

OPEN ARTIFICIAL INTELLIGENCE INTEGRITY SENTINEL

OpenAII

A Decentralized Trust, Governance and Safety Protocol for the Digital World

One universal protocol to establish provenance, govern digital trust, preserve accountability and enable verifiable safety across the Internet.

Tarun Sethi

White Paper | 25 October 2026

V4.0

Abstract

The Internet has made the creation, distribution and execution of digital information nearly universal, but it has not established an equally universal mechanism for verifying its origin, governing its authority, preserving responsibility or responding to threats against its integrity.

Digital signatures can authenticate claims, transparency systems can preserve records, identity frameworks can provide credentials, and security systems can report compromise. These mechanisms are valuable, yet the evidence they produce often remains separated across platforms, organizations and trust domains. As artificial intelligence and autonomous software increasingly create content and execute actions, the consequences of this separation become more significant.

We propose OpenAIIIS, a decentralized protocol for establishing digital provenance, governing trust, preserving accountability and enabling verifiable safety across the Internet.

OpenAIIIS combines cryptographically signed claims, typed attestations, verifiable consent, delegated authority, accountable responsibility, lifecycle governance and authenticated security evidence within a common, independently verifiable history.

A claim identifies what was asserted, by whom, about which exact digital subject and with what supporting evidence. Subsequent records preserve changes in authority, consent, responsibility, disputes, revocations and security-relevant events. Independent verifiers evaluate these records against their own trusted authorities and policies.

An append-only cryptographic history makes records auditable, while independently operated infrastructure is intended to remove dependence on any single organization. Optional Bitcoin anchoring can provide an additional historical commitment to authenticated registry checkpoints.

The result is a common foundation through which creators can preserve evidence of their work, institutions can govern digital authority, autonomous systems can demonstrate accountability, and security systems can exchange verifiable information about emerging threats.

OpenAIIIS proposes one universal protocol to establish provenance, govern digital trust, preserve accountability and enable verifiable safety across the Internet.

1. Introduction

The Internet was designed to exchange information between independent computers. It succeeded in making information globally accessible, but it left the interpretation of identity, authority, responsibility and trust largely to individual services.

A document can be copied without preserving its origin. An invention can be disclosed without producing independently verifiable evidence of its prior existence. A software release can carry a valid signature without demonstrating that the signer was authorized to publish it. An autonomous agent can execute an action without leaving a portable, independently verifiable chain identifying the organization and individual responsible for its authority.

The same limitations affect security.

A compromised signing key may continue to appear cryptographically valid to systems that have not received evidence of its compromise. A malicious software release may be distributed before independent

verifiers learn that its authority has been withdrawn. A security incident reported by one organization may have no authenticated, interoperable path into another organization's verification process.

In each case, information exists, but the evidence needed to establish trust is incomplete, fragmented or controlled by separate intermediaries.

The growth of artificial intelligence magnifies this problem. Systems increasingly produce software, research, media and decisions at a speed that makes manual verification impractical. Autonomous agents may act on behalf of organizations, operate other software, publish artifacts and interact with external systems.

A digital signature answers whether particular bytes were signed by a key. It does not, by itself, answer whether that key was permitted to authorize the action, whether its authority remained valid, or whether new evidence should prevent the action from proceeding.

The problem is therefore larger than verifying content.

The Internet requires a common mechanism for establishing what was claimed, who possessed authority, who accepted responsibility, what evidence supported the claim, and how subsequent events change its trustworthiness.

OpenAIIS proposes such a mechanism.

The name OpenAIIS stands for Open Artificial Intelligence Integrity Sentinel. Although artificial intelligence is a major motivation for its development, the protocol applies to digital activity generally, whether performed by a person, institution, software system or autonomous agent.

2. The OpenAIIS Principle

Consider a digital artifact published by an individual or organization.

The artifact may be a research paper, source-code release, invention disclosure, photograph, AI-generated work or signed declaration. Its publication raises several independent questions:

- Provenance: What exact artifact was declared, and who made the declaration?
- Governance: Which authority permitted the declaration, publication, transfer or action?
- Responsibility: Who accepted accountability for the relevant activity?
- Safety: Has any authenticated event subsequently affected the trustworthiness or permitted use of the artifact, actor or authority?
- Verification: Can an independent party establish these facts without relying exclusively on the original platform?

These questions are related, but none can replace the others.

An artifact may have valid provenance but invalid authorization. An authorized agent may publish content that is later identified as compromised. An authentic security report may be issued by an organization that a particular verifier does not recognize as authoritative. A legitimate claim may later be disputed without its historical existence becoming invalid.

OpenAIIS treats these as related parts of one verifiable system.

A signed claim establishes a cryptographic assertion. Governance records establish relevant authority and consent. Responsibility records bind accountable actors to defined roles and actions. Safety records communicate authenticated changes that can affect trust. A shared cryptographic history

preserves the sequence of these events.

Together, they allow an independent verifier to reconstruct a trust decision from evidence rather than accepting an unexplained declaration from a single intermediary.

This creates an important distinction:

The history of a claim can be immutable while its current trust state remains subject to legitimate change.

An invention claim can remain permanently recorded while a later dispute becomes visible. A signing key can be historically authentic while its future authority is revoked. An agent can have been authorized for an earlier release while being prohibited from publishing another.

The history remains available. The governing interpretation evolves through additional authenticated records.

This is the foundation of OpenAIIS.

3. Universal Signed Claims

The fundamental unit of OpenAIIS is a cryptographically signed, typed digital claim.

A claim binds a defined assertion to its issuer, subject, content commitment and supporting references.

For an artifact A, its content commitment may be expressed as:

$$h = \text{SHA256}(A)$$

A canonical claim body C contains the record type, issuer identity, subject identity, content digest, typed payload digest, declared time and references to relevant authority, consent, evidence or earlier records.

The issuer signs a domain-separated canonical representation of this body:

$$S = \text{Sign}(\text{issuer_private_key}, \text{domain} \parallel \text{Canonical}(C))$$

A verifier possessing an independently trusted issuer public key can verify that signature and detect modification of the signed bytes.

An OpenAIIS claim can represent an original work, an invention assertion, research, software, generated content, media, identity evidence, permission, responsibility or a lifecycle event.

These records use common structural principles while preserving their different meanings.

A claim that a document existed is distinct from an attestation that its author possesses a particular identity. An identity attestation is distinct from permission to execute software. A declaration of responsibility is distinct from proof that an action occurred.

The protocol therefore uses typed claims with explicit validation and authorization rules, rather than interpreting every valid signature as equivalent authority.

The initial claim vocabulary spans seventeen categories covering creative and technical artifacts, identity and artifact attestations, consent, authority, supporting evidence, revocation, withdrawal, disputes and artifact-attestation transfers.

Further claim types can extend the system through versioned definitions and independently verifiable semantics.

3.1. Linked evidence

A claim may refer to other signed claims, committed evidence or original content held outside the registry.

For example, a research claim may identify an exact manuscript and separately signed supporting observations. A software release may refer to an organizational authority record, a responsibility declaration and a release-evidence attestation.

The common registry preserves the cryptographic relationships between these records.

Original confidential material need not be published as part of every registration. A verifier with legitimate access to the original bytes can calculate their digest and compare it to the recorded commitment.

This creates a portable evidence structure in which a record can be verified independently of the service through which it was originally created.

4. Governance of Digital Trust

Trust is not a permanent characteristic of a signature. It depends on the identity and authority of its issuer, the scope of the claim, relevant consent and the lifecycle of the underlying relationships.

OpenAIIIS makes these relationships explicit and verifiable.

A governance record may describe an issuer's authority to make a particular attestation, an organization's delegation to an agent, a subject's consent to an identity claim or the withdrawal of previously granted permission.

The protocol distinguishes among actors rather than treating them as interchangeable.

An issuer signs an assertion. A subject is the entity or artifact about which the assertion is made. A principal may accept responsibility or delegate authority. An operator may control an agent. An evidence issuer may attest to supporting material. A relying party decides whether these records satisfy its own trust requirements.

4.1. Authority and consent

An organization may authorize one actor for a defined action without granting universal authority over its identity or assets.

The authorization can be restricted by subject, scope, permitted operation, duration and associated conditions.

Where a person's consent is required, the consent must be independently authenticated and bound to the relevant intent. Merely naming a person or referencing an unsigned consent statement is insufficient.

This separation enables differentiated trust.

A university may be trusted to attest to a researcher's institutional association without being authorized to transfer that researcher's intellectual property. A company may authorize a software agent to publish a release without authorizing it to delegate further permissions.

The validity of one relationship does not imply another.

4.2. Revocation, withdrawal and disputes

Authority can change after a record has been published.

OpenAIIS expresses such changes through additional signed, linked events.

An issuer may revoke its own attestation. A subject may withdraw consent under the applicable rules. An authorized party may raise a dispute. A dispute may subsequently receive an authenticated resolution.

These operations have different actors, meanings and consequences. A subject's withdrawal is not automatically the issuer's revocation, and a dispute is not automatically proof that the disputed claim was false.

By retaining the original records and their subsequent lifecycle, OpenAIIS permits independent verification of both historical assertions and current governance state.

The result is a continuously auditable trust history rather than a collection of isolated certificates.

5. Responsibility and Accountability

The delegation of digital authority introduces a problem that becomes especially important with autonomous systems: an action may be performed by software while its consequences belong to a human or organization.

A system can authenticate an agent's signature without establishing who authorized its behavior.

OpenAIIS addresses this by linking identity, operational control, accepted responsibility, scoped authorization, executed claims and supporting evidence.

Consider an organization called Company X.

Company X operates an autonomous software agent that prepares and publishes application releases.

A verifiable authorization chain can contain the following relationships:

1. Company X identifies the operating authority responsible for the agent.
2. The operator signs a declaration binding the agent to its operational control.
3. An accountable principal signs an explicit acceptance of responsibility.
4. The principal delegates a defined publication authority to the agent.
5. The agent signs a release claim committing to the exact software artifact.
6. An authorized evidence issuer signs evidence associated with that release.
7. The release and its authority/evidence references become independently verifiable through the common registry history.

An external verifier can then examine the complete relationship.

Was the release signed by the expected agent? Was the agent bound to the declared operator? Did a trusted principal accept responsibility? Was a valid delegation in force for this action? Does the exact software match the committed digest? Does the supporting evidence satisfy the applicable requirements?

This chain is more useful than a standalone statement that the software was signed successfully.

It establishes a verifiable basis for deciding whether the action was appropriately authorized and supported.

The relying party may consequently return one of three policy decisions:

ALLOW: The applicable authority, integrity, lifecycle and evidence requirements have been satisfied.

DENY: A required condition has failed, including a cryptographic failure or a known prohibition.

REVIEW: Available evidence is incomplete, contradictory, stale or insufficient for a safe automated decision.

The final decision follows the relying party's policy. The underlying proof remains independently inspectable.

The same model extends beyond AI.

It can apply to software maintainers, organizational delegates, research institutions, publishing authorities, automated infrastructure, robotic systems and other environments where the person or system performing an action differs from the authority responsible for it.

OpenAIIS makes delegated power traceable to authenticated authority and accepted responsibility.

6. Verifiable Safety and Security Response

Digital trust must remain responsive to evidence of compromise.

A historical signature can remain mathematically valid after a private key has been stolen. A software artifact can remain byte-for-byte identical after it has been identified as malicious. An agent can continue presenting an earlier authorization after that authorization has been revoked.

Integrity verification alone cannot resolve these conditions.

OpenAIIS connects safety-relevant evidence to the same authenticated claim and governance history used to establish provenance and authority.

This enables security information to participate directly in trust evaluation.

6.1. Authenticated security evidence

The safety model permits authorized actors to issue signed records concerning compromise, withdrawn authority, contested artifacts, affected releases and supporting security evidence.

Existing governance and evidence record types provide the basis for revocation, dispute and reassessment. Dedicated incident-advisory and compromise-reporting profiles can extend the vocabulary with explicit affected subjects, incident categories, issuer authority, observation windows and recovery relationships.

Such records must identify their issuer and precise subject, include verifiable supporting references where required, and remain subject to the same authorization and lifecycle checks as other claims.

The purpose is to make security-relevant assertions portable, attributable and independently inspectable.

6.2. Network-assisted threat response

Suppose Company X discovers that an agent's signing key has been compromised.

An attacker has used the key to produce a malicious software release that appears to carry a valid signature.

Under a conventional signature-only check, a consumer may accept the release because the signature verifies.

Under the OpenAIIS safety model, Company X can publish an authenticated key-compromise or authority-revocation record linked to the affected identity or authorization. A separately authorized security issuer may provide an incident advisory and supporting evidence identifying affected releases.

As these records become available through the shared verifiable history, participating verification systems can evaluate the new information.

A deployment system may deny further releases signed under the revoked authority. Another organization may flag an affected artifact for investigation. A software distributor may temporarily suspend publication. A security monitoring service may correlate the affected records with additional trusted incident evidence.

Each response is determined by that system's own configured trust roots and policies.

The common infrastructure allows independently operated systems to examine the same cryptographic evidence while maintaining their own responsibility for decisions.

6.3. Protection against malicious reporting

Security reports themselves can be instruments of attack.

An adversary may attempt to publish false compromise allegations, impersonate a security authority or suppress a competitor's software through fabricated incidents.

OpenAIIS therefore requires safety-sensitive assertions to retain explicit issuer identity, independently evaluated authority, precise scope, supporting evidence and a verifiable lifecycle.

An incident report from an unknown signer must not automatically override an organization's trusted authorization records. Conflicting evidence must remain visible. Revocation, correction, recovery and dispute resolution require authorized transitions rather than silent modification of earlier history.

This creates the basis for an independently governed, globally shareable safety-evidence network.

Its strength comes from combining authenticated reporting with transparent accountability, not from granting any single actor unquestionable authority to declare an artifact safe or unsafe.

7. A Shared, Verifiable History

For independent verification to be meaningful, signed claims and governance events must be bound to a history whose integrity can be checked.

OpenAIIS uses an append-only record model with cryptographic commitments.

Let the ordered registry records be:

$R_1, R_2, \dots, R_n$

A Merkle tree combines their cryptographic commitments into a root:

$M_n = \text{MerkleRoot}(R_1, R_2, \dots, R_n)$

The root commits to the ordered set of records represented by that tree.

An inclusion proof permits a verifier to establish that a particular record belongs to an authenticated tree state without retrieving the entire history.

A signed checkpoint binds the tree size and root to an identified transparency authority.

A portable evidence bundle can therefore contain the original signed claim, its inclusion proof and the associated signed checkpoint.

Verification is independent of the software interface used to produce that bundle.

Where a verifier possesses the expected public keys and trusted checkpoint information, it can independently establish signature integrity and inclusion.

7.1. Consistency and independent witnessing

An inclusion proof authenticates a record relative to a particular checkpoint. Long-term transparency also requires protection against attempts to show inconsistent histories to different observers.

Independently operated witnesses can retain authenticated checkpoints, compare successive tree states, examine consistency proofs and expose conflicting histories.

A mature distributed OpenAIIIS network must provide mechanisms for independent observation, checkpoint exchange, consistency verification and detection of equivocation.

The essential property is that the validity of a record should be checkable by an independent participant and that an operator's attempt to rewrite or selectively present history should be susceptible to cryptographic detection.

8. Decentralized Network Architecture

A universal trust protocol should not depend permanently on one organization's permission to register, inspect or verify digital records.

OpenAIIIS is designed for independently operated infrastructure, open verification and participation without a native cryptocurrency.

The decentralized network architecture extends the signed-claim and transparency model into a replicated system in which multiple operators validate and preserve a common registry history.

8.1. Independent operation

Participating nodes are intended to validate submitted records according to common protocol rules, maintain verifiable history and expose portable proofs.

Network participants should be capable of independently verifying record validity, ordering rules and cryptographic commitments rather than treating the operator that submitted a claim as a trusted authority over its correctness.

Claims about content, identity or authority remain subject to their own trust requirements even after network inclusion.

Consensus determines the accepted historical ordering of records; it does not convert every registered assertion into an authenticated real-world fact.

8.2. Consensus without a native token

OpenAIIIS's decentralized architecture is based on the principle that a common digital-trust infrastructure need not issue a speculative currency to operate.

The proposed permissionless network uses proof-of-work as a candidate foundation for independently verifiable ordering and resistance to unilateral history modification.

Nodes validate blocks and their contained records against deterministic rules. The network's fork-choice, difficulty adjustment, admission and reorganization rules must provide a common basis for identifying the accepted history.

Economic participation may be supported by the market for registry capacity, optional attestation services and fees denominated in existing payment assets, without requiring issuance of an OpenAIIIS token.

The final consensus and economic rules must align open participation with sustainable operation, resistance to spam and denial-of-service attacks, and incentives for independent infrastructure providers.

The consensus design is separate from the cryptographic validity of individual claims. Its purpose is to provide a common history that remains independently verifiable even when participants do not trust one another.

8.3. Universal participation

The registry is intended to serve individuals as well as institutions.

A creator seeking to preserve a signed authorship assertion should not be required to operate an autonomous agent or purchase an organizational identity credential.

A researcher should be able to register an evidence commitment without revealing a confidential manuscript.

An enterprise should be able to apply stronger issuer, consent, evidence and authorization requirements where the consequences demand them.

This distinction between open recording and differentiated trust assessment allows the same underlying network to support many applications without forcing all participants into a single identity, governance or commercial model.

9. Independent Historical Anchoring

A shared registry can gain an additional historical assurance by committing its authenticated checkpoints to an independently secured external system.

OpenAIIIS supports the use of OpenTimestamps and Bitcoin for this purpose.

Let Q_n represent an authenticated registry checkpoint covering the first n records.

A cryptographic commitment to Q_n may be submitted through OpenTimestamps and subsequently incorporated into the Bitcoin blockchain.

Once a valid Bitcoin attestation is available, a verifier can validate the commitment chain against an accepted Bitcoin block and apply an appropriate confirmation policy.

The relationship is:

Signed claims $\rightarrow$ authenticated registry checkpoint $\rightarrow$ OpenTimestamps commitment $\rightarrow$ Bitcoin historical evidence

This method allows many registry records to benefit from one checkpoint-level historical commitment.

It avoids requiring a separate Bitcoin transaction for every individual claim.

A Bitcoin-anchored checkpoint adds independently verifiable historical evidence about the committed registry state. The original OpenAIIIS signatures, authority checks, consent and governance records remain necessary to interpret the claims contained within it.

OpenAIIIS's own decentralized consensus and Bitcoin anchoring are complementary mechanisms. They provide distinct forms of assurance and can be verified independently.

10. Verification and Security

The security of OpenAIIIS depends on separating mathematical verification from the interpretation of claims.

A valid digital signature demonstrates that a particular signing key authenticated particular data. Real-world issuer identity and authority require independently trusted bindings.

An inclusion proof establishes membership in an authenticated registry state. Completeness, consistency and censorship resistance require further network and witnessing properties.

A historical commitment can establish that committed data existed by a verifiable historical bound. It cannot independently determine when a work was originally created or whether its author's factual assertions were correct.

These distinctions preserve the reliability of the protocol across different applications.

10.1. Threat model

An adversary may attempt to:

- Forge a claim or substitute a different artifact.
- Use a stolen or unauthorized signing key.
- Introduce false consent or delegated authority.
- Replay expired or revoked permissions.
- Submit fabricated evidence or security reports.
- Exploit contradictory governance records.
- Withhold, reorder or selectively disclose registry entries.
- Present inconsistent checkpoints.

- Infer confidential information from public commitments.

Cryptographic signatures, content commitments, typed authorization, verified consent, append-only history, independent witnessing and lifecycle checks provide complementary defenses.

A verifier must also establish that the history it relies upon is sufficiently complete and current for its policy. A valid historical receipt cannot establish the absence of later revocations when the verifier has not checked the relevant later history.

10.2. Privacy

OpenAIIS separates the public integrity history from original content and supporting evidence wherever possible.

This allows confidential works and sensitive records to be committed without automatically publishing their contents.

Nevertheless, a hash can reveal information when the underlying content is predictable, and public identifiers can expose relationships between actors.

Privacy-sensitive deployments must therefore apply appropriate disclosure, identifier, access-control and commitment strategies.

The protocol's verifiable history should preserve accountability without unnecessarily exposing the private material from which its evidence was derived.

10.3. Independent trust

The governing principle of verification is that trust must be established from identifiable evidence and configured authority.

An independent verifier should be able to distinguish record authenticity, content integrity, registered history, current authority, evidence sufficiency and external historical anchoring.

The resulting decision can be reproduced and challenged by another verifier possessing the relevant evidence and policy.

This provides a foundation for accountable disagreement rather than hiding different trust judgments behind a universal, unexplained approval mark.

11. Applications

OpenAIIS's universal record model allows one protocol to support a broad range of activities that currently rely on different forms of verification infrastructure.

11.1. Creators, inventors and researchers

An inventor can register a cryptographic commitment to an invention disclosure and preserve independently verifiable evidence of its recorded history.

A writer, photographer, designer or musician can register signed assertions concerning original work and supporting evidence.

A researcher can link a manuscript to datasets, observations, earlier disclosures and institutional

attestations.

Subsequent disputes or additional evidence can be linked without replacing the original historical record.

Such evidence can assist attribution, prior-art investigation, licensing discussions, research integrity and dispute resolution under the applicable legal and institutional standards.

11.2. Media and information integrity

A media publisher can associate content with signed origin, publication and editorial claims.

Later corrections, disputes, withdrawn attestations or security-relevant findings can be recorded as related events.

Independent platforms can verify the history of an artifact and apply their own editorial and trust policies.

This creates a foundation for verifiable media accountability even when the content travels beyond its original publishing platform.

11.3. Software supply chains

Software publishers can bind releases to exact content digests, authenticated issuers, delegated publishing authority and independently inspectable evidence.

Security teams can reassess releases when keys are compromised or authority changes.

Deployment infrastructure can require validated governance and evidence before accepting a release.

The resulting system supports continuous verification of both the artifact and the authority under which it was published.

11.4. Artificial intelligence and autonomous systems

An AI-generated artifact can carry a signed declaration identifying its asserted provenance and associated responsible actor.

Autonomous software can operate under scoped delegations linked to authenticated principals.

Verification systems can determine whether an action was within its authorized scope and whether subsequent revocations or security evidence should affect its execution.

The same accountability model can support agent-generated code, automated decisions and future machine-to-machine trust relationships.

11.5. Digital identity and institutional attestations

Trusted issuers can make signed identity or artifact attestations subject to appropriate consent, evidence and lifecycle rules.

Independent relying parties can assess those attestations against their own trust frameworks.

The common history permits attestations, withdrawals, revocations and disputes to remain linked and auditable over time.

12. Relationship to Existing Infrastructure

OpenAIIIS builds upon established principles of public-key cryptography, authenticated data structures, transparent logs, credentials and decentralized consensus.

Existing ecosystems have developed important mechanisms for specific domains.

Certificate Transparency established widely adopted techniques for auditing certificate issuance through append-only logs. Sigstore and Rekor apply signatures and transparency to software supply-chain integrity. C2PA defines signed provenance structures for digital media. Verifiable Credentials support portable issuer assertions. SCITT provides an architecture for transparent signed statements and receipts. OpenTimestamps supports independently verifiable historical commitments.

These mechanisms provide valuable capabilities and potential interoperability foundations.

OpenAIIIS's proposition is the integration of provenance, governance, accountable responsibility, safety-relevant lifecycle evidence and independent historical verification into one universal trust framework.

A media provenance assertion can participate in broader governance history. A software signature can be evaluated alongside delegated authority and accepted responsibility. An identity attestation can retain its consent and revocation lifecycle. A security report can be related to affected digital subjects and their existing trust evidence.

The significance lies in enabling these relationships to be expressed and independently evaluated through a common protocol, rather than requiring every application to reconstruct them within an isolated platform.

The architecture is therefore intended to complement existing infrastructure while providing a shared evidence and governance layer across domains.

13. Adoption and Open Verification

The usefulness of a universal protocol depends upon the independence and accessibility of its verification process.

OpenAIIIS is designed so that applications can exchange signed claims and portable proofs without requiring every participant to use the same identity provider, storage service, organizational policy or user interface.

Different systems can adopt the same cryptographic history while retaining their own decision-making responsibilities.

A university may recognize particular research attestations. An enterprise may require additional organizational authorization. A public software repository may apply different release-security requirements. An individual may use the registry to preserve evidence of an original work without obtaining institutional certification.

These applications can share the underlying record format and verification principles.

Open protocol specifications, independently executable verification, reproducible conformance checks

and transparent governance are necessary to establish confidence across independent implementations. The intended end state is infrastructure whose integrity can be tested by its participants rather than accepted solely on the reputation of its operator.

14. Conclusion

The Internet transformed the distribution of information, but digital trust remains fragmented across institutions, platforms and applications.

Artificial intelligence and increasingly autonomous software make this fragmentation more consequential. Content can be created without clear attribution, authority can be delegated without portable accountability, and security-relevant information can remain disconnected from the systems that need it.

OpenAIIS proposes a common foundation for addressing these problems.

Cryptographically signed claims establish provenance. Explicit authority and consent govern their meaning. Responsibility records connect actions to accountable actors. Authenticated lifecycle and security evidence allow trust to be reassessed as circumstances change. A shared, independently verifiable history preserves the record from which these decisions can be examined.

Decentralized operation is intended to make that history available beyond the control of any single institution, while optional Bitcoin anchoring provides an additional independent historical commitment.

The central contribution of OpenAIIS is the connection of digital provenance, governance, responsibility and safety into one independently verifiable system.

A digital world in which anyone can create, publish or automate consequential actions also needs a way for independent parties to examine their authority, evidence and accountability.

OpenAIIS is proposed as a foundation for that world.

References

- [1] S. Nakamoto, Bitcoin: A Peer-to-Peer Electronic Cash System, 2008. <https://bitcoin.org/bitcoin.pdf>
- [2] B. Laurie, E. Messeri and R. Stradling, Certificate Transparency Version 2.0, RFC 9162, 2021. <https://www.rfc-editor.org/rfc/rfc9162>
- [3] O. Steele, H. Birkholz, A. Delignat-Lavaud and C. Fournet, CBOR Object Signing and Encryption (COSE) Receipts, RFC 9942, 2026. <https://www.rfc-editor.org/rfc/rfc9942>
- [4] H. Birkholz et al., An Architecture for Trustworthy and Transparent Digital Supply Chains, RFC 9943, 2026. <https://www.rfc-editor.org/rfc/rfc9943>
- [5] Coalition for Content Provenance and Authenticity, C2PA Technical Specifications. <https://spec.c2pa.org/specifications/>
- [6] Sigstore Project, Rekor Transparency Log. <https://docs.sigstore.dev/logging/overview/>
- [7] World Wide Web Consortium, Verifiable Credentials Data Model v2.0, 2025. <https://www.w3.org/TR/vc-data-model-2.0/>

[8] OpenTimestamps Project, OpenTimestamps: A Timestamping Proof Standard.
<https://opentimestamps.org/>

[9] Community Cryptography Specification Project, Transparency Log Checkpoints.
<https://c2sp.org/tlog-checkpoint>

OpenAIIIS

CREATIVE COMMONS LICENSING NOTICE

OpenAIIIS Whitepaper - Version 4.0

Copyright (c) 2026 Tarun Sethi

This publication is licensed under the Creative Commons Attribution 4.0 International License (CC BY 4.0).

You may share and adapt it, including for commercial purposes, provided you follow the license conditions.

Attribution must identify the work and its version, credit the author, link to the license, and indicate whether changes were made.

License: <https://creativecommons.org/licenses/by/4.0/>

Legal terms: <https://creativecommons.org/licenses/by/4.0/legalcode.en>

The license covers copyright permissions in this document.

It does not grant patent rights or trademark rights.

Software is separately licensed under the Apache License 2.0.

Independent compatible implementations are welcomed; relevant patent, trademark and other rights remain separate from CC BY.

Official project: <https://openaiis.org>

This notice was added as a separate final page in October 2026.

The original publication pages remain unchanged.