Individual Submission
Internet-Draft
Intended status: Experimental
Expires: 22 August 2026

D. Condrey
WritersLogic

18 February 2026

Proof of Process (PoP): Forensic Appraisal and Security Model
draft-condrey-rats-pop-appraisal-04

Abstract

This document specifies the forensic appraisal methodology and quantitative security model for the Proof of Process (PoP) framework. It defines how Verifiers evaluate behavioral entropy, perform liveness detection, and calculate forgery cost bounds. Additionally, it establishes the taxonomy for Absence Proofs and the Writers Authenticity Report (WAR) format, as well as the Tool Receipt protocol for artificial intelligence (AI) attribution within the linear human authoring process.

Discussion Venues

This note is to be removed before publishing as an RFC.

Source for this draft and an issue tracker can be found at https://github.com/writerslogic/draft-condrey-rats-pop.

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 22 August 2026.

Table of Contents

1. Introduction

The value of Proof of Process (PoP) evidence lies in the Verifier's ability to distinguish biological effort from algorithmic simulation. While traditional RATS [RFC9334] appraisals verify system state, PoP appraisal verifies a continuous physical process. This document provides the normative framework for forensic appraisal, defining the logic required to generate a Writers Authenticity Report (WAR).

This document is a companion to [PoP-Protocol], which defines the Evidence Packet wire format and Attester procedures. The present document specifies the Verifier's appraisal logic, Attestation Result (WAR) wire format, and forensic methodology. Implementers of Verifier components require both documents.

At T3/T4 attestation tiers, platform integrity verification as described in the SEAT use cases [SEAT-UseCases] provides the trust anchor for PoP's hardware-bound claims. When PoP Evidence is delivered over an attested TLS channel [SEAT-EXPAT], the Verifier gains assurance that the Attesting Environment's platform was trustworthy during evidence generation.

2. Terminology

This document uses the following terms in addition to those defined in [RFC9334] and [PoP-Protocol]:

Synthetic Authoring: Content generated by AI or automated tools that is subsequently attributed to a human author.

Evidence Quantization: The process of reducing timing resolution in behavioral data to protect author privacy while maintaining forensic utility.

IKI (Inter-Keystroke Interval): The time elapsed between consecutive keystrokes, measured in milliseconds.

C_intra: Pearson correlation between pause duration and subsequent

edit complexity within a single checkpoint interval. Values near 0.0 indicate robotic pacing; values above 0.3 indicate human-like variable effort.

CLC (Cognitive Load Correlation): Statistical correlation between content semantic complexity and typing cadence, used to distinguish original composition from retyping.

SNR (Signal-to-Noise Ratio) Analysis: Spectral analysis of jitter intervals to distinguish biological motor noise patterns from synthetic injection.

3. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

4. Step-by-Step Verification Procedure

A Verifier MUST perform the following procedure to appraise a PoP Evidence Packet:

1. _Structural Validation:_ The Verifier MUST reject with verdict invalid (4) any Evidence Packet that: (a) fails CBOR decoding, (b) lacks CBOR tag 1347571280, (c) has version != 1, (d) is missing mandatory fields (keys 1-6 in evidence-packet, keys 1-9 in each checkpoint), or (e) contains CBOR types that do not match the CDDL schema.

2. _Chain Integrity:_ Verify the SHA-256 hash link between all checkpoints. Any break invalidates the entire Evidence Packet. The Verifier MUST set the verdict to invalid (4). The warnings field SHOULD include the checkpoint sequence number where the break was detected.

3. _Temporal Order:_ For each process-proof, recompute Argon2id from the declared seed to obtain state_0, then verify sampled Merkle proofs against the committed root (process-proof key 4, merkle-root). Verify that claimed-duration is within [0.5x, 3.0x] of the expected wall-clock time for the declared proof-params on reference hardware (defined as a system with DDR4 memory providing approximately 25 GB/s sustained bandwidth). Expected times are defined in [PoP-Protocol], Mandatory SWF Parameters section.

4. _Entropy Threshold:_ Independently estimate entropy from the jitter-binding intervals array using a standard entropy estimator (e.g., NIST SP 800-90B most common value estimator). Verify the independent estimate meets or exceeds 3.0 bits per inter-keystroke interval. The Attester's self-reported entropy-estimate field MUST NOT be relied upon. Low-entropy segments (below threshold) MUST be flagged as "Non-Biological."

5. _Entanglement:_ Verify the HMAC value (entangled-mac) over the combined document, jitter, and physical state.

6. _State Matching:_ Verify that the final checkpoint's content-hash matches the document-ref content-hash. Verify that the cumulative char-count from edit-deltas is consistent with the document-ref char-count.

7. _Channel Binding:_ If the Evidence Packet contains a channel-binding field and was received over TLS, verify that the binding-value matches the locally-computed TLS Exported Keying Material. Reject the Evidence Packet on mismatch.

Steps 4 and 5 apply only when jitter-binding and entangled-mac fields are present (ENHANCED and MAXIMUM profiles). For CORE Evidence Packets lacking these fields, the Verifier MUST skip Steps 4 and 5 and note in the WAR warnings that behavioral analysis was not performed.

5. Forensic Assessment Mechanisms

The appraisal logic is designed to detect "Synthetic Authoring" -- content generated by AI and subsequently "back-filled" with timing and hardware attestation.

SNR (Signal-to-Noise Ratio) Analysis: Verifiers MUST compute the power spectral density of jitter intervals. Human motor signals exhibit characteristic noise patterns consistent with biological motor control [Monrose2000]. Evidence exhibiting spectral flatness greater than 0.9 (indicating white noise rather than biological 1/f-like noise) MUST be flagged as potentially synthetic.

Cognitive Load Correlation (CLC): Verifiers MUST correlate timing

patterns with semantic complexity. Human authors exhibit increased inter-keystroke intervals (IKI) and pause frequency during composition of semantically complex segments compared to simple connective text. Verifiers MUST compute the Pearson correlation between segment semantic complexity and mean IKI. Evidence with r < 0.2 (or r < 0.1 in assistive mode) MUST be flagged as a Semantic Mismatch.

Mechanical Turk Detection: Verifiers MUST compute C_intra (Pearson correlation between pause duration and subsequent edit complexity within each checkpoint). C_intra values below 0.15 MUST be flagged as indicating robotic pacing, where an automated system maintains a machine-clocked editing rate independent of content demands.

Error Topology Analysis: Verifiers SHOULD analyze error patterns for consistency with human cognitive processing [Salthouse1986]: localized corrections near recent insertions, fractal self-similarity in revision patterns, and deletion-to-insertion ratios consistent with natural composition. Evidence exhibiting unnaturally low error rates (below 1 correction per 500 characters) or randomly distributed errors lacking positional correlation SHOULD be flagged.

QR Presence Challenge (OOB-PC): When presence-challenge structures are present in the Evidence Packet, Verifiers MUST verify that the response-time is within the corresponding checkpoint’s time window and MUST validate the device-signature. NOTE: The Attester-side procedure for issuing presence challenges is specified in [PoP-Protocol].

Session Consistency Analysis: Verifiers MUST analyze cross-checkpoint behavioral trends. IKI distributions should exhibit gradual drift consistent with fatigue effects. An abrupt change is defined as a shift in mean IKI between consecutive checkpoints exceeding 2 standard deviations of the session-wide IKI distribution. Verifiers MUST flag transitions exceeding this threshold as potential data source switching. Jitter-binding intervals across consecutive checkpoints MUST be checked for statistical independence (cross-checkpoint correlation below 0.3). Edit-delta patterns SHOULD be checked for non-stationarity consistent with human creative flow.

A conforming Verifier MUST evaluate all forensic mechanisms for which the Evidence Packet contains sufficient data. Any single triggered flag is sufficient to assign the suspicious verdict. Verifiers MAY implement additional analysis mechanisms beyond those defined in this specification.

5.1. SNR Computation (Informative)

The signal-to-noise ratio measures productive editing activity versus idle or mechanical noise within each evidence window:

```
SNR = 10 * log10(P_signal / P_noise)

where:
  P_signal = (keystroke_count + revision_count) / window_duration
  P_noise  = (pause_total_ms + idle_intervals) / window_duration
```

Typical ranges observed in human authorship:

* Human sessions: -3 dB to +12 dB, with variation reflecting cognitive processing cycles.

* Automated input (copy-paste, scripted typing): consistently above +15 dB due to minimal pause behavior.

* Sessions above +20 dB across all windows SHOULD be flagged as potentially non-human.

* Sessions below -10 dB across all windows indicate predominantly idle behavior and SHOULD be flagged as potentially fabricated padding.

The Verifier SHOULD compute per-window SNR and session-wide SNR statistics (mean, variance, trend) as forensic indicators.

5.2. CLC and IKI Computation (Informative)

The Compositional Lyapunov Coefficient (CLC) measures the rate at which writing complexity evolves over the session, analogous to Lyapunov exponents in dynamical systems:

```
CLC = (1/n) * sum_{i=1}^{n} ln(|delta_IKI[i]| / |delta_IKI[i-1]|)

where:
  delta_IKI[i] = IKI_mean[i] - IKI_mean[i-1]
  n = number of consecutive window pairs
```

The Incremental Kolmogorov Information (IKI) measures informational complexity added per window:

```
IKI[i] ~= compressed_size(delta_content[i]) / raw_size(delta_content[i])
```

Typical ranges: human authorship exhibits positive CLC values (0.01 to 0.5) reflecting natural creative divergence. CLC near zero indicates mechanical regularity. IKI values for human writing typically range from 0.3 to 0.8; values consistently near 1.0 suggest random content insertion, values near 0.0 suggest verbatim copying.

5.3. Mechanical Turk Scoring (Informative)

Indicators of mechanical turk behavior include:

* Paste-to-keystroke ratio exceeding 0.7 across a session.

* Burst insertion: more than 200 characters appearing in under 2 seconds, characteristic of clipboard paste operations.

* Low IKI variance: pasted content with uniformly high compressibility (IKI below 0.2), consistent with LLM-generated prose.

* Absence of cognitive pause patterns before and after complex sentences.

* Temporal clustering: paste events at regular intervals suggesting a prompt-copy-paste workflow.

Verifiers SHOULD compute a mechanical turk probability score from 0.0 (no indicators) to 1.0 (all indicators present). A score exceeding 0.6 SHOULD trigger a recommendation for tool receipt documentation.

5.4. Error Topology Model (Informative)

Error topology analysis constructs a directed graph of error and correction patterns. The error graph G = (V, E) has vertices V representing edit operations and edges E representing temporal succession. Human error topology exhibits:

* Power-law distribution of error cluster sizes.

* Short-range temporal locality (errors corrected within 5 seconds).

* Increasing error rates at cognitive load boundaries (end of paragraphs, section transitions).

* Fractal self-similarity in revision patterns.

Simulated error injection produces uniform error distribution, regular correction intervals, and no correlation between error rates and structural boundaries. A graph clustering coefficient below 0.1 combined with uniform correction latency is flagged as potentially synthetic.

6. Forgery Cost Bounds (Quantified Security)

Forgery cost bounds provide a Verifier with a lower bound on the computational resources required to forge an Evidence Packet. The cost (_C_total_) is computed as:

```
C_total = C_swf + C_entropy + C_hardware
```

6.1. Sequential Work Function Cost (C_swf)

The SWF cost component provides a lower bound on the computational time an adversary must expend:

```
C_swf >= n * t_checkpoint

where:
  n = number of checkpoints in the Evidence chain
  t_checkpoint = wall-clock time for one SWF computation
```

The memory-hard nature of Argon2id ensures that an adversary with k parallel processors achieves at most O(sqrt(k)) speedup due to memory bandwidth constraints. The minimum forgery time equals the sum of SWF claimed-durations across all checkpoints. At T1 tier without hardware binding, C_swf represents an economic cost only (the adversary must spend real time, but has no hardware constraint).

6.2. Behavioral Evidence Synthesis Cost (C_entropy)

The entropy cost component estimates the resources required to synthesize behavioral noise satisfying all forensic constraints:

```
C_entropy = O(d * n * log(1/epsilon))

where:
  d = number of independent forensic dimensions
  n = number of checkpoints
  epsilon = target false-negative rate
```

At T1/T2, only basic entropy and timing are checked (d = 2). For T3/T4, the full forensic assessment applies (d >= 7, including CLC, IKI, error topology, SNR dynamics, session consistency, and cross-checkpoint correlation), making synthesis exponentially more expensive in the number of correlated dimensions the adversary must simultaneously satisfy.

The cost of synthesizing behavioral noise that satisfies all forensic constraints is inherently uncertain and depends on adversary capability. Verifiers SHOULD set C_entropy conservatively. When the Verifier cannot independently assess AI synthesis costs, C_entropy SHOULD be set to 0 and the WAR warnings field SHOULD note that entropy cost was not estimated.

6.3. Hardware Attestation Cost (C_hardware)

* _T1/T2:_ C_hardware = 0. No hardware root of trust; keys are software-managed.

* _T3 (Hardware-Bound):_ Requires compromise of TPM or platform Secure Element. Estimated cost: USD 10,000-100,000 per device class, depending on the specific hardware and attack methodology.

* _T4 (Hardware-Hardened):_ Requires invasive hardware attacks, manufacturer collusion, or firmware exploits targeting PUF-bound keys. Estimated cost: USD 100,000 or more.

Verifiers MUST include these estimates in the WAR to allow Relying Parties to set trust thresholds based on objective economic risk.

The c-total field in the forgery-cost-estimate MUST equal the sum of c-swf, c-entropy, and c-hardware. All component costs within a single forgery-cost-estimate MUST be expressed in the same cost-unit.

7. Absence Proofs: Negative Evidence Taxonomy

Absence proofs assert that certain events did NOT occur during the monitored session. They are divided into categories based on verifiability:

Type 1: Computationally-Bound Claims Verifiable from the Evidence Packet alone (e.g., "Max single delta size < 500 bytes" or "No checkpoint timestamps out of order").

Type 2: Monitoring-Dependent Claims Require trust in the AE’s event monitoring (e.g., "No paste from unauthorized AI tool" or "No clipboard activity detected"). Trust in these claims MUST be weighted by the declared Attestation Tier (T1-T4).

Type 3: Environmental Claims Assertions about the execution environment (e.g., "No debugger attached" or "Hardware temperature remained within stable physical bounds").

Type 1 (Computationally-Bound) claims MUST be verified computationally by the Verifier from the Evidence Packet data alone. Type 3 (Environmental) claims SHOULD be evaluated against physical-state markers when present, and MUST be treated as unverifiable when physical-state is absent.

8. Attestation Result Wire Format

The Writers Authenticity Report (WAR) is a CBOR-encoded [RFC8949] Attestation Result identified by semantic tag 1463894560 (encoding ASCII "WAR "). The CDDL notation [RFC8610] defines the wire format:

```
pop-war = #6.1463894560(attestation-result)

attestation-result = {
    1 => uint,                    ; version (MUST be 1)
    2 => hash-value,              ; evidence-ref
    3 => verdict,                 ; appraisal verdict
    4 => attestation-tier,        ; assessed assurance level
    5 => uint,                    ; chain-length
    6 => uint,                    ; chain-duration (seconds)
    ? 7 => entropy-report,        ; entropy assessment (omit for CORE)
    ? 8 => forgery-cost-estimate, ; quantified forgery cost
    ? 9 => [+ absence-claim],     ; absence claims (1+ when present)
    ? 10 => [* tstr],             ; warnings
    11 => bstr,                   ; verifier-signature (COSE_Sign1)
    12 => pop-timestamp,           ; created (appraisal timestamp)
    * int => any,                 ; extension fields
}

verdict = &(
    authentic: 1,                 ; consistent with human authorship
    inconclusive: 2,              ; insufficient evidence
    suspicious: 3,                ; anomalies detected
    invalid: 4,                   ; chain broken or forged
)

entropy-report = {
    1 => float32,                 ; timing-entropy (bits/sample)
    2 => float32,                 ; revision-entropy (bits)
    3 => float32,                 ; pause-entropy (bits)
    4 => bool,                    ; meets-threshold
}
```

```
  forgery-cost-estimate = {
      1 => float32,                ; c-swf
      2 => float32,                ; c-entropy
      3 => float32,                ; c-hardware
      4 => float32,                ; c-total
      5 => cost-unit,              ; currency
  }

  cost-unit = &(
      usd: 1,
      cpu-hours: 2,
  )

  absence-claim = {
      1 => absence-type,           ; proof category
      2 => time-window,            ; claimed window
      3 => tstr,                   ; claim-id
      ? 4 => any,                  ; threshold/parameter
      5 => bool,                   ; assertion
  }

  absence-type = &(
      computationally-bound: 1,    ; verifiable from Evidence alone
      monitoring-dependent: 2,     ; requires trust in AE monitoring
      environmental: 3,            ; environmental assertions
  )

  time-window = {
      1 => pop-timestamp,          ; start
      2 => pop-timestamp,          ; end
  }

  ; Shared type definitions reproduced from [PoP-Protocol] for reader
  ; convenience. In case of conflict, [PoP-Protocol] is authoritative.
  pop-timestamp = #6.1(float32)      ; CBOR tag 1 (epoch-based, float32)
  hash-value = {
      1 => hash-algorithm,
      2 => bstr,
  }
  hash-algorithm = &(
      sha256: 1,
      sha384: 2,
      sha512: 3,
  )
  attestation-tier = &(
      software-only: 1,            ; T1: AAL1
      attested-software: 2,        ; T2: AAL2
      hardware-bound: 3,           ; T3: AAL3
```

```
    hardware-hardened: 4,         ; T4: LoA4
  )
```

The evidence-ref field MUST contain a hash-value computed as SHA-256 over the CBOR-encoded evidence-packet structure (including CBOR tag 1347571280), excluding any COSE_Sign1 wrapper. This binds the Attestation Result to a specific Evidence Packet.

In the absence-claim structure, claim-id is a unique textual identifier for the claim (e.g., "no-paste-event", "max-delta-below-500"). The assertion field is true if the claim holds and false if the Verifier determined it does not hold. The time-window specifies the temporal scope of the claim within the Evidence Packet's session.

When appraising CORE Evidence Packets that lack jitter-binding data, the Verifier SHOULD omit the entropy-report field from the Attestation Result and include a warning indicating that behavioral entropy analysis was not performed.

The created field (key 12) MUST contain the timestamp at which the Verifier completed the appraisal. Relying Parties use this field to evaluate the freshness of the Attestation Result.

8.1. Entropy Report Computation

The Verifier MUST compute entropy-report fields as follows:

timing-entropy: Shannon entropy of quantized jitter intervals across all checkpoints, expressed in bits per sample.

revision-entropy: Shannon entropy of edit-delta sizes (chars-added values) across all checkpoints, expressed in bits.

pause-entropy: Shannon entropy of inter-checkpoint pause durations, expressed in bits.

meets-threshold: True if and only if timing-entropy is at or above the minimum threshold (3.0 bits per sample) AND revision-entropy is at or above 3.0 bits AND pause-entropy is at or above 2.0 bits. These thresholds are calibrated for the NIST SP 800-90B most common value estimator. Implementations using alternative entropy estimators MUST provide equivalent assurance levels.

8.2. Verdict Assignment

The Verifier MUST assign the verdict based on the appraisal outcome:

authentic (1): All verification steps passed. Evidence is

consistent with human authorship. No forensic flags triggered.

inconclusive (2): Verification steps passed but insufficient behavioral data available for forensic assessment (e.g., CORE profile without jitter-binding).

suspicious (3): One or more forensic flags triggered (low entropy, failed CLC correlation, mechanical pacing detected) but chain integrity is intact. When multiple forensic checks produce contradictory results, the Verifier MUST assign the more conservative verdict (suspicious over authentic).

invalid (4): Chain integrity broken, SWF verification failed, or structural validation error. Evidence cannot be trusted.

9. Tool Receipt Protocol (AI Attribution)

NOTE: This section is informational. The complete CDDL wire format for Tool Receipts, including signature algorithms and binding mechanisms, will be specified in a future revision. Implementations SHOULD treat this section as guidance only.

When external tools (LLMs) contribute content, the framework enables a "compositional provenance" model:

1. Receipt Signing: The Tool signs a "Receipt" containing its tool_id, an output_commit (SHA-256 hash of generated text), and an optional input_ref (SHA-256 hash of the prompt).

2. Binding: The human Attester records a PASTE event in the transcript referencing the Tool Receipt’s output_commit.

3. Countersigning: The Attester binds the Receipt into the next human-driven checkpoint, anchoring the automated work into the linear human effort.

Verifiers appraise the ratio of human-to-machine effort based on these receipts and the intervening SWF-proved intervals.

10. Adversary Model

This document inherits the adversary model defined in the Threat Model section of [PoP-Protocol]. The appraisal procedures defined herein assume the adversarial Attester capabilities and constraints specified there. The primary threat is an adversarial Attester -- an author who controls the Attesting Environment and seeks to generate Evidence for content they did not authentically author.

The following adversary tiers characterize the appraisal-specific threat landscape. Each tier defines the adversary capabilities that the corresponding Attestation Tier is designed to resist:

Tier 1 Adversary (Casual): Can manipulate system clocks and intercept local IPC. Cannot perform real-time behavioral simulation exceeding basic cadence matching. The T1 appraisal policy accepts the risk of basic retype attacks; SWF time-binding provides the primary defense.

Tier 2 Adversary (Motivated): Can invest computational resources up to the cost of a high-end workstation and study the verification algorithm to craft evidence targeting specific thresholds. The T2 appraisal policy defends through multi-dimensional behavioral analysis (SNR + CLC + mechanical turk detection).

Tier 3 Adversary (Professional): Has access to custom hardware (FPGAs, specialized ASICs) for SWF acceleration and sophisticated behavioral models trained on human authorship data. The T3 appraisal policy defends through HAT cross-validation and advanced forensic metrics (CLC, IKI, error topology, and SNR dynamics).

Tier 4 Adversary (Nation-State): Has all Tier 3 capabilities plus: can potentially compromise hardware manufacturer endorsement chains, deploy large-scale parallel computation, and employ teams of human operators for sophisticated retype attacks. The T4 appraisal policy defends through the combined cost of SWF sequentiality, multi-dimensional behavioral evidence synthesis (d >= 7 correlated dimensions), and hardware attestation integrity. Even a Tier 4 adversary faces a minimum forgery cost equal to the claimed authorship duration plus the hardware compromise cost.

11. Privacy Considerations

11.1. Evidence and Attestation Result Privacy

High-resolution behavioral data poses a stylometric de-anonymization risk [Goodman2007]. Implementations SHOULD support Evidence Quantization, reducing timing resolution to a level that maintains forensic confidence while breaking unique author fingerprints.

The entropy-report in Attestation Results (timing-entropy, revision-entropy, pause-entropy) may enable cross-document author identification by Relying Parties. Verifiers SHOULD quantize entropy-report values to reduce fingerprinting precision while preserving forensic utility. Relying Parties MUST NOT correlate entropy reports across multiple Attestation Results to identify or track authors.

11.2. Evidence Quantization Requirements

Attestation Results MUST quantize forensic indicator values to the following resolutions:

* Cadence (IKI) values: millisecond resolution. Sub-millisecond data MUST NOT be included.

* Entropy values: 0.01 bit resolution (two decimal places).

* SNR values: 0.5 dB resolution.

* CLC and IKI metric values: two decimal places.

These quantization levels are calibrated to preserve the forensic utility of all assessment mechanisms defined in Section 5 while limiting the precision available for stylometric fingerprinting.

11.3. Data Retention and Behavioral Profiles

Verifiers MUST NOT maintain per-author behavioral profile databases. Attestation Results SHOULD NOT include raw forensic indicator values; tier-level pass/fail determinations are sufficient for Relying Parties. Evidence retention SHOULD NOT exceed 90 days (the default validity period). Implementations SHOULD support anonymous Evidence submission to prevent linking authorship sessions to real-world identities.

12. Accessibility and Assistive Modes

Verifiers MUST NOT automatically reject evidence based solely on atypical timing patterns. Implementations MUST support "Assistive Modes" that adjust SNR and CLC thresholds for authors with motor disabilities or those using assistive technologies (eye-tracking, dictation).

To signal assistive mode usage, the Attester SHOULD include an assistive-mode indicator in the profile-declaration structure of the Evidence Packet. When this indicator is present, Verifiers MUST apply adjusted thresholds as follows:

12.1. Eye-Tracking Mode

Eye-tracking input produces IKI ranges of 500-3000 ms (versus 100-300 ms for keyboard). Adjusted thresholds:

* Entropy: 2.0 to 4.0 bits/sample (reduced from 3.0 minimum)

* SNR: -5 dB to +5 dB (narrower than keyboard range). SNR anomaly threshold: +15 dB.

* CLC correlation: r > 0.1 (reduced from r > 0.2)

* Error topology: Adjusted for gaze drift corrections, which produce characteristic error patterns distinct from keyboard errors.

12.2. Dictation Mode

Dictation input produces burst patterns with higher cadence variance than keyboard. Adjusted thresholds:

* SNR: -8 dB to +8 dB (wider range reflecting speech pauses)

* CLC correlation: r > 0.1 (range 0.1 to 0.8)

* Paste-to-keystroke ratio threshold: disabled (dictation engines produce burst insertions by design)

* Error topology: waived (dictation corrections follow speech-recognition patterns, not typing patterns)

12.3. Additional Accommodations

* Switch-access input: minimum event count per checkpoint reduced to 1 (from default of 5).

* Head-tracking and mouth-stick input: apply eye-tracking thresholds.

* When assistive mode thresholds produce anomalous results, the Verifier SHOULD flag the inconsistency in the WAR warnings rather than reject the Evidence.

The WAR MUST indicate when assistive mode thresholds were applied. Assistive mode is signaled through the profile-declaration structure in the Evidence Packet. Implementations MAY include an assistive-mode feature flag (value 60) in the feature-flags array. The following values are defined: 0 (none), 1 (motor-disability), 2 (eye-tracking), 3 (dictation). A future revision of [PoP-Protocol] will formalize this signaling mechanism.

13. IANA Considerations

This document has no IANA actions. All IANA registrations for the PoP framework are defined in [PoP-Protocol].

14. Security Considerations

This document defines forensic appraisal procedures that inherit and extend the security model from [PoP-Protocol]. The broader RATS security considerations [Sardar-RATS] also apply. Implementers should consider the following security aspects:

14.1. Entropy Manipulation Attacks

An adversary may attempt to inject synthetic jitter patterns that satisfy entropy thresholds while lacking biological origin. The use of multi-dimensional analysis (SNR, CLC, Error Topology) rather than single metrics provides defense-in-depth against high-fidelity simulation.

14.2. Verifier Trust Model

The forensic assessments defined in this document produce probabilistic confidence scores, not binary determinations. Relying Parties MUST understand that forgery cost bounds represent economic estimates, not cryptographic guarantees. Trust decisions SHOULD incorporate the declared Attestation Tier (T1-T4) and the specific absence proof types claimed.

14.3. Stylometric De-anonymization

High-resolution behavioral data (keystroke timing, pause patterns) can enable author identification even when document content is not disclosed. Implementations SHOULD support Evidence Quantization to reduce timing resolution while maintaining forensic utility. The trade-off between forensic confidence and privacy should be documented for Relying Parties.

14.4. Assistive Mode Abuse

Adversaries may falsely claim assistive technology usage to bypass behavioral entropy checks. Verifiers SHOULD require consistent assistive mode declarations across sessions and MAY request additional out-of-band verification for mode changes. The WAR should indicate when assistive modes were active, as specified in the accessibility section above.

15. References

15.1. Normative References

[PoP-Protocol]
Condrey, D., "Proof of Process (PoP): Architecture and Evidence Format", Work in Progress, Internet-Draft, draft-condrey-rats-pop-protocol-05, February 2026, <https://datatracker.ietf.org/doc/html/draft-condrey-rats-pop-protocol-05>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC5869] Krawczyk, H. and P. Eronen, "HMAC-based Extract-and-Expand Key Derivation Function (HKDF)", RFC 5869, DOI 10.17487/RFC5869, May 2010, <https://www.rfc-editor.org/info/rfc5869>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC8610] Birkholz, H., Vigano, C., and C. Bormann, "Concise Data Definition Language (CDDL): A Notational Convention to Express Concise Binary Object Representation (CBOR) and JSON Data Structures", RFC 8610, DOI 10.17487/RFC8610, June 2019, <https://www.rfc-editor.org/info/rfc8610>.

[RFC8949] Bormann, C. and P. Hoffman, "Concise Binary Object Representation (CBOR)", STD 94, RFC 8949, DOI 10.17487/RFC8949, December 2020, <https://www.rfc-editor.org/info/rfc8949>.

[RFC9334] Birkholz, H., Thaler, D., Richardson, M., Smith, N., and W. Pan, "Remote ATtestation procedureS (RATS) Architecture", RFC 9334, DOI 10.17487/RFC9334, January 2023, <https://www.rfc-editor.org/info/rfc9334>.

15.2. Informative References

[Goodman2007]
Goodman, A. and V. Zabala, "Using Stylometry for Biometric Keystroke Dynamics", 2007, <https://doi.org/10.1007/978-3-540-77343-6_14>.

[Monrose2000]
Monrose, F. and A. Rubin, "Keystroke dynamics as a biometric for authentication", 2000, <https://doi.org/10.1145/351427.351438>.

[RFC9052] Schaad, J., "CBOR Object Signing and Encryption (COSE): Structures and Process", STD 96, RFC 9052, DOI 10.17487/RFC9052, August 2022, <https://www.rfc-editor.org/info/rfc9052>.

[Salthouse1986] Salthouse, T.A., "Perceptual, Cognitive, and Motoric Aspects of Transcription Typing", Psychological Review 93(3), 303-319, 1986, <https://doi.org/10.1037/0033-295X.93.3.303>.

[Sardar-RATS] Sardar, M.U., "Security Considerations for Remote ATtestation procedureS (RATS)", Work in Progress, Internet-Draft, draft-sardar-rats-sec-cons-02, February 2026, <https://datatracker.ietf.org/doc/html/draft-sardar-rats-sec-cons-02>.

[SEAT-EXPAT] Sardar, M.U., Fossati, T., Reddy, T., Sheffer, Y., Tschofenig, H., and I. Mihalcea, "Remote Attestation with Exported Authenticators", Work in Progress, Internet-Draft, draft-fossati-seat-expat-01, January 2026, <https://datatracker.ietf.org/doc/html/draft-fossati-seat-expat-01>.

[SEAT-UseCases] Mihalcea, I., Sardar, M.U., Fossati, T., Reddy, T., Jiang, Y., and M. Chen, "Use Cases and Properties for Integrating Remote Attestation with Secure Channel Protocols", Work in Progress, Internet-Draft, draft-mihalcea-seat-use-cases-01, January 2026, <https://datatracker.ietf.org/doc/html/draft-mihalcea-seat-use-cases-01>.

Verification Constraint Summary

The following constraints summarize the verification requirements defined in the preceding sections:

Structural Integrity

1. Chain Integrity: SHA-256 hash chain is unbroken from genesis to final checkpoint.

2. Temporal Monotonicity: All checkpoint timestamps strictly exceed their predecessors.

3. SWF Continuity: Recompute Argon2id from seed; verify sampled Merkle proofs.

4. Content Binding: Final document hash matches document-ref in Evidence Packet.

Behavioral Analysis (ENHANCED/MAXIMUM profiles)

1. Entropy Threshold: Independent entropy estimate >= 3.0 bits per inter-keystroke interval per checkpoint.

2. SNR Analysis: Jitter exhibits characteristic biological noise patterns, not periodic or spectrally flat patterns.

3. CLC Correlation: Semantic complexity correlates with timing (r > 0.2, or r > 0.1 for assistive mode).

4. Error Topology: Correction patterns consistent with human cognitive processing.

5. Mechanical Turk Detection: No robotic pacing (machine-clocked editing rate).

Absence Proof Validation

1. Type 1 Claims: Verify computationally from Evidence Packet (delta sizes, timestamp ordering).

2. Type 2 Claims: Weight by Attestation Tier (T1-T4).

3. Type 3 Claims: Evaluate environmental assertions against physical-state markers.

Tool Receipt Validation (when present)

1. Verify Tool signature over Receipt.

2. Verify PASTE event references correct output_commit.

3. Calculate human-to-machine effort ratio from SWF-proved intervals.

Per-Tier Verification Constraints

This appendix summarizes the verification thresholds and constraints for each Attestation Tier. These values are the normative defaults; deployment profiles MAY adjust them within the ranges specified.

T1 (Software-Only) Constraints

* Chain integrity: prev-hash linkage required.

* Temporal ordering: monotonic timestamps required; SWF claimed-duration within [0.5x, 3.0x] of expected time.

* Entropy: minimum 3.0 bits/sample when jitter-binding is present. No upper bound enforced.

* Entanglement: jitter seal presence required when jitter-binding is present; HMAC verification SHOULD be performed (MAC key derivable from public merkle-root).

* State matching: final content hash match required.

* Forensic assessment: SNR computation OPTIONAL; CLC, error topology, and mechanical turk detection RECOMMENDED for ENHANCED+ profiles.

* Forgery cost bound: C_total = C_swf + C_entropy (no hardware component). Physical-state fields are self-reported and provide no additional assurance.

T2 (Attested Software) Constraints

* Chain integrity: all T1 requirements.

* Temporal ordering: all T1 requirements; SWF claimed-duration within [0.5x, 3.0x] of expected time on reference hardware.

* Entropy: 3.0 to 6.0 bits/sample per checkpoint. Values above 6.0 suggest injected randomness and SHOULD be flagged.

* Entanglement: jitter seal and entangled-mac presence required for ENHANCED+ profiles. HMAC verification SHOULD be performed.

* State matching: final content hash match required; intermediate content hash progression SHOULD be verified for monotonic growth.

* Forensic assessment: SNR, CLC, and mechanical turk detection required. Error topology OPTIONAL.

* Forgery cost bound: C_total = C_swf + C_entropy. Minimum forgery time equals the sum of SWF claimed-durations.

T3 (Hardware-Bound) Constraints

* Chain integrity: all T2 requirements. COSE_Sign1 signature MUST verify against hardware-bound key.

* Temporal ordering: all T2 requirements; HAT delta cross-validation SHOULD be performed when TPM monotonic counter data is available.

* Entropy: 3.0 to 5.5 bits/sample, reflecting tighter calibration against verified human authorship baselines.

* Entanglement: HMAC verification MUST be performed. Device attestation certificate chain SHOULD be validated against known Endorser roots.

* State matching: all T2 requirements; intermediate content hash progression MUST be verified for monotonic growth. Non-monotonic changes (document size decreasing by more than 50% between consecutive checkpoints) MUST be flagged.

* Forensic assessment: all T2 requirements plus error topology analysis required. QR presence challenge OPTIONAL.

* Forgery cost bound: C_total = C_swf + C_entropy + C_hardware. Hardware compromise cost estimated at USD 10,000-100,000.

T4 (Hardware-Hardened) Constraints

* Chain integrity: all T3 requirements.

* Temporal ordering: all T3 requirements; HAT delta cross-validation MUST be performed; HAT-SWF agreement within 5% tolerance required.

* Entropy: 3.0 to 5.0 bits/sample; entropy trajectory standard deviation MUST exceed 0.1 bits across the session. A constant-entropy session is a strong indicator of synthetic generation.

* Entanglement: all T3 requirements; timing vector entropy consistency check required (within 0.5 bits of reported entropy-estimate).

* State matching: all T3 requirements.

* Forensic assessment: all T3 requirements; cross-correlation analysis between entropy and SNR required. QR presence challenge RECOMMENDED.

* Forgery cost bound: C_total = C_swf + C_entropy + C_hardware. Hardware compromise cost estimated at USD 100,000 or more. Total minimum forgery cost exceeds sum of claimed-durations plus hardware procurement.

Acknowledgements

The author thanks the participants of the RATS working group for their ongoing work on remote attestation architecture and security considerations that informed this specification.

Author’s Address

David Condrey
WritersLogic Inc
San Diego, California
United States
Email: david@writerslogic.com