SILENCE LABORATORIES

Enterprise Policy Engine

A rules-based control layer that enforces compliance and risk policy on distributed ledger transactions.

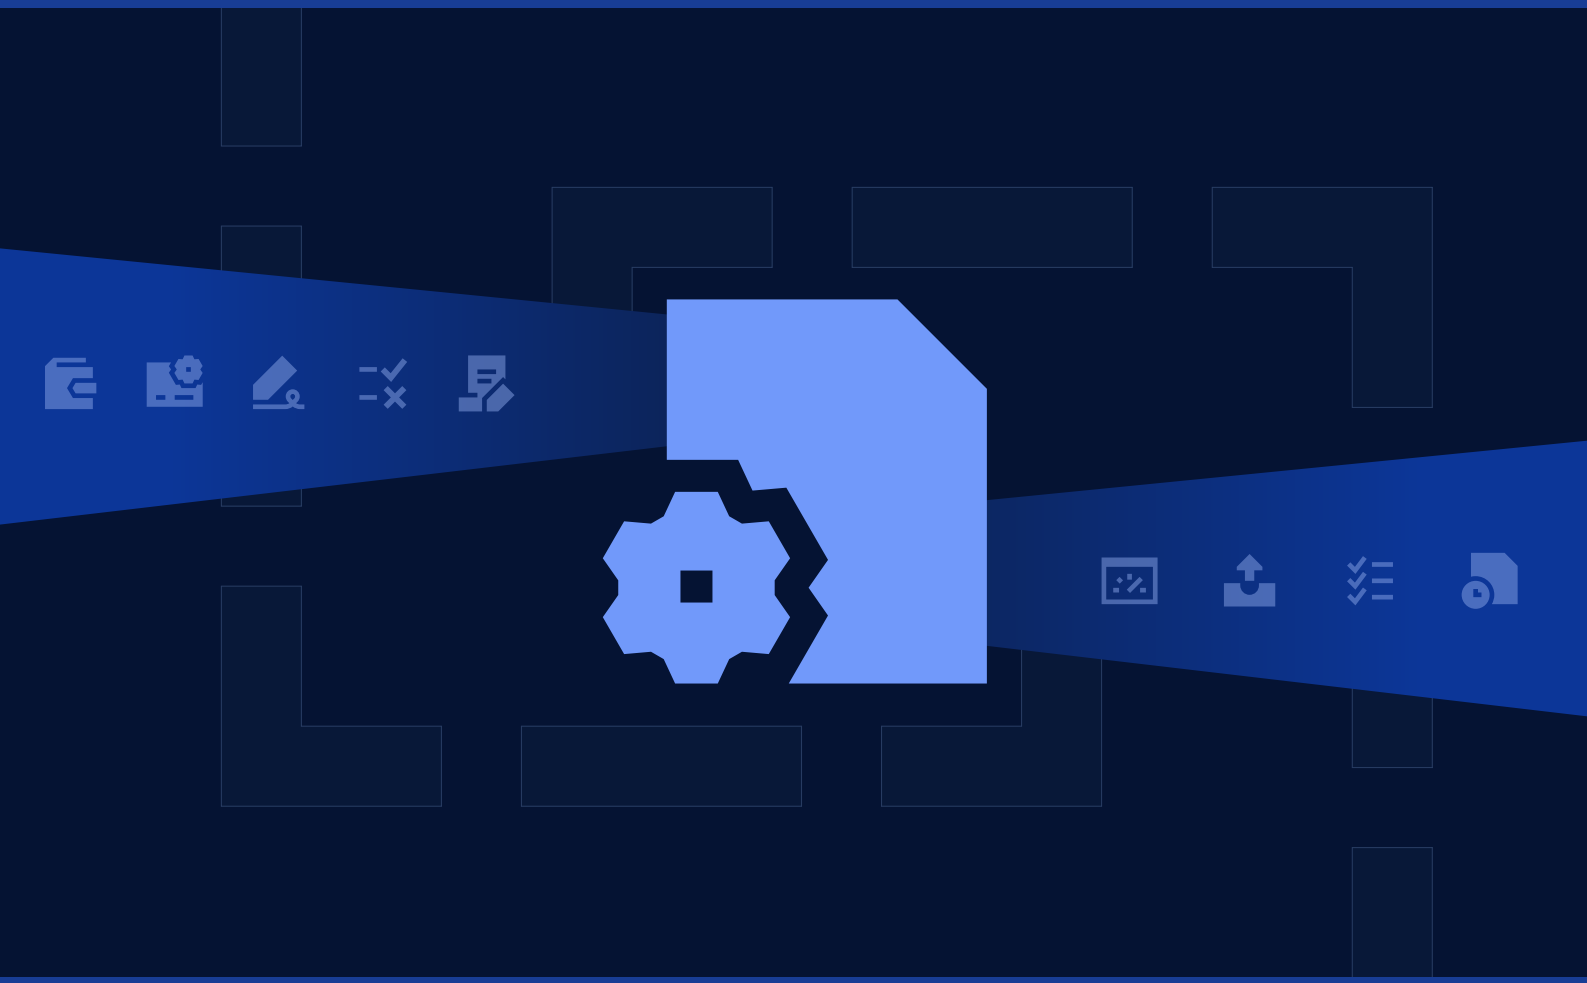

TABLE OF CONTENTS

1. Executive summary

This document describes the architecture, capabilities, and operating model of Silence's enterprise policy engine — a rules-based control layer that sits between transaction initiation and signing across blockchain wallets and accounts.

It covers the end-to-end architecture and flow from decoding through policy evaluation, approval orchestration, and signing handoff; the role-based access model governing who can author, approve, and audit policy; and the stateless and stateful rule types used to enforce whitelisting, velocity, and value-based controls.

It also details external compliance integrations (KYC/AML, sanctions, and risk scoring), the channels through which the engine is used (web console, mobile app, and APIs), supported deployment models (on-prem, distributed on-prem, and SaaS sandbox), the simulation capability for pre-publish testing, and the reporting and explainability features that keep every decision auditable.

Together, these capabilities let enterprises enforce compliance and risk controls on blockchain transactions without the policy engine ever holding custody of key material.

2. High level architecture

The policy engine is built as a modular architecture: every incoming transaction passes through a sequence of purpose-built components before it is signed.

The diagram below outlines each component in the order a transaction flows through them, from initial submission to final audit logging.

Web Console

Policy authoring, approvals, address book & audit search — role-based views.

API Interface

Full policy CRUD, submission, approval & audit query via REST — keyed on state-changing calls.

Transaction Decoder

Ingests the proposed transaction from the user or application and extracts the relevant values for policy evaluation.

External Signal Fetch

Parallel calls to external systems with a bounded timeout; results are cached and attached to the intent.

Approval Orchestration

If escalated, the engine creates an approval case, notifies the relevant approvers/groups, & collects cryptographically authenticated approvals until quorum is met or case expires.

Policy Evaluation

The rule engine evaluates the active policy version against the enriched intent.

Returns Allow Block Require-Approval

Reporting Engine

Every step is written to the immutable audit log and emitted as an event for ledger, reconciliation & monitoring systems.

Signing Handoff

Once allowed, the still-unsigned transaction is forwarded to the MPC/HSM Signing Network via an adapter interface.

The Policy Engine never holds key material.

Policy Database

Enterprise-grade store for policies, metadata, logs, and the full audit trail.

3. Key concepts & terms

Before describing the engine's rules in detail, it is useful to define the core building blocks referenced throughout this document.

Cluster

A logical grouping of multiple wallets — e.g. a business unit or a country jurisdiction. Policies, address books, users, and audit logs are logically segregated per Cluster, letting business units or legal entities manage accounts per their own requirements and regulations.

Wallet

A blockchain account (EVM, Solana, Cosmos and/or others) with a set of public and private keys. Policies can be applied to one or many wallets within a Cluster.

Rules

A group of ordered, configurable rules that check a proposed transaction before it is passed to the Signing Network.

Policies

The constituent components of policies, following a pattern of value, condition and action. Every rule resolves to exactly one of: Allow, Block, or Require-Approval.

Signing Network

The underlying MPC (TEE or HSM) network used to manage private keys for each Wallet.

4. Role-based access

Silence's enterprise policy follows a role-based access model to ensure compliance & security.

The engine supports the following roles and personas.

ROLE	PRIMARY FUNCTION	TYPICAL ORG PERSONA
Cluster Admin (or quorum)	Authors and reviews policy Defines velocity/exposure limits Configures AML/KYC triggers Reviews the audit trail	Compliance / risk officer of BU
Wallet User	Initiates transactions (payments, rebalancing, etc.) subject to policy.	Trader / treasury operator
Programmatic Initiator	An external application that triggers transactions subject to the same policy rules as human-initiated ones.	Trading bot / payments engine / tokenization service
Transaction Approver (or quorum)	Reviews and approves/rejects transactions, often as part of an M-of-N group.	Head of trading / treasury desk
Auditor	Reviews historical policy versions, approvals, and rationale, without the capability to alter records.	Compliance officer / regulator / accounting
System Admin (or quorum)	Maintenance and security of the platform: sets up new accounts, integrations, etc.	Ops / IT systems lead

Separation of duties is enforced: the system prevents a transaction initiator from counting as a required approver on the same transaction, and prevents a single admin from both authoring and unilaterally publishing sensitive policy changes

5. Types of rules

Every rule resolves to one of three outcomes.

Allow
Txn proceeds to signing.

Block
Txn is rejected.

Require-Approval
Escalated to approver quorum.

The engine supports two categories of rules. STATELESS STATEFUL

STATELESS Rules evaluated against the transaction payload

NAME	DESCRIPTION	OPTIONS
Address / account whitelisting	Permit or deny interaction with specific accounts / wallets.	Allow all · Block all · Allow list · Block list
Contract whitelisting	Permit or deny specific smart contracts.	Allow all · Block all · Allow list · Block list
Contract function check	Permit or deny specific smart-contract methods or selectors.	Allow all · Block all · Allow list · Block list
Token whitelisting	Validate calldata fields so allowed functions can only be called with permitted parameter values.	Allow all · Block all · Allow list · Block list
Per-transaction amount	Set per-transaction value thresholds.	All logical operators · < · > · = · <= · >=
Custom rule(s)	Set value thresholds for a custom signal/ indicator pulled from an external service	All logical operators · < · > · = · <= · >=

STATEFUL

Aggregate or velocity limits across a wallet, user, or actor.

NAME	DESCRIPTION	OPTIONS
Wallet velocity transactions	Limit the number of transactions allowed with a wallet in a time period.	`Rolling 24h / 7d / 1mo` `Custom range` `Per token & user/role` `Limit by recipient`
Wallet velocity value	Limit the value of transactions allowed with a wallet in a time period.	`Rolling 24h / 7d / 1mo` `Custom range` `Per token & user/role` `Limit by recipient`
User aggregate transactions	Limit the number of transactions allowed by a user across all wallets in a time period.	`Rolling 24h / 7d / 1mo` `Custom range`
User aggregate value	Limit the value of transactions allowed by a user across all wallets in a time period.	`Rolling 24h / 7d / 1mo` `Custom range`
Time-based blackout	Restrict transactions by operating hours or maintenance windows.	`Rolling blocks per day` `Custom date & time range`
Custom rule(s)	Set aggregate thresholds for a custom signal/ indicator pulled from an external service	`All logical operators` `<` `>` `=` `<=` `>=`

Default / fallback rule

Every policy has a system-enforced, non-deletable default rule. Its action is configurable at the cluster level but defaults to never Allow.

Auto-reject

Pending approvals auto-expire and auto-reject after a configurable timeout, avoiding stale authorizations lingering.

6. External integrations

The policy engine is designed to plug into an organisation's existing compliance and risk infrastructure. The integrations below extend policy decisions with external checks, alerts, and data feeds

KYC / AML / Travel Rule

Require and validate KYC/AML/Travel Rule checks for originator & beneficiary accounts via API integrations, performed before a transaction is evaluated against any policy.

Alerting service

Emit alerts for denied transactions, compliance hits, failed webhooks, and failed operations — deliverable to the dashboard, mobile app, or channels like Slack, Telegram, SMS, and email.

External scores & data

Ingest signals from external services (e.g. Chainalysis, Elliptic, Pyth) or internal risk scoring systems and use them as conditions in any rule — e.g. checking a beneficiary's risk score to allow, block, or require approval

Freeze on Alert

A configurable action that freezes an account, address, or specific transaction pending manual compliance review when a risk signal arrives from an external monitoring system.

7. Interaction options

Users and systems interact with the engine through several channels, ranging from human-facing consoles to programmatic access for automated systems.

Web Console
Policy authoring with diff/preview, approvals, address-book management, and audit search — with rolebased views.

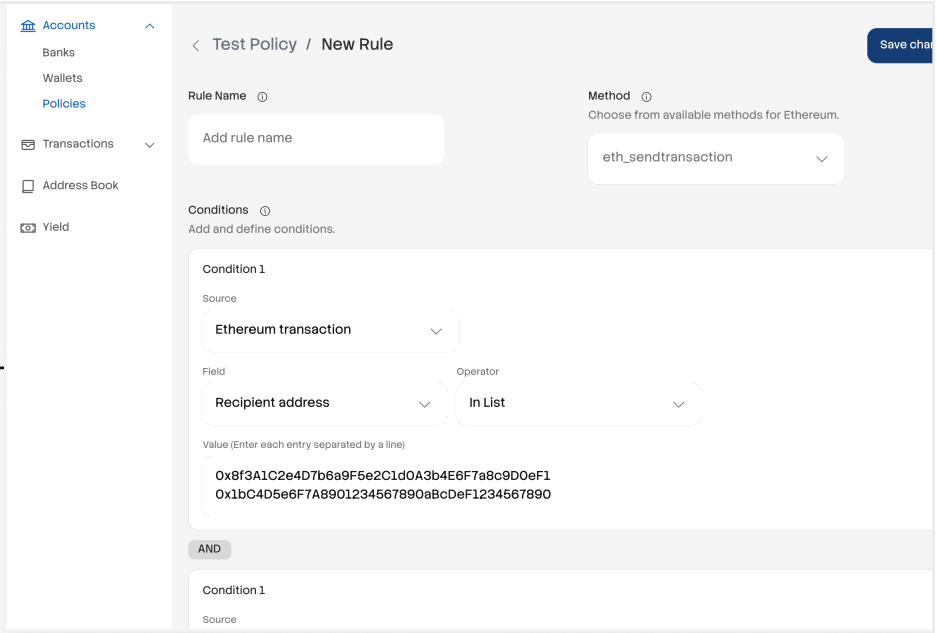

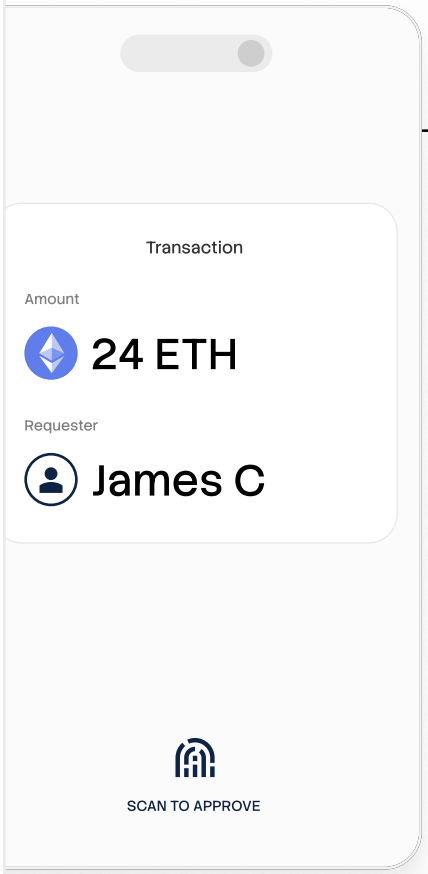

Mobile App / SDK
Native iOS/Android app supporting biometric-gated approval actions and push notifications

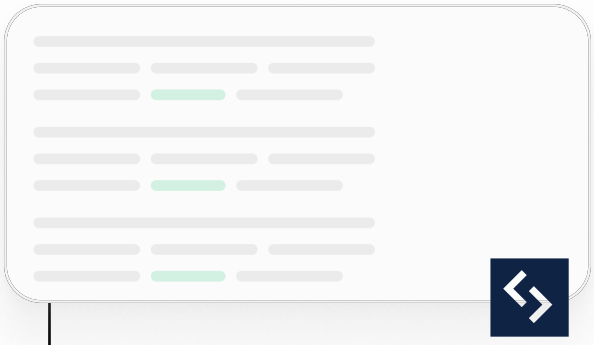

External Systems
Full policy CRUD, submission, approval, and audit query via REST and/or gRPC — with idempotency keys on all statechanging calls.

8. Deployment models

The engine can be deployed in the configuration that best matches an organisation's data-residency, security, and operational requirements — from fully managed to fully self-hosted.

On-Prem
Supports data-residency requirements and compliance rules, removing reliance on Silence infrastructure.

Best for
Bank-grade enterprise

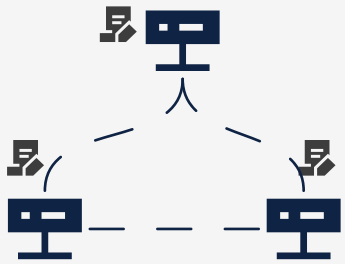

On-Prem (Distributed)
Deployed across multiple MPC network nodes for redundancy, decentralization, and censorship resistance.

Best for
Consortiums / payment networks

SaaS Sandbox
Hosted by Silence as an external service, with the fastest path to a running environment.

Best for
PoC / experimental use

9. Simulation

To reduce the risk of unintended outcomes going live, **the engine provides simulation at both the transaction and policy level** — validating behaviour before anything is enforced in production.

Transaction simulation

Before signing, the engine runs a full preapproval simulation to determine the exact effect on wallet balances and token holdings

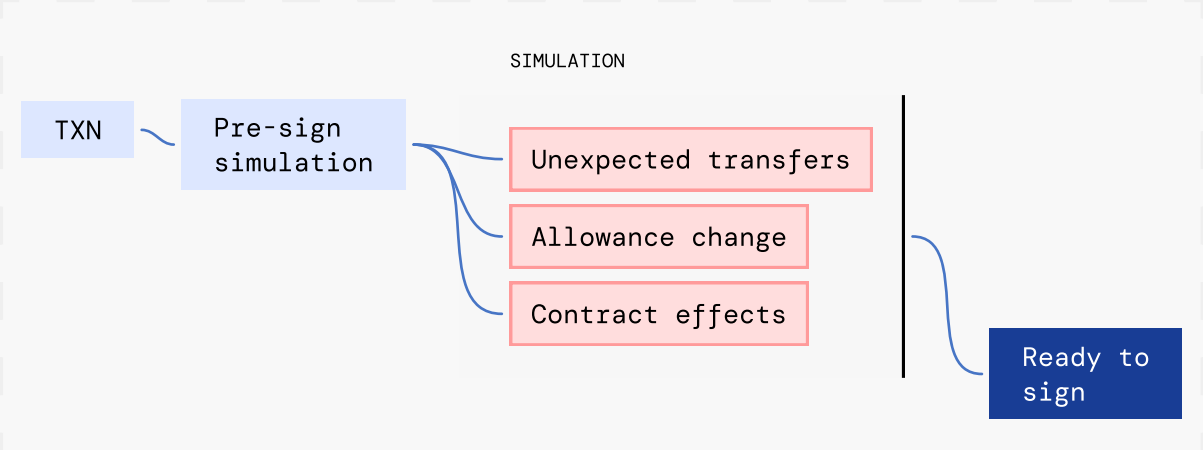

Policy simulation

Before publishing, admins replay a new or amended policy against historical transactions or synthetic cases to see which rule would match each transaction.

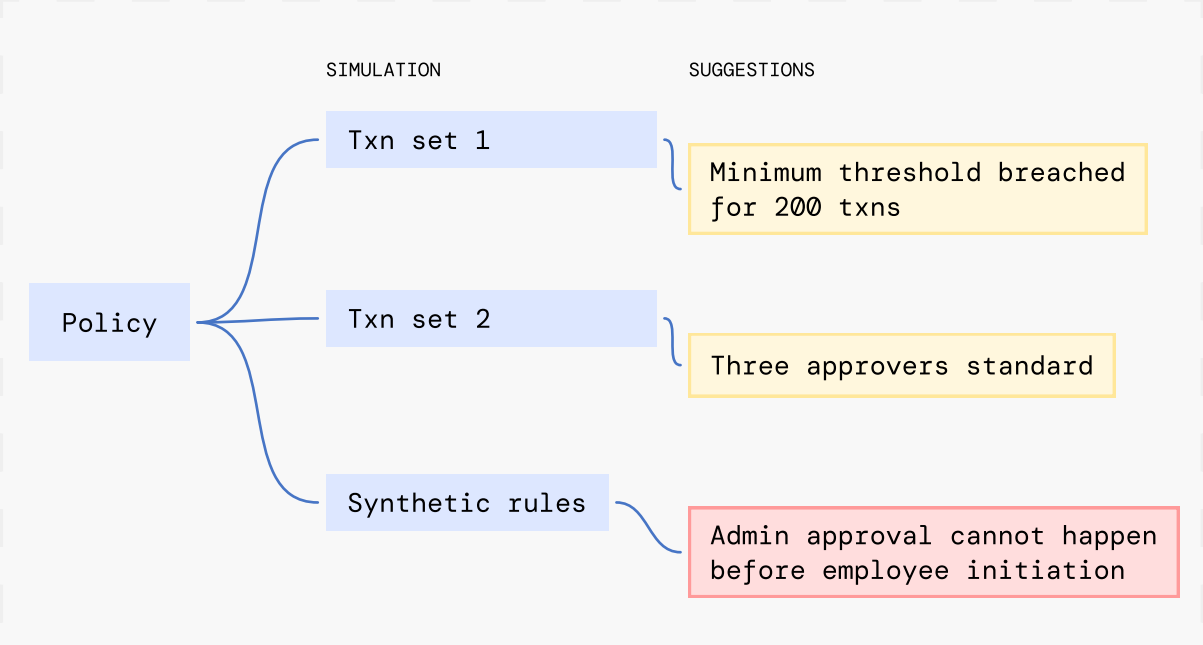

10. Reporting

Every decision the engine makes is designed to be traceable after the fact, giving administrators, auditors, and regulators visibility into what happened and why.

Immutable Logs

Every evaluation, approval, rejection, freeze, override and configuration change is recorded in an append-only, tamper-evident log with actor, timestamp and full context.

Logs may optionally be anchored on-chain.

Explainability

For every evaluated transaction, the engine returns:

- Exact rule ID
- Matched conditions
- Reason for the outcome

Rejections are understandable and auditable.

Export to External Systems

Audit data can be exported in structured formats for regulators and auditors and retained according to jurisdiction-specific record-keeping requirements (typically 5–7+ years).

Monitoring

Operational dashboards provide visibility into:

- Pending approvals
- Policy hit rates
- Denial reasons
- Approval SLA adherence

SILENCE
LABORATORIES

EMAIL

info@silencelaboratories.com

WEBSITE

silencelaboratories.com