SHANGO MID · TECHNICAL WHITEPAPER · v1.0

# Governance as Enforcement

## Cryptographically verifiable write-governance for enterprise AI agents

Architecture, attestation, and a 1,000-receipt production proof
that anyone can independently verify — server-side and offline.

**TBN attests. Shango enforces.**

Two independent layers — peers, not dependency.

**Ishaan Ghosh · Founder & CEO, Shango MID**

June 2026 · shango.in

`Every figure in this document is reproducible on demand.`

ABSTRACT

Enterprise AI agents increasingly write to systems of record — CRMs, ticketing systems, knowledge bases, financial ledgers. Most “AI governance” observes these writes after they occur. Shango MID takes a different position: governance must be **enforcement at the write-path**, applied before an action binds, and every decision must be reconstructable from cryptographic evidence. This paper describes Shango’s architecture — a write-governance layer, a hash-chained audit trail (ShangoVault), and a Bind-Point that collapses policy, execution, and attestation into a single verifiable event — and reports a production run of **1,000 governed write decisions** (400 allowed, 600 blocked). Each decision produced a real RSA-PSS-SHA256 receipt issued via TBN Protocol, independently verifiable both server-side and offline. The two systems’ records reconcile exactly at 1,000. We also state plainly what is *not* yet proven — because for a governance vendor, the discipline of only-claim-what-you-can-reproduce is the product.

SECTION 01

## The problem: ungoverned writes

An agent that can read is a convenience. An agent that can **write** is a liability surface. Once an autonomous system can modify a record of consequence, three questions become unavoidable: Was this action permitted? Can it be proven after the fact? And could a regulator, auditor, or counterparty verify that proof *without trusting the vendor that produced it*?

The prevailing answer is logging — capture what happened and review it later. But a log written by the same system that performed the action is not oversight; it is documentation. It cannot stop a bad write, and it can be edited by whoever controls the log. **Oversight without interception is just post-hoc documentation.**

SECTION 02

## Design principle: enforcement, not observation

Shango is built on one inversion: the governance decision happens **before** the write binds, and the record of that decision is signed by an independent party so it cannot be quietly rewritten. Three properties follow:

| Property | What it means |
|---|---|
| **Pre-emptive** | A blocked action never reaches the system of record. Enforcement is a gate, not a report. |
| **Attributable** | Every decision is bound to who acted, what was attempted, and which frameworks it was evaluated against. |
| **Independently verifiable** | The cryptographic receipt is checkable by anyone, against a published key, without trusting Shango’s servers. |

SECTION 03

## Architecture

Shango sits in the write-path between an agent and its system of record. A write request is evaluated by the governance engine, the decision and its evidence are chained into an append-only audit trail, and the event is attested by TBN Protocol. The boundary is deliberate: **TBN attests; Shango enforces.** Two independent companies, two layers, named as peers.

The write-path

```
agent write → Shango governance engine (allow / block) → Bind-Point (policy_hash +
execution_hash + audit_hash) → TBN attestation (RSA-PSS receipt) → ShangoVault
(append-only, SHA-256 chained)
```

**Governance engine.** Deterministic policy evaluation decides allow or block. Constitutional reasoning (an LLM-backed layer, with a deterministic keyword fallback when no model key is present) can be consulted for ambiguous cases. The decision is the gate — a blocked write does not proceed.

**Bind-Point.** The novel primitive. Rather than scatter a policy log, an execution log, and an audit log across systems, the Bind-Point collapses the policy hash, the execution hash, and the audit hash — together with TBN's attestation — into **one** verifiable event. Authority, action, and proof are a single object, not four that must be correlated later.

**ShangoVault.** An append-only audit trail with a SHA-256 hash chain: each entry commits to the previous one, so any tampering breaks the chain. The cumulative trail holds 225,574 entries with zero chain breaks (see §07 for the full, labelled composition).

SECTION 04

## Cryptographic attestation and verification

Each governed decision is attested by TBN Protocol, which issues a receipt signed with a 2048-bit **RSA-PSS-SHA256** key. Verification does not require trusting Shango or TBN at runtime — the signature can be checked against TBN's published public key.

Server-side verification

```
GET https://tbn.hardinai.co.uk/api/v1/verify/{receipt_id} → { "valid": true,
"receipt": {...} }
```

Offline verification (no network trust)

The signature covers a canonical serialization of eight fields — `receipt_id, attestation_id, agent_id, decision, decision_hash, timestamp, frameworks, version` — encoded as JSON with sorted keys and no whitespace. Anyone can fetch the public key (`/api/signing/public-key`) and verify RSA-PSS / SHA-256 / MGF1 over those bytes. A receipt is therefore replayable and attributable *even with both companies offline.*

Every receipt also records the frameworks the action was evaluated against — in production these are `["EU_AI_ACT", "ISO_42001"]` — so the compliance context travels *inside* the signed evidence,

not in a separate claim.

SECTION 05

## Production validation: 1,000 receipts

On 4 June 2026, Shango ran 1,000 governed write decisions against the live TBN production API. The run produced 1,000 unique receipts with no duplicates. A random sample of 40 receipts was independently re-verified against TBN's endpoint; all 40 returned `valid: true`. Shango's recorded count and TBN's issued count reconcile exactly.

| Metric | Result |
|---|---|
| Governed write decisions | **1,000** |
| Unique TBN receipts issued (tbn_vr_) | **1,000** |
| Allowed | 400 |
| Blocked | 600 |
| Duplicate receipts | 0 |
| SHA-256 audit-chain breaks | 0 |
| Random sample independently re-verified | **40 / 40 valid: true** |
| Reconciliation (Shango count = TBN count) | **1,000 = 1,000** |

This is a deliberately modest, fully reproducible claim. It is not a throughput benchmark: the run was sequential, and the elapsed time reflects real per-receipt signing, not system speed. The point is not speed — it is that one thousand independent governance decisions each carry a signature anyone can check.

SECTION 06

## Transparency by construction

A governance vendor's credibility *is* its product. So Shango labels every number by its pedigree — production, test, or simulation — and never blends them into a single impressive figure.

| Record set | Nature | Count |
|---|---|---|
| Production attestations (run 2026-06-04) | TBN-issued, signed | **1,000** |
| ShangoVault audit entries (cumulative) | append-only, 0 chain breaks | 225,574 |
| — of which cryptographically attested (TBN test API) | test | 182,040 |
| — of which labelled local simulation | simulation | 43,534 |

**On latency, honestly.** The governance decision is sub-millisecond and local. The cryptographic attestation is a real ~1-second RSA-PSS signature. We report both separately and never present the signing time as a hardware-speed figure — the second *is* the value, because it is a genuine signature you can verify rather than a number you must take on faith.

SECTION 07

## Standards alignment

Shango is designed against recognised governance frameworks. We describe these as *mapped to* or *designed for* — not as certifications, which would require independent audit we have not yet completed.

| Framework | Relationship |
|---|---|
| EU AI Act, Article 14 (human oversight) | Mapped / designed-for — replayable, attributable decision trails |
| ISO/IEC 42001 (AI management) | Recorded inside each production receipt's signed frameworks field |
| Constitutional Memory Architecture (arXiv:2603.04740) | An early production implementation inspired by CMA |
| Append-only, hash-chained audit | SHA-256 chain; tamper-evident by construction |

SECTION 08

## Limitations and roadmap

In keeping with the transparency principle, the current honest boundary of the system:

| Area | Status |
|---|---|
| Independent red-team | Not yet performed. Internal unit tests pass; we do not call these red-team. |
| Customer-deployment telemetry | None yet. Behavioural-drift analytics are instrumented, awaiting a design-partner pilot. |
| Throughput | Validated for correctness, not for high-concurrency scale; concurrency hardening is on the roadmap. |
| Tenant isolation & auth hardening | In progress alongside rate-limiting and key rotation. |

Stating these openly is not a weakness in a governance product — it is the demonstration of the discipline the product exists to provide.

SECTION 09

## Conclusion

Governance that cannot stop an action, or cannot be verified by an outside party, is theatre. Shango's wager is that the winning posture in enterprise AI is the opposite: enforce at the write-path, attest every decision with an independent signature, and publish a record that reconciles to the cryptography. The 1,000-receipt run is a first, deliberately verifiable proof of that posture — every figure in this paper is reproducible on demand.

> *"TBN attests. Shango enforces." Don't trust the claim — check it. Verify any receipt at* `tbn.hardinai.co.uk/api/v1/verify/{receipt_id}`

**REFERENCES & VERIFICATION**

[1] TBN Protocol public verification endpoint — tbn.hardinai.co.uk/api/v1/verify/{receipt_id}
[2] TBN signing public key — tbn.hardinai.co.uk/api/signing/public-key
[3] Constitutional Memory Architecture, arXiv:2603.04740 (2026)
[4] EU AI Act, Regulation (EU) 2024/1689, Article 14 — Human oversight
[5] ISO/IEC 42001:2023 — AI management systems
[6] Shango MID — shango.in