Security Overview

Sentrik Inc. | Shadow Mode | May 2026

Executive Summary

Sentrik is an AI agent governance and authorization platform. In shadow mode, Sentrik observes and evaluates agent-initiated financial actions in parallel with normal transaction flow — nothing in the existing payment stack is touched or intercepted.

This document summarizes Sentrik's current security posture, data handling architecture, and relevant certifications for the purposes of evaluating a shadow mode pilot integration.

> **Key Architectural Point**
> Sentrik never receives payment credentials, card numbers, bank account details, or PII from financial transactions. The evaluate endpoint accepts metadata only — action types, amounts, agent IDs, and context fields. This is a deliberate security design decision, not a product limitation.

Security Posture at a Glance

Data Minimization	**Production-Grade**
Access Controls & Auth	**Production-Grade**
Tamper-Evident Audit Records	**Production-Grade**
Encryption at Rest	**Covered via Supabase Pro**
Encryption in Transit	**TLS Enforced**
SOC 2 Certification *infrastructure provider is SOC 2 Type 2*	**In Progress**
Data Retention Policy	**Pilot-Configurable**
Data Processing Agreement	**Available on Request**

Data Minimization

Sentrik is architected from the ground up to handle metadata, not money. The evaluate endpoint is validated and enforced at the API layer to reject the following:

- Payment card numbers (16-digit sequences detected and rejected via Zod validation)
- Bank account and routing numbers
- Social Security Numbers and government ID formats
- Any field resembling payment credentials

The fields Sentrik accepts are limited to: action type, amount, currency, agent ID, vendor ID, customer status flag, human approval status, source system identifier, and an external event ID for reconciliation. No financial instrument data is transmitted, stored, or logged.

Access Controls & Authentication

- Row Level Security (RLS) on every database table — organizations can only access their own data
- API key authentication with SHA-256 hashing — raw API keys are never stored; only hashed representations are retained
- Service role separation — webhook endpoints use isolated service roles; UI sessions use user-scoped JWTs
- Role-based access control — admin, observer, and reviewer roles enforced at the application layer

Tamper-Evident Authorization Records

Every authorization record produced by Sentrik is SHA-256 hash-chained at creation time. The chain includes:

- Event hash — a hash of the incoming evaluation payload
- Decision hash — a hash of Sentrik's evaluation output
- Previous proof hash — links each record to the prior record in the chain

Records are effectively immutable once created. The hash chain provides cryptographic proof of record integrity and ordering, enabling independent verification of the audit trail.

Encryption

At Rest	All data encrypted at rest via Supabase Pro (AES-256). Supabase holds SOC 2 Type 2 certification.
In Transit	HTTPS enforced across all API endpoints. TLS 1.2+ required. All API communication encrypted in transit via Railway infrastructure.

Infrastructure & Certifications

- Database: Supabase Pro — SOC 2 Type 2 certified, GDPR compliant, encrypted at rest
- API hosting: Railway — TLS enforced, isolated container environments
- API documentation: Publicly available at shadow.getsentrik.com/api-docs

Pilot Data Handling

For shadow mode pilots, Sentrik operates as follows:

- All shadow events are scoped to your organization — no cross-tenant data access is possible
- Authorization records are retained for the duration of the pilot and configurable deletion windows can be agreed upon in writing
- A Data Processing Agreement (DPA) is available and can be executed prior to any data flowing through the evaluate endpoint
- Pilot integrations require only one additional API call per agent-initiated action — existing payment flows are not modified

> **Integration Footprint**
> Shadow mode integration requires: (1) one-time agent registration, (2) one-time delegated authority configuration, and (3) one POST call to /api/shadow/evaluate per agent action. Average response time under 100ms. No changes to existing payment rails, gateways, or merchant-facing systems.

Contact

For security questions, DPA requests, or pilot scoping:

Sentrik Inc. | getsentrik.com
Delaware C-Corp | Formed April 2026