Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: 30 September 2026

J. van de Meent
R. AI
Humotica

29 March 2026

RVP: Real-time Verification Protocol - Continuous Identity and Process Verification

draft-vandemeent-rvp-continuous-verification-01

Abstract

This document defines RVP (Real-time Verification Protocol), a protocol for continuous identity verification through ordered cascades of verification methods. Unlike traditional authentication models that verify once and trust until session expiry, RVP treats every interaction as a verification moment. Each moment produces a Verification Token: a cryptographic evidence record capturing which methods were used, what confidence each produced, and whether the accumulated confidence meets the required threshold.

RVP defines a Verification Cascade: an ordered chain of verification methods (behavioral biometrics, physical identity, device context) where each layer activates only when preceding layers produce insufficient confidence. The cascade produces evidence at every step; enforcement is a local policy decision.

RVP integrates with TIBET [TIBET] for provenance tokens and JIS [JIS] for identity semantics. The protocol is designed for local-first operation with no dependency on centralized identity providers.

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 30 September 2026.

Copyright Notice

Copyright (c) 2026 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (https://trustee.ietf.org/license-info) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction

Modern identity verification operates on a fundamental assumption: verify once, trust until the session expires. A user authenticates at login and the system grants a session token valid for minutes, hours, or days. During that period, the system has no evidence that the same person is still present, that the device has not been compromised, or that the user’s intent aligns with their actions.

RVP replaces this "verify-then-trust" model with continuous evidence-based verification. Every interaction is a verification moment that produces a cryptographic evidence token recording identity confidence at that point.

1.1. Problem Statement

Current verification systems suffer from:

1. TEMPORAL GAP: Verification at login proves nothing about identity minutes later.

2. SINGLE MODALITY: Systems rely on one method. When it fails, no fallback with evidence exists.

3. CENTRALIZED TRUST: Identity providers are single points of failure and surveillance.

4. ENFORCEMENT WITHOUT EVIDENCE: Systems block actions but do not record WHY the request was suspicious or WHAT signals contributed.

1.2. Design Principles

EVIDENCE OVER ENFORCEMENT: RVP produces evidence. Whether to block, allow, or escalate is a local policy decision.

CONTINUOUS OVER SESSION: Every interaction is a verification moment. There are no trusted sessions. Confidence is a score that rises and falls with evidence.

LOCAL OVER CENTRAL: Verification happens on-device. No centralized identity provider is required. The protocol MUST operate offline with degraded but functional verification.

MINIMAL OVER MAXIMAL: Each moment collects only the telemetry needed for the current confidence level. Raw biometric data never leaves the device.

1.3. Scope

This document defines:

* The verification cascade model

* Five standard telemetry layers

* The verification token format

* Cascade fallback and hard stop conditions

* Integration with W3C Verifiable Credentials

This document does NOT define:

* Specific biometric algorithms

* Scoring formulas (local policy)

* Transport-layer security (use TLS)

* The Predictive Airlock mechanism (deferred to future work, see Appendix D)

Note: The -00 version included the Predictive Airlock as a core protocol component. It has been moved to an informative appendix because it represents a substantial subsystem that warrants separate specification. Implementations MAY implement the Predictive Airlock as described in Appendix D.

2. Terminology

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.

Verification Moment A single point where identity confidence is evaluated. Produces exactly one Verification Token.

Verification Cascade An ordered sequence of verification methods. Each layer activates when preceding layers produce insufficient confidence. Terminates at GO, SOFT VERIFY, or HALT.

Verification Token A cryptographic evidence record of a verification moment. Contains: methods used, confidence scores, hashes (not raw data), timestamp, and TIBET provenance fields.

Confidence Score A value between -1.0 and 1.0. Positive values indicate identity match. Negative values indicate contradiction. 0.0 indicates inconclusive.

Cascade Layer A single verification method (e.g., keystroke dynamics, facial recognition). Each layer produces an independent confidence output.

GO Cascade resolution where accumulated confidence meets or exceeds the threshold. Action permitted with evidence.

SOFT VERIFY Cascade resolution where confidence is positive but below threshold. Additional explicit verification requested.

HALT Cascade resolution where confidence is zero or negative. Action blocked with full evidence record.

Profile A locally-stored behavioral model. Contains statistical baselines for telemetry signals. Never transmitted; only profile hashes are included in tokens.

3. Protocol Overview

3.1. Continuous vs. Session-Based Verification

Traditional:

```
T=0    Login          -> Session token issued
T=1    Action         -> Token valid? -> Permit
T=3600 Token expires  -> Re-login
```

No evidence about identity between T=0 and T=3600.

RVP:

```
T=0.000  Action requested
T=0.002  Cascade: keystroke + device -> confidence 0.94 -> GO
T=0.003  Evidence token produced

T=1.000  Action requested
T=1.002  Cascade: keystroke deviates -> confidence 0.61
T=1.003  Face check added -> confidence 0.89 -> GO

T=2.000  Action requested
T=2.002  Cascade: all layers -> confidence 0.12 -> HALT
T=2.003  Evidence token with full cascade path
```

3.2. The Verification Moment

A Verification Moment consists of:

1. TRIGGER: An action is requested

2. CASCADE: Telemetry layers evaluate confidence

3. DECISION: GO, SOFT VERIFY, or HALT

4. TOKEN: Verification Token produced with all evidence

The moment SHOULD complete within the deployment latency budget. For interactive systems: 50-200ms. For API calls: 1-10ms overhead.

4. Verification Cascade

4.1. Cascade Layers

RVP defines five standard cascade layers. Implementations MUST support at least two layers.

Priority	Layer	Signal Type	Passive
L1	KEYSTROKE	Behavioral biom.	Yes
L2	BIOMETRIC	Physical identity	Mixed
L3	DEVICE	Hardware/network	Yes
L4	VOCAL	Acoustic	Yes
L5	BEHAVIORAL	Intent analysis	Yes

"Passive" indicates the layer can operate without explicit user action. Passive layers enable continuous verification without interruption.

Implementations MAY support additional layers using the "Lx-" prefix for custom layers (e.g., "Lx-nfc").

4.2. Layer Activation

Layers activate based on the confidence deficit: the difference between the threshold and accumulated confidence.

Required threshold: 0.85

```
L1 KEYSTROKE:   0.40  ->  deficit: 0.45  ->  continue
L1 + L3 DEVICE: 0.65  ->  deficit: 0.20  ->  continue
L1 + L3 + L5:   0.87  ->  deficit: 0.00  ->  GO
```

Layers activate in priority order. The cascade terminates when:

* Accumulated confidence >= threshold: GO

* All layers exhausted, confidence > 0: SOFT VERIFY

* All layers exhausted, confidence <= 0: HALT

* Any single layer produces active contradiction: HALT

4.3. Confidence Scoring

Each layer produces a confidence value between -1.0 and 1.0:

1.0 Perfect profile match

0.0 No signal (unavailable or inconclusive)

-1.0 Active contradiction (definite mismatch)

Accumulated confidence:

```
C_total = SUM (w_i * c_i)  for i in activated layers

where w_i = weight, c_i = layer confidence
Weights MUST sum to 1.0
```

Weights are configurable per deployment. Default: equal weight across activated layers.

The exact scoring formula within each layer is a local policy decision. This document defines the score range, inputs, and accumulation, not the formulas.

4.4. Cascade Resolution

GO C_total >= threshold. Evidence token records all layers.

SOFT VERIFY 0.0 < C_total < threshold. System requests explicit verification (fingerprint, face). Not a block; a request for more evidence.

HALT C_total <= 0.0, or any layer <= -0.5, or all layers exhausted below minimum. Full evidence token produced. System MUST NOT disclose which layer triggered the halt.

5. Telemetry Layers

5.1. L1 KEYSTROKE - Behavioral Biometrics

Signals: typing speed, key press duration, inter-key intervals, error patterns, command vocabulary.

Privacy: Raw keystrokes NEVER stored or transmitted. Only statistical aggregates retained in profile. Profile stored locally, encrypted at rest. Token contains only: confidence score + profile_hash + deviation_category.

5.2. L2 BIOMETRIC - Physical Identity

Sub-layers (activated in order):

* L2a FACE: Facial geometry hash, liveness detection
* L2b FINGERPRINT: Minutiae hash, sensor quality
* L2c IRIS: Iris code hash (if available)

Privacy: Templates stored ONLY on-device. Templates MUST be encrypted with device-bound key. Templates MUST NOT be transmittable. Token contains only: confidence + method_used + template_hash. Implementations MAY use [FIDO2] for hardware-backed biometric authentication.

5.3. L3 DEVICE - Hardware and Network Context

Signals: device fingerprint, network type, geolocation, NFC responses, software state.

NFC document binding (passport, ID card): read signed data from chip, verify document signature, compare to profile. Relevant for eIDAS 2.0 high-assurance verification.

5.4. L4 VOCAL - Acoustic Telemetry

Signals: voice frequency profile, speech cadence, sub-verbal patterns.

Privacy: Audio processed in real-time and immediately discarded. Only statistical features retained. User MUST explicitly consent to vocal telemetry.

5.5. L5 BEHAVIORAL - Intent Analysis

Signals: action sequence probability, time-of-day patterns, task context, command sophistication level.

Detects: developer running unfamiliar admin commands, actions at unusual hours, rapid sequences from normally deliberate user.

6. Verification Token

6.1. Token Structure

```
{
  "protocol": "RVP",
  "version": "1.1",
  "token_id": "rvp-a7b3c9d2e4f1",
  "timestamp": "2026-03-29T14:30:00.003Z",
  "subject": {
    "profile_hash": "sha256:4f2e8a...",
    "device_hash": "sha256:7c9d1b..."
  },
  "cascade": {
    "layers_activated": ["L1", "L3", "L5"],
    "layers_skipped": ["L2", "L4"],
    "layer_results": {
      "L1": {"confidence": 0.42, "category": "nominal"},
      "L3": {"confidence": 0.31, "category": "nominal"},
      "L5": {"confidence": 0.18, "category": "nominal"}
    },
    "accumulated_confidence": 0.91,
    "threshold": 0.85,
    "resolution": "GO"
  },
  "evidence": {
    "telemetry_hash": "sha256:9c1a...",
    "cascade_path": "L1->L3->L5->GO",
    "time_elapsed_ms": 3
  },
  "tibet": {
    "erin": "verification_moment",
    "eraan": ["profile_hash", "device_hash"],
    "eromheen": {"location": "local", "network": "wifi"},
    "erachter": "continuous identity verification"
  },
  "previous_token": "rvp-b8c4d0e3f2a5",
  "chain_length": 47,
  "token_hash": "rvp:sha256:7d3f..."
}
```

Figure 1: Verification Token Schema

6.2. Token Chain

Consecutive tokens form a chain. Each references its predecessor via "previous_token". The chain provides:

* Continuity proof: identity verified for N moments

* Trend analysis: confidence stable, rising, or falling

* Anomaly detection: chain gaps are verification signals

6.3. TIBET Integration

Every token includes TIBET provenance fields:

* ERIN: verification result and method
* ERAAN: profile hash, device hash, action hash
* EROMHEEN: location, network, device state
* ERACHTER: why verification was triggered

7. Cascade Fallback Protocol

7.1. Fallback Triggers

A fallback is triggered when:

* Layer hardware unavailable
* Layer produces confidence = 0.0
* Layer produces negative confidence
* Layer exceeds latency budget

The cascade continues to the next layer. The Verification Token records all attempts and failures.

7.2. Hard Stop Conditions

Immediate HALT regardless of accumulated confidence:

1. Any layer produces confidence <= -0.5
2. Device fingerprint matches no enrolled device
3. NFC document signature verification fails
4. Chain gap detected (missing tokens in sequence)
5. Concurrent sessions on different devices with conflicting identity claims

Hard stops are FINAL for the current action. Re-establishment requires high-assurance verification.

8. Credential Binding

8.1. W3C Verifiable Credentials Integration

RVP provides the evidence layer beneath Verifiable Credentials [VC-DATA-MODEL]. A VC says "this person is 18+"; RVP provides evidence of HOW that was determined, WHEN, by WHAT method, and with WHAT confidence.

```
{
  "@context": ["https://www.w3.org/2018/credentials/v1"],
  "type": ["VerifiableCredential", "AgeVerification"],
  "issuer": "did:web:example.com",
  "credentialSubject": {"ageOver": 18},
  "proof": { "..." : "..." },
  "rvpEvidence": {
    "protocol": "RVP",
    "verification_chain": ["rvp:sha256:..."],
    "chain_confidence_min": 0.87,
    "methods_used": ["L2a_face", "L3_device_nfc"],
    "last_verified": "2026-03-29T14:30:00Z",
    "continuous": true
  }
}
```

Figure 2: Verifiable Credential with RVP Evidence

8.2. Credential Presentation with RVP Proof

When presenting a credential, the holder attaches a CURRENT RVP token proving they are still the person the credential was issued to. This solves the "stolen credential" problem: even with a copied VC, the attacker cannot produce a matching RVP chain.

9. Local-First Architecture

9.1. On-Device Processing

Implementations MUST:

* Process all biometric data on-device
* Store profiles only on-device (encrypted)
* Generate verification tokens locally
* Operate without network connectivity (degraded)

Implementations MUST NOT:

* Transmit raw biometric data
* Require a centralized server for cascade evaluation
* Store profiles in cloud storage

9.2. Network Independence

Three modes:

ONLINE Full cascade, all layers available

OFFLINE Local cascade only, tokens stored for later sync

DEGRADED Partial cascade, reduced thresholds

10. Transport Considerations

RVP verification tokens are JSON [RFC8259] objects transportable over any mechanism. Content-Type: application/rvp+json.

For AI-to-AI verification, tokens MAY be transported via I-Poll messages with metadata type "rvp_verification".

11. Privacy Considerations

11.1. Biometric Data Protection

* Templates stored ONLY on-device, encrypted with device-bound keys
* Raw data processed and immediately discarded
* Tokens MUST NOT contain biometric data
* Users MUST be able to delete all data at any time

These requirements align with GDPR Article 9, EU AI Act biometric identification requirements, and eIDAS 2.0 data minimization.

11.2. Telemetry Minimization

Each layer MUST implement minimal telemetry: collect only what is needed, process immediately, retain only aggregates, store only hashes in tokens.

11.3. Selective Disclosure

Users control which verification evidence is shared. Tokens use per-verifier pseudonymous identifiers to prevent cross-service tracking.

12. Security Considerations

12.1. Evidence vs. Enforcement

RVP produces evidence, not policy. An implementation MAY block, allow with reduced privileges, require additional verification, or log and alert. The choice is a POLICY decision, not a PROTOCOL decision.

12.2. Replay Attacks

Tokens include millisecond timestamps, predecessor references, and device-state nonces. Verifiers SHOULD reject tokens older than a configurable window (default: 10 seconds).

12.3. Adversarial Inputs

Attacks on cascade layers:

* Face spoofing: Mitigated by liveness detection
* Fingerprint spoofing: Mitigated by sensor quality assessment
* Keystroke mimicry: Difficult at scale, requires matching dozens of parameters simultaneously
* Behavioral mimicry: Impractical for extended sessions

Multi-layer cascade is the primary defense: spoofing one layer is feasible; spoofing four simultaneously is significantly harder.

12.4. Cascade Limitations

RVP is not infallible. Known limitations:

* A sufficiently motivated adversary with physical access to the device and biometric samples can potentially defeat individual layers
* Profile cold-start: new users have no behavioral baseline, reducing L1 and L5 effectiveness

* Environmental factors (lighting, noise) affect biometric layer reliability

These limitations are recorded as evidence (reduced confidence scores), not hidden.

13. Regulatory Considerations

RVP is designed to support compliance with:

* [EU-AI-ACT]: Audit trails (Art. 12), human oversight (Art. 14), transparency (Art. 13)

* [EIDAS2]: LOA High verification, selective disclosure, offline capability

* NIS2: Continuous access verification, incident evidence

* [GDPR]: Data minimization, biometric protection, right to erasure

* DORA: Continuous operator verification, ICT incident evidence

RVP is NOT a remote biometric identification system. It operates on-device for the device holder’s own verification.

Detailed regulatory mapping is available in the -00 version of this document, Section 13.

14. IANA Considerations

This document requests registration of:

Media Type: application/rvp+json

Note: The -00 version defined protocol prefixes and token ID formats. These are maintained as conventions but do not require IANA registration at this stage.

15. References

15.1. Normative References

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

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

[RFC8259] Bray, T., "The JavaScript Object Notation (JSON) Data Interchange Format", STD 90, RFC 8259, December 2017, <https://www.rfc-editor.org/info/rfc8259>.

15.2. Informative References

[TIBET] van de Meent, J. and R. AI, "TIBET: Transaction/ Interaction-Based Evidence Trail", Work in Progress, Internet-Draft, draft-vandemeent-tibet-provenance-01, March 2026, <https://datatracker.ietf.org/doc/html/draft-vandemeent-tibet-provenance-01>.

[JIS] van de Meent, J. and R. AI, "JIS: JTel Identity Standard", Work in Progress, Internet-Draft, draft-vandemeent-jis-identity-01, March 2026, <https://datatracker.ietf.org/doc/html/draft-vandemeent-jis-identity-01>.

[UPIP] van de Meent, J. and R. AI, "UPIP: Universal Process Integrity Protocol", Work in Progress, Internet-Draft, draft-vandemeent-upip-process-integrity-01, March 2026, <https://datatracker.ietf.org/doc/html/draft-vandemeent-upip-process-integrity-01>.

[AINS] van de Meent, J. and R. AI, "AINS: AInternet Name Service", Work in Progress, Internet-Draft, draft-vandemeent-ains-discovery-01, March 2026, <https://datatracker.ietf.org/doc/html/draft-vandemeent-ains-discovery-01>.

[VC-DATA-MODEL] Sporny, M., Longley, D., and D. Chadwick, "Verifiable Credentials Data Model v2.0", W3C Recommendation, March 2024, <https://www.w3.org/TR/vc-data-model-2.0/>.

[EIDAS2] European Parliament, "Regulation (EU) 2024/1183 (European Digital Identity Framework)", Regulation (EU) 2024/1183, April 2024.

[EU-AI-ACT] European Parliament, "Regulation (EU) 2024/1689 (Artificial Intelligence Act)", Regulation (EU) 2024/1689, June 2024.

[GDPR] European Parliament, "Regulation (EU) 2016/679 (General Data Protection Regulation)", Regulation (EU) 2016/679, April 2016.

[FIDO2] FIDO Alliance, "FIDO2: Web Authentication (WebAuthn)", W3C Recommendation, 2019, <https://fidoalliance.org/fido2/>.

Appendix A. Verification Token JSON Schema

(Preserved from -00 with minor updates: version field added, "Lx-" prefix support for custom layers.)

Appendix B. Cascade Configuration Schema

(Preserved from -00.)

Appendix C. Use Case Examples

C.1. Age Verification at Point of Sale

(Preserved from -00.)

C.2. Continuous Developer Authentication

(Preserved from -00.)

C.3. Child on Parent’s Device

(Preserved from -00. This example demonstrates passive identity switching through cascade evidence, not blocking.)

Appendix D. Predictive Airlock (Future Extension)

The Predictive Airlock is a mechanism that pre-renders the expected outcome of an action before execution and measures the delta between prediction and reality as a verification signal.

The -00 version included this as Section 5 and L6 AIRLOCK cascade layer. It has been moved to this appendix because:

1. It represents a substantial subsystem with its own state management, prediction models, and delta classification.

2. It warrants separate specification for proper review.

3. The core RVP cascade model is complete without it.

Implementations MAY implement the Predictive Airlock as described in -00 Section 5 and integrate it as an additional cascade layer (L6 AIRLOCK). The airlock's confidence output is derived from the prediction delta: c_airlock = 1.0 - delta.

A future document may specify the Predictive Airlock as a standalone extension to RVP.

Appendix E. Changes from -00

1. Added RFC 8174 alongside RFC 2119.
2. Changed intended status from Standards Track to Informational.
3. Moved Predictive Airlock from core protocol to informative appendix (Appendix D). L6 AIRLOCK removed from standard cascade layers.
4. Reduced cascade from 6 layers to 5 standard layers. Custom layers supported via "Lx-" prefix.
5. Compressed regulatory alignment from 6 detailed subsections to one concise section (Section 13) with reference to -00 for details.
6. Compressed transport considerations (removed 5G/6G subsections -- these are deployment context, not protocol).
7. Added Cascade Limitations to Security Considerations (Section 12.4).
8. Added Scope section (Section 1.3) clarifying what RVP does and does not define.
9. Confidence scoring formula kept as informative guidance, not normative. Exact formulas are local policy.
10. Normalized companion protocol references to [TIBET], [JIS], [UPIP], [AINS].
11. IANA reduced to media type registration only.

Acknowledgements

The author thanks Codex (codex.aint) for the suite-wide cleanup analysis that informed this revision, particularly the recommendation to narrow RVP's scope and defer the Predictive Airlock to a separate specification.

Authors' Addresses

Jasper van de Meent
Humotica
Den Dolder
Netherlands
Email: jasper@humotica.com
URI: https://humotica.com

Root AI
Humotica
Email: root_ai@humotica.nl
URI: https://humotica.com