None
Internet-Draft
Intended status: Standards Track
Expires: 24 July 2026

J. Ramirez
20 January 2026

Real Human Verification (RHV): A Hardware-Rooted Cryptographic Standard for Human-Origin Visual Evidence
draft-ramirez-rhv-core-00

Abstract

This document defines Real Human Verification (RHV), a public cryptographic standard for establishing a constitutional root of trust for direct human-origin visual evidence in the post-AI era. RHV specifies cryptographic primitives, hardware-backed capture requirements, and public transparency logging mechanisms that allow digital photos and videos to prove that they originate from a physically authenticated human capture pipeline. RHV does not judge content truth; it defines cryptographically verifiable human-origin capture, tamper-evident integrity, and publicly auditable provenance chains. By anchoring visual evidence in hardware roots of trust and append-only transparency logs, RHV provides a neutral, publicly verifiable foundation for judicial, institutional, and societal reliance on digital visual evidence across operating systems, cameras, drones, and sensor-based capture devices.

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 24 July 2026.

Table of Contents

1. Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79. This document is not an Internet Standard. It is a work in progress and is subject to change. It may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use this document as reference material or to cite it other than as a “work in progress.”

2. Introduction

Digital visual evidence has historically served as a foundational pillar for legal, institutional, and societal trust. For over a century, photographs and videos have functioned as primary instruments for documenting reality, establishing accountability, and preserving historical truth. The emergence of generative artificial intelligence has fundamentally altered this trust model. Synthetic images and videos can now be generated that are visually indistinguishable from direct human-origin capture, rendering traditional visual verification methods unreliable. As a result, the evidentiary value of digital visual media is undergoing rapid erosion across judicial systems, public institutions, media organizations, and digital platforms. At present, there exists no public, neutral, and cryptographically verifiable constitutional root of trust for establishing whether digital visual content originated from a physically authenticated human capture pipeline. Existing authenticity mechanisms focus primarily on content provenance, metadata, or platform-level attestations, but do not provide a universal, hardware-backed, and publicly auditable guarantee of direct human-origin capture. Real Human Verification (RHV) defines a constitutional cryptographic layer that addresses this structural deficiency. RHV introduces a public, vendor-neutral, and hardware-rooted framework for establishing verifiable human-origin capture, tamper-evident integrity, and publicly auditable provenance chains for digital photos and videos. By anchoring visual evidence in hardware roots of trust and append-only transparency logs, RHV establishes a foundational trust primitive upon which operating systems, hardware manufacturers, browsers, platforms, and legal institutions can build interoperable systems for visual evidence verification in the post-AI era.

3. Terminology and Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119]. The following terms are used throughout this specification:

RHV (Real Human Verification): The cryptographic standard defined in this document for establishing a human-origin root of trust for visual digital evidence.

Human-Origin Capture: A capture process in which photons or sensor signals originate from the physical environment and are acquired by a physically authenticated sensor pipeline backed by a hardware root of trust.

CAPTURED Asset: A photo or video that has been directly acquired through a Human-Origin Capture process and for which an RHV Proof has been generated by a hardware-backed secure capture pipeline.

DERIVED Asset: Any photo or video that has been transformed, edited, re-encoded, cropped, composited, or otherwise modified from a CAPTURED asset or another DERIVED asset. DERIVED assets are not considered direct human-origin captures, but MAY maintain cryptographically verifiable provenance links to their originating assets.

RHV Proof: A cryptographic statement generated by a secure capture pipeline that binds a media hash, capture time, and hardware-backed human-origin attestation into a signed structure.

Hardware Root of Trust (HRoT): A tamper-resistant hardware security component (e.g., Secure Enclave, Trusted Execution Environment (TEE), StrongBox, Trusted Platform Module (TPM), Secure Element) capable of generating and protecting cryptographic keys used to attest that media originated from a physically authenticated sensor pipeline.

Secure Capture Pipeline: A hardware-backed data path from sensor to signing environment that ensures authenticity, integrity, and temporal correctness of captured visual data.

Transparency Log: A publicly verifiable, append-only, cryptographically auditable log used to anchor RHV Proofs and provide public auditability and immutability.

Chain of Provenance: A cryptographically verifiable sequence of links connecting a DERIVED asset to its originating CAPTURED asset through signed transformation statements.

RHV Anchor ID: The unique identifier assigned to an RHV Proof once it has been anchored into a Transparency Log.

4. Problem Statement

The trustworthiness of digital visual evidence is undergoing a systemic collapse. Advances in generative artificial intelligence have enabled the creation of synthetic images and videos that are visually indistinguishable from direct human-origin capture. As a result, traditional methods for assessing the authenticity of visual media—based on visual inspection, metadata analysis, or platform-level attestations—are no longer reliable indicators of physical reality. Judicial systems, insurers, journalists, scientific institutions, and digital platforms increasingly face the inability to establish whether a given visual asset represents a physically

captured event or a synthetic fabrication. This erosion of visual evidentiary trust undermines legal due process, public accountability, and the integrity of historical records. Current provenance and authenticity frameworks primarily focus on recording transformations, authorship, or platform-level content attestations. While these approaches provide valuable context, they do not establish a universal, hardware-backed, publicly auditable root of trust capable of proving that visual content originated from a physically authenticated human capture pipeline. Without such a root of trust, it is cryptographically impossible to distinguish direct human-origin capture from synthetic generation in a universally verifiable manner. This structural deficiency constitutes a fundamental gap in the trust architecture of the modern digital ecosystem. RHV is designed to address this gap by defining a constitutional cryptographic root of trust for human-origin visual evidence.

5. Design Goals

RHV is designed as a constitutional trust layer for digital visual evidence. Its design is guided by the following foundational goals.

5.1. Hardware-Anchored Human-Origin Proof

RHV MUST anchor human-origin verification in hardware roots of trust. Proof of human-origin capture MUST be derived from physically authenticated sensor pipelines backed by tamper-resistant secure elements.

5.2. Public Verifiability

RHV MUST enable any third party to independently verify proofs without reliance on proprietary platforms, closed services, or vendor-specific trust anchors.

5.3. Chain of Provenance

RHV MUST support cryptographically verifiable provenance chains linking DERIVED assets to their originating CAPTURED assets, enabling full auditability of transformations.

5.4. Tamper Evidence

Any modification to an RHV-protected asset MUST be detectable through cryptographic integrity verification.

5.5. Vendor Neutrality

RHV MUST NOT depend on any single platform vendor, operating system, or hardware manufacturer. The standard MUST be implementable across heterogeneous ecosystems.

5.6. Privacy Preservation

RHV MUST minimize the exposure of personal data. Human-origin attestation MUST NOT require user identification, biometric data, or persistent personal identifiers.

5.7. Scalability

RHV MUST be capable of supporting planetary-scale adoption with minimal storage, bandwidth, and computational overhead.

5.8. Forward Compatibility

RHV MUST support extensibility to accommodate future cryptographic algorithms, hardware roots of trust, and capture device classes.

6. System Overview

RHV defines a layered, hardware-anchored architecture for establishing and verifying human-origin visual evidence. The system is composed of three primary layers: Secure Capture, Public Anchoring, and Verification.

6.1. Secure Capture Layer

The Secure Capture Layer operates within devices capable of Human-Origin Capture (e.g., smartphones, cameras, drones, and sensor-based capture hardware). It includes:

A Secure Capture Pipeline that authenticates sensor-originated data. A Hardware Root of Trust (HRoT) that generates and protects attestation keys. A signing environment that produces RHV Proofs binding media hashes, capture time, and human-origin attestation.

This layer is responsible for producing cryptographically verifiable statements asserting that an asset was directly captured by a physically authenticated sensor pipeline.

6.2. Public Anchoring Layer

The Public Anchoring Layer provides Transparency Logs that:

Are append-only and publicly auditable. Anchor RHV Proofs to produce globally verifiable RHV Anchor IDs. Enable third-party auditability, immutability, and independent verification.

Transparency Logs form the public memory of RHV, ensuring that proofs cannot be retroactively altered or suppressed without detection.

6.3. Verification Layer

The Verification Layer enables any third party to:

Validate RHV Proof signatures. Confirm anchoring in Transparency Logs. Reconstruct Chains of Provenance linking DERIVED assets to their originating CAPTURED assets. Assess tamper evidence and transformation history.

This layer is intentionally vendor-neutral and does not require trusted platform intermediaries.

6.4. Storage and Scalability Model

RHV DOES NOT store visual content.
RHV stores only cryptographic proofs and transparency anchors.

Each RHV Proof consists exclusively of cryptographic material, including:

* Media Hash: 32 bytes

* Timestamp: 8 bytes

* Hardware Attestation Hash: 32 bytes

* Digital Signatures (dual-root): 128 bytes

* Log metadata and indexing: approximately 56 bytes

Total (rounded): *256 bytes per proof*.

Thus, each RHV Proof occupies approximately 256 bytes (0.00025 MB).

RHV Proofs bind cryptographic hashes to a specific canonical representation of the captured media. Any modification to the media bitstream, including accidental corruption or non-canonical metadata changes, SHALL invalidate the corresponding RHV Proof unless such modifications are explicitly recorded as a DERIVED asset with a signed derivation statement.

RHV does not attempt to preserve proof validity across arbitrary or untracked modifications. This behavior is intentional and ensures strong tamper-evidence, forensic integrity, and unambiguous origin verification.

Implementations MAY define canonicalization procedures for media hashing, such as excluding non-semantic metadata or normalizing container-level fields, provided that such procedures are explicitly specified, deterministic, and consistently applied during both proof generation and verification.

6.4.1. Planetary-Scale Feasibility

At a global adoption scale:

* *100 million captures/day* → 36.5 billion proofs/year → ~9.3 TB/year

* *1 billion captures/day* → 365 billion proofs/year → ~93 TB/year

* *10 billion captures/day* → 3.65 trillion proofs/year → ~930 TB/year

Using current commodity cloud storage pricing (approximately USD 20 per TB per month), the total annual storage cost for planetary-scale RHV operation remains under:

USD 300,000 per year for the entire world.

These estimates are order-of-magnitude approximations.

6.4.2. Economic Implication

RHV storage requirements are several orders of magnitude smaller than conventional content storage and are economically trivial compared to:

* social media video storage

* CDN image caches

* legal discovery archives

* surveillance and logging systems

RHV therefore introduces no material economic barrier to planetary-scale deployment.

6.4.3. Constitutional Implication

Because RHV stores only cryptographic proofs and not visual content, RHV:

* avoids data hoarding
* avoids privacy risk
* avoids censorship risk
* avoids cost centralization
* remains globally replicable

This ensures RHV can function as a permanent constitutional layer of Internet memory.

7. Core Cryptographic Primitive

RHV defines a minimal cryptographic primitive that establishes the foundational proof of human-origin capture and enables public auditability and provenance tracking.

For any visual media asset, the following primitive is defined:

H = HASH(media) P = SIGN(HASH_ID || SENSOR_CLASS_ID || H || CAPTURE_TIME || HUMAN_ORIGIN_ATTESTATION) ID = ANCHOR(P)

Where:

HASH_ID is an identifier for the cryptographic hash algorithm used, as defined by the RHV Cryptographic Parameters registry.

SENSOR_CLASS_ID identifies the class of capture sensor and secure capture pipeline (e.g., smartphone camera, drone camera, body-worn camera, industrial sensor), as asserted by the Hardware Root of Trust.

Including HASH_ID and SENSOR_CLASS_ID in the signed payload cryptographically binds the proof to a specific cryptographic algorithm and physical capture class, preventing downgrade attacks, algorithm substitution, or sensor class ambiguity.

The RHV Proof (P) is generated and cryptographically signed prior to anchoring.

The RHV Anchor ID is generated exclusively by the Transparency Log as a result of the ANCHOR(P) operation. The Anchor ID is an externally assigned reference used for public lookup and auditability and SHALL NOT be considered part of the signed RHV Proof payload.

The resulting Anchor ID uniquely and immutably references the RHV Proof within a publicly auditable Transparency Log, establishing a constitutional cryptographic root of trust for human-origin visual evidence.

8. Secure Capture Pipeline

The Secure Capture Pipeline defines the mandatory requirements for generating RHV Proofs from direct human-origin capture.

8.1. Pipeline Requirements

Implementations SHALL use a hardware-backed secure capture pipeline that ensures that sensor-originated data cannot be injected, replayed, or modified prior to attestation. A Secure Capture Pipeline MUST: Authenticate sensor-originated data using hardware-enforced mechanisms (e.g., secure MIPI CSI channels, secure ISP validation). Execute cryptographic operations within a Hardware Root of Trust (HRoT). Generate cryptographically verifiable capture timestamps. Prevent post-capture modification prior to RHV Proof generation.

8.2. Attestation Generation

For each CAPTURED asset, the Secure Capture Pipeline SHALL generate a HUMAN ORIGIN ATTESTATION that asserts: _ _ The asset was captured by a physically authenticated sensor pipeline. The capture occurred within a secure hardware environment. The capture timestamp was generated by a trusted time source. The attestation SHALL be signed using hardware-protected keys that are non-exportable and tamper-resistant.

8.3. Non-Spoofability

Implementations SHALL reject any media that cannot be cryptographically bound to an authenticated sensor pipeline. Media injected, replayed, or synthesized outside the secure capture pipeline MUST NOT produce valid RHV Proofs.

8.4. Physical Signal Integrity and Analog Validation

To mitigate sensor injection and replay attacks, RHV-compliant implementations SHOULD employ physical layer authentication mechanisms. Secure Sensor Channel: Raw sensor data MUST be authenticated and/or encrypted at the sensor interface (e.g., MIPI CSI-2 secure modes) prior to transmission to the SoC. Analog Noise and Jitter Analysis: The Secure ISP SHALL perform statistical analysis of sensor-specific physical noise characteristics (including fixed-pattern noise and thermal noise) to detect synthetic or replayed signals. Timing Attestation: The TEE SHALL validate that frame timing and jitter correspond to physical sensor behavior. Signals failing physical consistency validation MUST be rejected.

9. Secure Capture Pipeline

The Secure Capture Pipeline defines the mandatory requirements for generating RHV Proofs from direct human-origin capture.

9.1. Pipeline Requirements

Implementations SHALL use a hardware-backed secure capture pipeline that ensures sensor-originated data cannot be injected, replayed, or modified prior to attestation.

A Secure Capture Pipeline MUST:

* Authenticate sensor-originated data using hardware-enforced mechanisms (e.g., secure MIPI CSI channels, secure ISP validation).

* Execute cryptographic operations within a Hardware Root of Trust (HRoT).

* Generate cryptographically verifiable capture timestamps.

* Prevent post-capture modification prior to RHV Proof generation.

9.2. Attestation Generation

For each CAPTURED asset, the Secure Capture Pipeline SHALL generate a *HUMAN_ORIGIN_ATTESTATION* asserting that:

* The asset was captured by a physically authenticated sensor pipeline.

* The capture occurred within a secure hardware execution environment.

* The capture timestamp was generated by a trusted time source.

The attestation SHALL be signed using hardware-protected keys that are non-exportable and tamper-resistant.

9.3. Non-Spoofability

Implementations SHALL reject any media that cannot be cryptographically bound to an authenticated sensor pipeline.

Media injected, replayed, or synthesized outside the secure capture pipeline MUST NOT produce valid RHV Proofs.

9.4. Physical Signal Integrity and Analog Validation

To mitigate sensor injection and replay attacks, RHV-compliant implementations SHOULD employ physical layer authentication mechanisms.

9.4.1. Secure Sensor Channel

Raw sensor data MUST be authenticated and/or encrypted at the sensor interface (e.g., MIPI CSI-2 secure modes) prior to transmission to the System on Chip (SoC).

9.4.2. Analog Noise and Jitter Analysis

The Secure Image Signal Processor (ISP) SHALL perform statistical analysis of sensor-specific physical noise characteristics, including fixed-pattern noise and thermal noise, to detect synthetic or replayed signals.

9.4.3. Timing Attestation

The Trusted Execution Environment (TEE) SHALL validate that frame timing and jitter correspond to physical sensor behavior.

Signals failing physical consistency validation MUST be rejected.

10. Hardware Root of Trust

RHV relies on Hardware Roots of Trust (HRoT) to anchor human-origin attestation in tamper-resistant physical hardware.

10.1. Root of Trust Requirements

An HRoT MUST:

* Protect private attestation keys against extraction.

* Execute cryptographic operations in an isolated, tamper-resistant environment.

* Provide secure key storage and attestation services.

* Support secure boot and firmware integrity validation.

Examples of acceptable HRoT implementations include Secure Enclave, Trusted Execution Environments (TEE), StrongBox, Trusted Platform Modules (TPM), and dedicated Secure Elements.

10.2. Device Attestation Keys

Each compliant device SHALL possess one or more hardware-protected attestation keys used exclusively for RHV Proof generation.

These keys:

* MUST NOT be exportable.

* MUST be generated and stored within the HRoT.

* MUST support revocation and rotation mechanisms.

10.3. Physical Capture Assertion

The HRoT SHALL provide a cryptographic assertion that the signed media originated from a physically authenticated sensor pipeline and was not synthetically generated or externally injected.

10.4. Dual Public Root Attestation

RHV SHALL employ a dual-root attestation model.

Each RHV Proof SHALL be signed by:

* A device Hardware Root of Trust attesting physical capture.

* A public RHV network root attesting public anchoring.

This dual attestation prevents isolated device forgeries, private log suppression, and unverifiable local-only proofs.

11. Transparency Log

RHV defines a public Transparency Log as a foundational component for anchoring RHV Proofs and enabling independent auditability.

11.1. Log Properties

A Transparency Log MUST be:

* Append-only.
* Publicly readable.
* Cryptographically verifiable.
* Resistant to retroactive modification.
* Replicable across independent operators.

The log SHALL support inclusion proofs and consistency proofs to allow third parties to independently audit log integrity.

11.2. Anchoring Operation

The ANCHOR(P) operation appends a previously generated and signed RHV Proof P into a Transparency Log and returns a globally unique RHV Anchor ID.

The RHV Anchor ID is generated exclusively by the Transparency Log upon successful inclusion of the signed proof. The Anchor ID is an externally assigned reference used for public lookup and auditability and SHALL NOT be considered part of the RHV Proof signature or payload.

11.3. Public Auditability

Any party MAY audit the Transparency Log to verify:

That a given RHV Proof has been anchored. That the log has not been modified, truncated, or forked. That inclusion proofs and consistency proofs are valid.

12. Capture Modes

RHV defines two normative capture modes to distinguish direct physical capture from subsequent transformations.

12.1. CAPTURED

A *CAPTURED* asset is a photo or video that originates from a Secure Capture Pipeline and has an RHV Proof anchored in a Transparency Log.

CAPTURED assets represent the cryptographic root of a provenance chain and are considered direct human-origin captures.

Only CAPTURED assets establish new human-origin roots.

12.2. DERIVED

A *DERIVED* asset is any asset that results from transformation of a CAPTURED asset or another DERIVED asset, including but not limited to:

* Cropping
* Editing
* Compositing
* Re-encoding
* Resizing
* Color correction
* Montage
* Synthesis operations

DERIVED assets MUST NOT be treated as direct human-origin captures. However, they MAY maintain cryptographically verifiable provenance links to their originating assets.

Each DERIVED asset SHALL produce a signed derivation statement linking it to its parent asset, forming a publicly auditable Chain of Provenance.

12.3. Human-Derived Provenance Lineage

RHV defines a *Lineage Chain (LC)* as a sequence of cryptographic assertions anchored to a *Root Human Event (RHE)*.

12.3.1. Root Human Event (RHE)

The initial CAPTURED asset signed within a Hardware Root of Trust.

12.3.2. Derived Assertions

Any subsequent modification (including but not limited to cropping, color grading, temporal editing, re-encoding, compositing, or synthesis) MUST be recorded as a new derivation statement referencing the parent asset hash.

RHV DOES NOT evaluate the semantic intent or legitimacy of transformations. Instead, it provides a cryptographically verifiable traceability path.

In judicial or institutional contexts, an auditor SHALL be able to reconstruct the full lineage of an asset.

If the lineage chain is broken or contains unauthenticated synthetic generation steps, the asset MUST be classified as *UNVERIFIED*.

13. Verification Process

RHV defines a public verification process that allows any party to independently validate human-origin capture, integrity, and provenance of a visual asset.

13.1. Proof Validation

To verify an RHV-protected asset, a verifier SHALL perform the following steps:

Validate the RHV Proof structure and digital signatures, including verification of the hardware-backed attestation and the public anchoring signature.

Confirm anchoring of the RHV Proof in a Transparency Log using inclusion and consistency proofs.

Upon successful validation of the RHV Proof, compute H = HASH(media).

Verify that the computed media hash matches the hash bound within the validated RHV Proof.

13.2. Provenance Reconstruction

For DERIVED assets, a verifier SHALL reconstruct the Chain of Provenance by:

* Validating each signed derivation statement.

* Confirming that each link references a valid parent asset.

* Ensuring that the chain terminates in a CAPTURED asset.

13.3. Determination of Origin Class

Upon successful verification, an asset SHALL be classified as one of the following:

CAPTURED — Direct human-origin capture. *DERIVED-N* — Nth-order derivation from a CAPTURED asset.

14. Privacy Considerations

RHV is designed to establish verifiable human-origin capture while preserving individual privacy and minimizing personal data exposure.

14.1. Data Minimization

RHV MUST NOT require the collection, storage, or disclosure of: Personally identifiable information (PII) Biometric identifiers User account identifiers Persistent device identifiers Location data beyond cryptographically necessary timing information RHV Proofs SHALL contain only the minimum cryptographic material required to establish origin, integrity, and provenance.

14.2. Pseudonymous Hardware Attestation

Hardware-backed attestation keys SHALL be pseudonymous and non-linkable across independent capture events, except where explicitly required for revocation or security maintenance. Implementations SHOULD support key rotation to prevent long-term device correlation.

14.3. No Behavioral Tracking

RHV MUST NOT introduce behavioral tracking, profiling, or persistent surveillance capabilities. Transparency Logs record cryptographic proofs only and do not store user activity histories, content payloads, or personal metadata.

14.4. Selective Disclosure

Implementations MAY support selective disclosure mechanisms allowing holders of RHV-protected assets to disclose only those provenance elements necessary for verification in specific legal, institutional, or contractual contexts.

15. Threat Model

RHV defines a comprehensive threat model addressing adversarial attempts to forge, alter, suppress, or misrepresent human-origin visual evidence.

15.1. Synthetic Media Injection

An adversary may attempt to generate synthetic images or videos and falsely present them as human-origin captures.

Mitigation:

RHV relies on Hardware Roots of Trust (HRoT) to ensure that only media originating from physically authenticated sensor pipelines can generate valid RHV Proofs. Media generated outside a secure capture pipeline MUST NOT produce valid proofs.

15.2. Sensor Injection Attacks

An adversary may attempt to inject replayed or synthetic data directly into the sensor data path.

Mitigation:

Implementations SHALL use secure sensor channels, secure Image Signal Processors (ISPs), and hardware-backed attestation mechanisms to reject unauthenticated sensor-originated data.

15.3. Local-Only Proof Forgery

An adversary may attempt to create locally valid but publicly unverifiable proofs, bypassing public auditability.

Mitigation:

RHV SHALL employ a dual-root attestation model. Each RHV Proof SHALL be signed by:

A device Hardware Root of Trust attesting physical capture A public RHV Network Root attesting public anchoring

This dual attestation model prevents isolated device forgeries and unverifiable local-only proofs.

15.4. Log Suppression or Forking

An adversary may attempt to suppress, truncate, or fork Transparency Logs.

Mitigation:

Transparency Logs SHALL support public replication, inclusion proofs, and consistency proofs, enabling independent detection of suppression, truncation, or divergence.

15.5. Replay and Substitution Attacks

An adversary may attempt to reuse valid RHV Proofs for different media assets.

Mitigation:

Each RHV Proof cryptographically binds the media hash, capture time, and hardware-backed attestation, preventing substitution, replay, or reassignment to different media.

16. Security Considerations

RHV establishes a constitutional cryptographic root of trust for human-origin visual evidence by combining hardware-backed attestation, public transparency anchoring, and cryptographically verifiable provenance chains.

16.1. Hardware Integrity

All cryptographic operations used to generate RHV Proofs MUST be executed within a Hardware Root of Trust (HRoT). Private attestation keys MUST NOT be exportable and MUST be protected by tamper-resistant hardware.

16.2. Dual-Root Trust Enforcement

RHV Proofs MUST include dual-root attestation consisting of: - A hardware root of trust attesting physical capture - A public RHV network root attesting public anchoring Both roots are REQUIRED for an RHV Proof to be considered valid.

16.3. Cryptographic Agility

Implementations SHALL support algorithm agility to allow migration to successor cryptographic primitives without invalidating existing proofs.

16.4. Log Replication and Survivability

Transparency Logs SHALL be replicable across independent operators to ensure survivability, censorship resistance, and global auditability.

16.5. Revocation

RHV SHALL support revocation mechanisms for compromised devices, attestation keys, or Transparency Log operators. Revocation events MUST be publicly auditable. Revocation in RHV applies exclusively to cryptographic attestation keys or trust roots, not to users, personal identities, or persistent device identifiers. RHV does not require globally persistent device identifiers. Revocation is achieved through the invalidation of compromised attestation keys, key epochs, or trust roots, consistent with established public-key infrastructure and hardware attestation models. This approach preserves privacy while enabling effective and publicly auditable revocation.

17. Governance & Neutrality

RHV is maintained by a private neutral steward organization responsible for preserving the integrity, neutrality, and continuity of the RHV specification. The RHV specification is a public, vendor-neutral constitutional standard. No single vendor, platform, government, or commercial entity has privileged control over the RHV protocol. All changes to the RHV specification SHALL be subject to public technical review, cryptographic scrutiny, and documented consensus processes. Implementations of RHV MAY be open-source or proprietary. The RHV specification itself is public and freely implementable by any party without licensing restrictions. The RHV steward organization MAY provide commercial services, infrastructure, verification endpoints, certification programs, and compliance tooling, provided such services do not compromise the neutrality, openness, or auditability of the RHV protocol.

18. IANA Considerations

This document defines an optional registry request. The authors propose the creation of a new IANA registry: Registry Name: RHV Cryptographic Parameters Registration Procedure: Specification Required (RFC 8126) Initial Entries: 0x01 — SHA-256 (Mandatory) 0x02 — BLAKE3 (Optional) 0x03 — Ed25519 (Digital Signature)

19. References

Normative References RFC 2119 — Key words for use in RFCs to Indicate Requirement Levels RFC 6962 — Certificate Transparency FIPS 186-4 — Digital Signature Standard NIST SP 800-53 — Security and Privacy Controls ISO/IEC 27001 — Information Security Management WebAuthn (W3C) C2PA Specification Informative References Secure Enclave Architecture (Apple) Android StrongBox / TEE Architecture MIPI CSI-2 Secure Channel Specification TPM 2.0 Specification Certificate Transparency Logs

20. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", RFC 2119, 1997, <https://www.rfc-editor.org/rfc/rfc2119>.

Author’s Address

Jonathan Ramrez
Email: jonathan.ramirezf@gmail.com