Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 18, 2026

K. Singh
Apex Intelligence Empire

March 17, 2026

Proof of Sovereign Integrity (PSI): A Cryptographic Protocol for Verifiable AI Regulatory Compliance

draft-singh-psi-00

Abstract

This document specifies the Proof of Sovereign Integrity (PSI) Protocol, version 1.2, a cryptographic framework enabling organizations to prove compliance with AI regulations (including the EU AI Act 2024/1689, NIST AI RMF, UK AI Safety Institute guidelines, and equivalent frameworks) without disclosing proprietary model architectures, training data, or inference logic.

PSI achieves this through a combination of SHA-256 hash-chained audit trails, Ed25519 digital signatures, Merkle inclusion proofs, Groth16-compatible zero-knowledge commitments over BN128 fields, and a 3-node Multi-Party Computation (MPC) consensus mechanism with 2/3 threshold verification.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on September 18, 2026.

Copyright Notice

Table of Contents

1. Introduction 2
2. Terminology 3
3. Protocol Overview 3
4. Cryptographic Primitives 4
5. Verification Pipeline 5
6. Deterministic Pre-Flight 6
7. Merkle Tree Construction 7

8. MPC Consensus Layer 7
9. Zero-Knowledge Commitments 8
10. Sovereign Tribunal 9
11. Predicate Registry 9
12. Proof Bundle Format 10
13. Legal-to-Technical Mapping 11
14. Security Considerations 11
15. IANA Considerations 12
16. Orbital Integrity Protocol (OIP) 13
17. References 14
Authors' Addresses 15

1. Introduction

The proliferation of artificial intelligence systems across critical sectors has created an urgent need for verifiable compliance mechanisms. The EU AI Act mandates technical conformity assessment for high-risk AI systems. Existing compliance approaches suffer from IP exposure risk, non-verifiability, and reliance on trust.

The PSI Protocol addresses these limitations through cryptographic verification primitives that enable mathematical proof of compliance without disclosing protected intellectual property. PSI v1.2 introduces Deterministic Mode (blocking non-compliant actions before commit) and the Sovereign Tribunal (human ratification).

2. Terminology

Commit: An atomic action submitted for compliance verification.

Predicate: A machine-readable regulatory requirement (e.g. EU AI Act Article 14).

Commit Hash: SHA-256(JCS(action || predicate_id || timestamp)).

Merkle Root: The root hash of the binary tree containing ledger entries.

MPC Node: One of three independent verification nodes.

Proof Bundle: A JSON document containing all cryptographic artifacts for verification.

Sovereign Tribunal: A panel of 5 independent auditors providing human ratification.

PIL: Protocol Intervention Layer (runtime execution gate).

3. Protocol Overview

The PSI Protocol operates as a 4-stage pipeline:

Stage 1 - COMMIT: Input action and predicate. Run deterministic pre-flight check. Compute commit_hash and merkle_leaf.

Stage 2 - CHALLENGE: Evaluate commit against predicate violation patterns. Compute challenge_hash and record status.

Stage 3 - PROVE: Generate Merkle inclusion proof. Generate Groth16-compatible ZK commitment over BN128 fields.

Stage 4 - VERIFY: 3 MPC nodes independently verify the proof. 2/3 consensus required. Sign result with Ed25519.

4. Cryptographic Primitives

4.1. Hash Function

PSI uses SHA-256. Input MUST be canonicalized using RFC 8785 (JCS) before hashing.

4.2. Digital Signatures

Ed25519 (RFC 8032) is used for non-repudiation of Merkle roots and auditor verdicts.

4.3. Merkle Trees

Binary Merkle trees are constructed from leaf hashes. Deterministic ordering and leaf duplication for odd counts are required.

4.4. Zero-Knowledge Commitments

ZK commitments use BN128 finite field arithmetic. Proof elements (pi_A, pi_B, pi_C) demonstrate knowledge of action satisfying the predicate without revealing the action.

4.5. Sequence Counter

A monotonic sequence counter with gap detection prevents omission attacks on the ledger.

5. Verification Pipeline

1. Commit: Generate ID, JCS canonical hash, and persist.
2. Challenge: Match against violation patterns; assign status.
3. Prove: Generate Merkle proof path and ZK commitment.
4. Verify: MPC nodes audit integrity and sign root hash.

6. Deterministic Pre-Flight

The PIL evaluates actions BEFORE commit. If action matches HIGH or UNACCEPTABLE risk patterns, the PIL blocks execution and prevents ledger entry.

7. Merkle Tree Construction

Trees are computed by recursively hashing lexicographically sorted sibling pairs. ProofPath includes sibling hashes and positions.

8. MPC Consensus Layer

Verification is distributed across Alpha, Beta, and Gamma nodes. Consensus ensures that no single entity can forge a compliance certificate.

9. Zero-Knowledge Commitments

BN128 field operations provide the privacy layer. The system uses Groth16 elements to prove predicates without data disclosure.

10. Sovereign Tribunal

Human oversight is provided by 5 auditors. 3/5 threshold is required for final ratification of automated verdicts.

11. Predicate Registry

Contains machine-readable rules for EU AI Act, MiFID II, DORA, NIST AI RMF, and AU Privacy Act.

12. Proof Bundle Format

Standardized JSON including merkle_root, ed25519_signature, and zk_proof points.

13. Legal-to-Technical Mapping

Article 12: Immutable logging + RFC 8785.
Article 14: Human Oversight PIL + Sovereign Tribunal.
Article 15: MPC consensus + ZK-Integrity.

14. Security Considerations

Protocol mitigates log tampering, false-negatives, and single points of failure. IP protection is guaranteed by ZK mathematics.

15. IANA Considerations

Requests registration of application/psi-proof+json and the psi:// URI scheme.

16. Orbital Integrity Protocol (OIP)

Extends PSI to space-based compute. Defines LAM levels for radiation-tolerant proof generation on satellite hardware.

17. References

RFC 2119, RFC 8032, RFC 8785, EU 2024/1689.

Authors' Addresses

Kawaljeet Singh
Apex Intelligence Empire
Balaclava, Victoria, Australia
Email: kawaljeet.singh3008@gmail.com
URI: https://apex-infrastructure.com