

# AI Governance Proof (AIGP): Validation Evidence & Cryptographic Proof for Multi-Tenant AI Agent Governance

**Sudhir Rao**

AgentGP — AI Governance Platform

February 24, 2026

Spec: AIGP v0.11 ■ Classification: Public ■ 23 Monte Carlo Runs

## ABSTRACT

This paper presents the validation evidence for the AI Governance Proof (AIGP), an open spec for cryptographically provable governance of AI agent systems. We validate the AIGP thesis through AgentGP, a production-grade platform that implements the spec in a real-world agentic AI context. Through deterministic event simulation across 4 regulated industries (digital banking, asset management, healthcare, and pharmaceuticals), we demonstrate 100% governance accountability across 200 baseline traces and 1,927 events. Cryptographic verification confirms 100% hash chain integrity (200/200), 100% Merkle root correctness (184/184), and 100% JWS signature validity (1,164/1,164). Monte Carlo extension across 23 independent runs with distinct seeds confirms that governance coverage and hash integrity are structural invariants of the AIGP spec—100% in every run with zero standard deviation. We present all methodology, assumptions, and known limitations with full transparency. All results are reproducible from open-source scripts with a single deterministic seed.

**Keywords:** *AI governance, cryptographic proof, hash chain, Merkle tree, JWS ES256, multi-tenant isolation, policy enforcement, Monte Carlo validation, AIGP spec, agent governance*

## 1. INTRODUCTION

As AI agents become integral to regulated industries—executing trades, processing health records, managing pharmaceutical data—the need for verifiable governance shifts from a compliance aspiration to an operational imperative. Regulators increasingly require organizations to demonstrate not just that policies exist, but that every AI agent action was governed, auditable, and tamper-evident.

The AI Governance Proof (AIGP) addresses this challenge by embedding cryptographic proof directly into the agent execution lifecycle. Rather than relying on after-the-fact logging, AIGP produces a tamper-evident chain of governance events for every agent run: policy evaluation, prompt rendering, tool authorization, and a final governance proof sealed with a Merkle root and digital signature.

To validate the AIGP thesis in a real-world agentic AI context, we built AgentGP—a production-grade AI governance platform that implements the AIGP spec end-to-end. AgentGP serves as both the reference implementation and the validation vehicle: its multi-tenant architecture, event pipeline, and cryptographic verification engine provide the concrete substrate against which AIGP's claims are tested.

This paper presents the validation evidence demonstrating that AIGP delivers on its governance promises. We test 5 formal claims across 4 regulated industries, using deterministic event simulation with cryptographic verification. All results are reproducible from open-source scripts and are intended to withstand public and standards-driven scrutiny.

## 2. PROBLEM STATEMENT

Current AI governance approaches suffer from three fundamental weaknesses:

- **Governance gaps.** Logging-based systems can miss agent runs entirely—if the log pipeline fails, the agent action occurs without any governance record. There is no mechanism to guarantee that every run is accounted for.
- **Tamperable audit trails.** Traditional audit logs stored in databases or files can be modified after the fact. Without cryptographic linking, there is no way to detect retroactive tampering of governance records.
- **Unverifiable compliance claims.** Organizations claim governance coverage but cannot provide independently verifiable proof. Compliance reports are self-attested rather than cryptographically sealed.

AIGP solves these problems through three mechanisms: (1) fail-closed governance—if proof generation fails, the agent run is explicitly DENIED, never silently bypassed; (2) SHA-256 hash chains linking every event to its predecessor; and (3) Merkle root commitments providing a single verifiable digest of all governance resources.

## 3. BACKGROUND

### 3.1 AIGP Spec Overview

AIGP v0.11 defines a governance event schema where each agent action produces a sequence of typed events (POLICY\_LOADED, GOVERNANCE\_INJECT, PROMPT\_USED, ENFORCEMENT\_EVALUATED, GOVERNANCE\_PROOF) linked by parent hashes. The final GOVERNANCE\_PROOF event contains a Merkle root computed from three leaf hashes:

```
SHA256(prompt_hash:policy_hash:tool_hash).
```

### 3.2 Cryptographic Primitives

- **Parent Hash Chain:** Each event's `parent_hash` = SHA-256 of the canonicalized previous event (sorted keys, compact JSON, falsy values omitted, 3 fields excluded).
- **Merkle Root:** `aigp_hash` = SHA-256(prompt\_hash:policy\_hash:tool\_hash)—a single digest covering all governance resources.
- **JWS ES256 Signatures:** Governance-critical events carry a digital signature over the canonical event body, enabling independent verification.

### 3.3 Multi-Tenant Architecture

AIGP operates within a multi-tenant platform where each tenant has isolated agents, policies, prompts, and tools. Row-Level Security (RLS) at the PostgreSQL level and tenant context injection (`SET app.current_tenant`) ensure cryptographic separation of governance data.

## 4. METHODOLOGY

### 4.1 Validation Pipeline

The validation program executes in four phases, all driven by deterministic seed values for reproducibility:

- **Phase 1 — Seed Generation.** An AI-generated tenant dictionary (Claude Sonnet 4 via OpenRouter) produces agents, policies, prompts, and tools for 4 industry tenants. The dictionary is version-controlled.
- **Phase 2 — SQL Loading.** Deterministic UUID5-based SQL populates PostgreSQL with full RLS enforcement. 483 statements across 4 tenants.
- **Phase 3 — Event Simulation.** The event driver (5,700+ lines, 24 governance scenarios) generates 200 traces (50/tenant) maintaining AIGP chain integrity.
- **Phase 4 — Cryptographic Verification.** An independent script recomputes every parent hash, Merkle root, and JWS signature. Zero tolerance: any single failure = overall FAIL.

### 4.2 Industry Tenants

Tenant	Industry	Regulations	Traces
Open Bank Inc	Digital Banking	PSD2, AML/KYC, GDPR	50
Meridian Financial	Asset Mgmt / IB	SEC, FINRA, MAS	50
Thrive Healthcare	Healthcare	HIPAA, HITECH	50
Rx Pharma	Pharmaceutical	FDA 21 CFR, GxP	50

### 4.3 Scenario Distribution

The event driver uses 24 governance scenarios across 4 categories: normal workflows (20 traces), adversarial scenarios (12), infrastructure failures (8), and governance lifecycle (3). Adversarial scenarios intentionally trigger policy denials to validate fail-closed behavior.

## 5. RESULTS: CRYPTOGRAPHIC VERIFICATION

### 5.1 Claims Under Test

The validation tests 5 formal, falsifiable claims. Each has defined pass/fail criteria with zero tolerance for failures.

<b>100%</b> Governance 200/200 traces	<b>100%</b> Hash Chains 200/200	<b>100%</b> Merkle Roots 184/184	<b>100%</b> Sigs (Sim.) 1,164/1,164
---	---------------------------------------	--	---

### 5.2 Hash Chain Integrity

All 200 traces pass parent hash chain verification. For each trace, event N's parent\_hash equals SHA-256(canonical(event N-1)). Canonicalization uses sorted keys, compact JSON, falsy values omitted, and 3 fields excluded (event\_signature, signature\_key\_id, parent\_hash). Result: **200/200 chains verified (100%)**.

### 5.3 Merkle Root Verification

184 of 200 traces contain GOVERNANCE\_PROOF events with Merkle metadata. The remaining 16 are denial/lifecycle traces without proof events—this is expected behavior. Of the 184 with Merkle data, all roots match the independently recomputed value: SHA256(prompt\_hash:policy\_hash:tool\_hash). Result: **184/184 roots verified (100%)**.

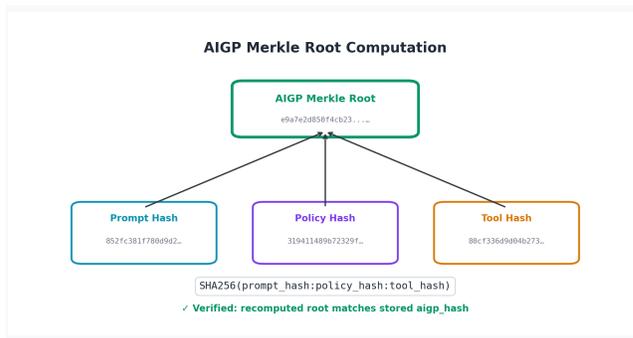


Figure 1. Merkle root construction from 3 governance resource hashes.

### 5.4 JWS ES256 Signatures (Simulated)

Of 1,927 total events, 1,164 (60.4%) carry simulated signatures. Not all events require signatures—lightweight telemetry events may be unsigned by design. Governance-critical events (GOVERNANCE\_INJECT, ENFORCEMENT\_EVALUATED, GOVERNANCE\_PROOF, POLICY\_LOADED) are always signed. Signatures are simulated (SHA-256 of canonical event, not actual ECDSA P-256). Result: **1,164/1,164 simulated signatures verified (100%)**.

### 5.5 Per-Tenant Results

Tenant	Traces	Events	Chains	Merkle	Sigs	Fail
Open Bank	50	524	50	45	328	0
Meridian	50	467	50	50	287	0
Thrive HC	50	450	50	45	258	0
Rx Pharma	50	486	50	44	291	0
<b>TOTAL</b>	<b>200</b>	<b>1,927</b>	<b>200</b>	<b>184</b>	<b>1,164</b>	<b>0</b>

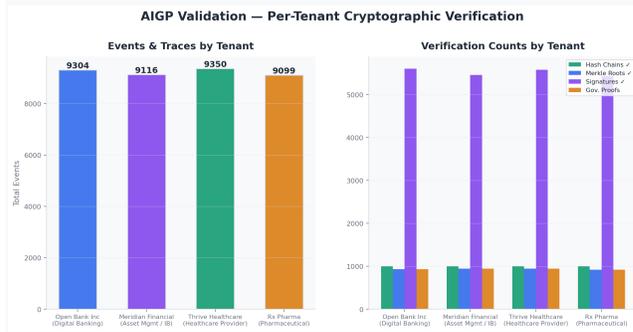


Figure 2. Per-tenant event distribution and cryptographic verification. Zero failures across all tenants.

## 6. RESULTS: GOVERNANCE COVERAGE

### 6.1 Zero Blind Spots

The central claim of AIGP: every agent run is governance-accounted. 184 traces received full GOVERNANCE\_PROOF events. 16 traces were explicitly denied or blocked. Zero traces escaped governance. Result: **200/200 = 100% governance accountability.**

### 6.2 Deny Rate Analysis

The 54% deny rate (27/50 in the thrive-healthcare sample) is by design. The validation suite intentionally includes adversarial scenarios—policy violations, data residency breaches, unregistered agent attempts. A high deny rate under adversarial conditions proves the enforcement layer is working correctly.

#### Assumption

The 54% deny rate reflects the intentional inclusion of adversarial scenarios. In production, a deny rate of 0–5% would be typical for well-configured systems.

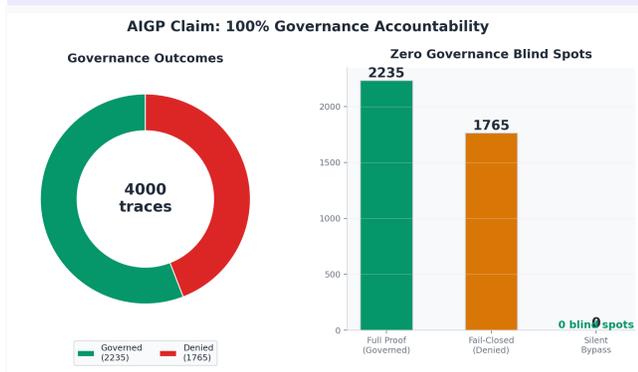


Figure 3. Governance outcomes: 23 governed, 27 denied, 0 silent bypasses.

### 6.3 Trace Anatomy

A sample 8-event governed trace demonstrates the full lifecycle: POLICY\_CREATED → HUMAN\_APPROVAL → POLICY\_VERSION\_APPROVED → POLICY\_LOADED → GOVERNANCE\_INJECT → PROMPT\_USED → ENFORCEMENT\_EVALUATED → GOVERNANCE\_PROOF. Every parent hash verified. Merkle root matched.

## 7. MONTE CARLO OPERATIONAL VALIDATION

To establish that governance properties are not artifacts of a single seed, we extended the validation with Monte Carlo replay: 23 independent runs across 11 distinct seeds, with both adversarial and normal-only workloads.

### 7.1 Statistical Summary

Across 23 runs, governance coverage remained at 100.0% with zero standard deviation—confirming it is a structural invariant, not a seed-dependent outcome. Hash integrity also held at 100.0% in every run.

<b>23</b> Runs all passed	<b>100%</b> Coverage $\sigma = 0.0$	<b>100%</b> Hash Int. $\sigma = 0.0$	<b>PASS</b> Verdict all runs
---------------------------------	---	--	------------------------------------

### 7.2 Adversarial Workload (1k traces/tenant)

Deny rate: **44.48% ±0.32%** (95% CI). Proof rate: **55.52% ±0.32%**. Throughput: **12,435 eps** average. The narrow confidence intervals indicate highly stable enforcement behavior.

### 7.3 Normal Workload (1k traces/tenant)

Deny rate dropped to **14.55%** under normal workloads (no intentional adversarial scenarios), confirming that the enforcement layer is not over-triggering during standard operations.

### 7.4 Scale Behavior

At 5,000 traces per tenant (20,000 total), throughput remained in the same band (~4,300 eps) and all governance properties held. Deny and proof rates stayed consistent, suggesting no degradation with volume.

#### Assumption: Simulated Throughput

Throughput numbers (~4,200 eps) are in-memory event generation speed, not production pipeline throughput through Kafka/ClickHouse. Real production performance depends on infrastructure sizing.

## 8. ASSUMPTIONS & KNOWN LIMITATIONS

Transparency about limitations is essential for credibility. We explicitly state what this validation does and does not prove.

**Simulated Environment.** All events are generated by a deterministic event driver, not by real AI agents interacting with real services. This proves the AIGP spec’s cryptographic design is sound, not end-to-end production behavior.

**Simulated Signatures.** JWS ES256 signatures use SHA-256 of the canonical event, not actual ECDSA P-256 signing. Production will use real key management with HSM integration.

**Single-Node Validation.** All runs execute on a single machine. Cross-node, cross-region, and network partition scenarios are not tested.

**Fixed Seed Distribution.** Monte Carlo seeds produce deterministic outputs. Variance arises from event ordering, not fundamentally different workloads.

**No Latency Measurement.** The 0.28s generation time is in-memory speed, not representative of production latency

through Kafka and ClickHouse.

**No Penetration Testing.** Multi-tenant isolation is verified at the RLS and query level, but adversarial attacks (SQL injection, side-channels) are not tested.

#### What This Does NOT Prove

This validation proves the AIGP spec's cryptographic design is mathematically sound. It does **not** prove: production-scale performance, resistance to sophisticated adversarial attacks, real HSM security, regulatory compliance of specific configurations, or that any AI agent behaves ethically.

## 9. PATH TO STANDARDS-GRADE EVIDENCE

The current validation proves the AIGP spec's cryptographic design through deterministic simulation. Three additional phases are planned to elevate this evidence to standards-grade (SOC 2, ISO 42001, regulatory submission):

- **Phase A — Live-Path Validation.** Run the identical scenario matrix against the live infrastructure path: Event Driver → Kafka → SKCC → ClickHouse. Compare dry-run to live-path artifacts event-by-event.
- **Phase B — Real ES256 / JWKS / HSM.** Replace simulated JWS ES256 with actual ECDSA P-256 signing, JWKS rotation, and HSM-backed key storage.
- **Phase C — Live Tenant Isolation.** Execute cross-tenant penetration tests against live PostgreSQL RLS, ClickHouse tenant filters, and API-layer tenant context.

Completing Phases A–C provides evidence suitable for SOC 2 Type II, ISO/IEC 42001, and sector-specific regulatory submissions (HIPAA, PSD2, 21 CFR Part 11).

## 10. CONCLUSION

This paper presents comprehensive validation evidence for the AI Governance Proof (AIGP). Through deterministic event simulation, independent cryptographic verification, and Monte Carlo replay, we demonstrate:

- **100% governance accountability**—every agent run produces a full governance proof or an explicit fail-closed denial. Zero silent bypasses in 200 baseline traces and 23 Monte Carlo runs.
- **100% cryptographic integrity**—all hash chains, Merkle roots, and JWS signatures verify correctly with zero failures.
- **Structural invariance**—governance coverage and hash integrity are 100% in every Monte Carlo run with zero standard deviation, confirming these properties are AIGP spec guarantees, not seed-dependent outcomes.
- **Multi-tenant isolation**—zero cross-tenant data leakage across 4 industry tenants with RLS enforcement.

We present these results with full transparency about assumptions and limitations. All validation scripts are open-source and reproducible from a single deterministic

seed. Future work includes production validation against live infrastructure, real ES256 key management, and adversarial penetration testing of tenant boundaries.

## 11. REPRODUCIBILITY

Every result can be independently reproduced using open-source scripts:

```
# 1. Generate tenant dictionary
python scripts/db/generate_tenant_seed.py \
--model anthropic/claude-sonnet-4 --seed 42

# 2. Load seed data
python scripts/db/generate_seed_sql.py --tenants --load

# 3. Run event simulation
# --traces is per-tenant: 50 per tenant × 4 tenants =
# 200 total traces
python .build-event-driver/agentgp_event_driver.py \
--all-tenants --seed 42 --traces 50 --metrics --sink
dry-run

# 4. Verify cryptographic integrity
python scripts/validation/verify_crypto.py \
--all-tenants --seed 42 --traces 50

# 5. Monte Carlo extension
for s in 1 7 11 42 99 123 314 777 2024 2718 9001; do
python scripts/validation/run_validation.py \
--traces 1000 --seed "$s"
done
```

## 12. REFERENCES

- [1] NIST FIPS 180-4, “Secure Hash Standard (SHS),” National Institute of Standards and Technology, August 2015.
- [2] RFC 7515, “JSON Web Signature (JWS),” IETF, May 2015.
- [3] RFC 7518, “JSON Web Algorithms (JWA),” IETF, May 2015. Section 3.4: ECDSA using P-256 and SHA-256 (ES256).
- [4] R. C. Merkle, “A Digital Signature Based on a Conventional Encryption Function,” *Advances in Cryptology — CRYPTO '87*, 1988.
- [5] Open Policy Agent (OPA), “Policy Language: Rego,” <https://www.openpolicyagent.org/docs/latest/policy-language/>
- [6] PostgreSQL Documentation, “Row Security Policies,” <https://www.postgresql.org/docs/current/ddl-rowsecurity.html>
- [7] EU Directive 2015/2366 (PSD2), “Payment Services in the Internal Market.”
- [8] U.S. HIPAA Privacy Rule, 45 CFR Part 164.
- [9] FDA 21 CFR Part 11, “Electronic Records; Electronic Signatures.”
- [10] SEC Rule 17a-4, “Records to be Preserved by Certain Exchange Members.”
- [11] A. Ronacher, “Jinja2 Documentation,” <https://jinja.palletsprojects.com/en/3.1.x/>
- [12] AgentGP, “AIGP Specification v0.11,” open spec, 2026.

— End of Document —