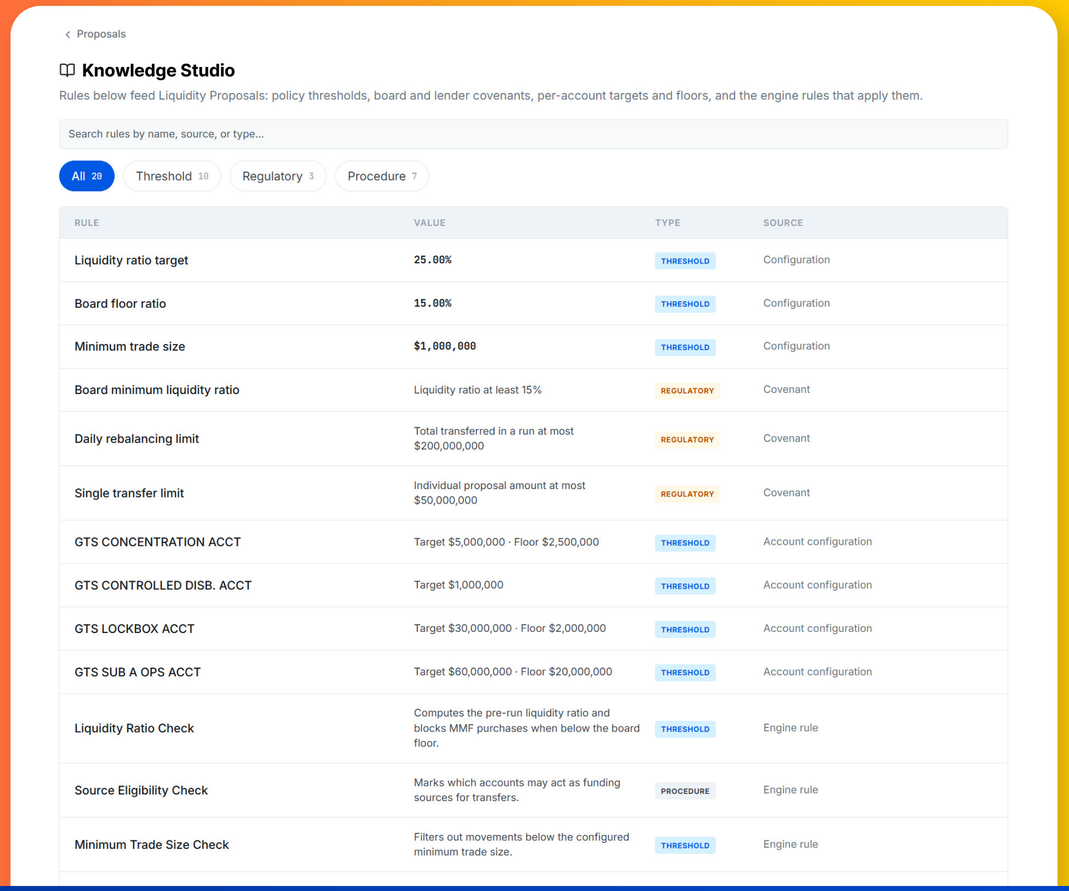

How Knowledge Studio works

Knowledge Studio holds your treasury policy as structured, machine-readable controls: floors, ceilings, targets, counterparty limits, intercompany rules and funding cascade logic. Not as a document. Not as a config file. As live constraints the agents operate within on every cycle.

Every GSmart agent reads from Knowledge Studio before it acts. Every proposal it generates cites the specific clause that authorized the move. Every override your team makes is recorded with its reason. At the close of each cycle, Knowledge Studio reads back what was enforced, what was overridden and where the policy may need to be refined.

Knowledge Studio meets your team wherever your policy actually lives today. Some teams have a formal policy document that just needs to be systemized. Some have rules scattered across spreadsheets, emails and shared drives. Some have a funding playbook that exists only in the head of the person who has run the morning rebalance for fifteen years. Knowledge Studio handles all three.

The longer the platform runs, the sharper the policy becomes. Every cycle adds to the picture of how your team actually decides, not just what the policy says.

GSmart

Key capabilities

Live policy controls
Set your floors, ceilings, counterparty limits, concentration rules and intercompany parameters directly as machine-readable controls. No document required to start. The controls go live immediately and every GSmart agent inherits them.

Tribal knowledge capture
For the positioning logic, conservatism adjustments and funding rules your team has never written down, Knowledge Studio provides guided setup to encode those controls directly. The judgment that lives in someone's head gets captured in the platform before it walks out the door.

Control-cited audit trail
Every agent action is linked to the Knowledge Studio control that authorized it, captured at the moment of the decision. When an auditor asks how a move was made, the answer is already there.

Override capture and pattern accumulation
Every time your team overrides an agent proposal, the reason is recorded alongside the action. Over time, Knowledge Studio builds a picture of the gap between what your policy says and how your team actually decides. Where a control should change, it proposes the change and waits for your review.

GSmart

Business outcomes

- Make your treasury policy enforceable by the platform, not dependent on the person who holds it in their head
- Capture institutional knowledge before it leaves when someone moves on
- Replace month-end audit trail reconstruction with a clause-cited record captured at the moment of every decision
- Ensure every new connector and counterparty relationship automatically inherits existing policy controls
- Give your audit committee a living record of what was enforced, what was overridden and why, across every agent cycle
- Build a sharper picture of how your treasury function actually decides, compounding with every cycle the platform runs

GSmart

AI security and compliance information

AI is a powerful amplifier of value for CFOs and Treasury teams. This power comes with great responsibility. So, trust in AI used must be earned. At Ripple Treasury, we believe transparency is the best first step to take in earning that trust.

How GSmart protects your data

Auditability and observability

Full auditability, every AI interaction is logged with a unique **trace_id** and captured by our observability platform, giving complete visibility into every prompt, response, and decision. Observability data is protected by the same encryption and access controls as the rest of the platform.

Explainability

No Black Box AI. All AI outputs are fully traceable to their originating data, with transparent reasoning steps and justifications included for every insight generated. Each customer's data and context are processed in isolation, ensuring insights are explainable and auditable, never mixed across clients. Detailed observability is maintained for every AI interaction, so you always know how and why each decision is made.

Risk classification

Active evaluation of AI use cases against the **EU AI Act** framework to ensure compliance and maintain low-risk classification.

Zero trust security

GSmart employs a comprehensive **Zero Trust Security Architecture**, anchored by Azure Managed Identity. This approach eliminates risks from hardcoded credentials and enforces strict role-based access controls (RBAC). All resources and interactions are continuously authenticated and validated. The principle of least privilege ensures systems are granted only the minimum access necessary, significantly reducing the attack surface and strengthening security resilience. All interactions are meticulously logged and monitored, enabling swift detection and response to any anomalies or unauthorized access attempts.

Agentic AI

Agentic capabilities in GSmart are securely constrained within customer-defined boundaries. Agents interact solely with customer-specific, isolated datasets and do not utilize or train on information from other clients. Each agent operates under stringent permission sets, with comprehensive auditability, and system-defined security guardrails, empowering users to innovate safely and confidently without sacrificing data privacy or control.

Security testing

GSmart employs rigorous, multi-layered security testing integrated into Ripple Treasury's CI/CD pipelines. This includes proactive "red team" testing against carefully curated golden datasets, validating inputs and outputs of prompts, comprehensive static and dynamic security scans on all code, and systematic benchmarking of AI models for vulnerabilities, compliance, and performance.

Encryption

All data is encrypted using advanced cryptographic standards. Data in transit is secured via TLS 1.2+ protocols, employing modern cipher suites such as AES-GCM, ECDHE, and SHA-384 hashing algorithms. Data at rest utilizes AES-256 encryption, with key management rigorously controlled and regularly rotated per Azure's enterprise-grade standards. Our encryption practices fully align with SWIFT Customer Security Controls Framework (CSCF).

Data usage for AI

AI data usage is governed by Inference-Only policies, with audit logs and verification mechanisms ensuring compliance.

Client control over AI

Clients have full control over AI capabilities through feature flags and can restrict which datasets are accessible for AI processing.

Data privacy & residency

AI is now available in every region of the platform, so storage and AI processing follow your business operating needs. Client data is encrypted in transit and at rest, stays in your selected region, and is never used to train models.

GSmart

About Ripple Treasury

Founded in 2012, Ripple is the leading provider of blockchain-based enterprise solutions across traditional and digital finance. Its solutions span global payments, custody, liquidity, and treasury management, serving as a one-stop shop for moving, storing, exchanging, and managing value. Ripple's stablecoin, RLUSD, and the cryptocurrency XRP underpinning these solutions allow Ripple and its customers to shape the modern financial system.

treasury.ripple.com

treasury-sales@ripple.com

ripple treasury