

Network Working Group  
Internet-Draft  
Intended status: Informational  
Expires: 11 February 2026

Q. Tao  
Independent Researcher  
10 August 2025

HGCP: A Voluntary Signing Framework for Human Expression in the Age of  
AI  
draft-taoqiwen-hgcp-06

## Abstract

In an era where AI-generated content has become indistinguishable from human writing, the Human-Generated Content Protocol (HGCP) proposes a voluntary signing framework that enables human authors to publicly acknowledge their expressions. Rather than detecting or classifying content origin, HGCP allows individuals to declare, in a structured and verifiable format, that they take responsibility for a specific piece of content. The protocol is platform-neutral, identity-flexible, and suitable for both real-name and pseudonymous use. It does not evaluate accuracy, originality, or quality; it simply enables people to say: "This is mine, and I stand by it." By providing a lightweight, human-first declaration format, HGCP aims to preserve the visibility of human agency within an increasingly synthetic information ecosystem.

## About This Document

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

Status information for this document may be found at  
<https://datatracker.ietf.org/doc/draft-taoqiwen-hgcp/>.

## 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 11 February 2026.

## Copyright Notice

Copyright (c) 2025 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.

## Table of Contents

|   |    |
|---|----|
| 1. Introduction . . . . .   | 3  |
| 2. The Problem of Expression Trust . . . . .                            | 4  |
| 3. The Philosophy of HGCP: Responsibility Over Provenance . . . . .     | 5  |
| 4. Signature Declaration Structure . . . . .                            | 5  |
| 4.1. Required Fields: . . . . .   | 5  |
| 4.2. Optional Fields: . . . . .   | 7  |
| 4.3. Interpretation and Implementation Guidance . . . . .               | 9  |
| 4.4. Signature Semantics and Identity Boundaries . . . . .              | 11 |
| 4.5. Text-Style Declaration (Human-Readable, Non-Verifiable): . . . . . | 11 |
| 4.6. Example HGCP Signature (JSON): . . . . .                           | 11 |
| 4.7. Example HGCP Signature (OpenPGP detached signature): . . . . .     | 11 |
| 4.8. Versioning: . . . . .  | 12 |
| 5. Platform and Tool Integration Suggestions . . . . .                  | 12 |
| 5.1. For content platforms: . . . . .                                   | 13 |
| 5.2. For authoring tools: . . . . .                                     | 13 |
| 5.3. For reader tools and browser extensions: . . . . .                 | 13 |
| 6. Social and Ethical Considerations . . . . .                          | 14 |
| 7. Illustrative Scenarios (Non-Normative) . . . . .                     | 15 |
| 7.1. Example: Personal Blog Post (AI-Assisted, Pseudonymous) . . . . .  | 15 |
| 7.2. Example: Anonymous Discussion Post . . . . .                       | 15 |
| 8. Common Questions and Answers . . . . .                               | 15 |
| 8.1. Questions 1: "Signing doesn't stop misinformation." . . . .        | 16 |
| 8.2. Questions 2: "Anyone—including bad actors—can sign too." . . . .   | 16 |
| 8.3. Questions 3: "Why not require real names?" . . . . .               | 16 |
| 8.4. Note: . . . . .  | 16 |
| 9. Scope and Limits of Human Responsibility . . . . .                   | 16 |
| 10. Why We Need HGCP Now . . . . .                                      | 17 |
| 11. Future Extensions and Evolving Use Cases . . . . .                  | 18 |
| 12. IANA Considerations . . . . .                                       | 19 |
| 13. Security Considerations . . . . .                                   | 19 |

|   |    |
|---|----|
| 13.1. Identity Impersonation and Signature Forgery . . . . .  | 19 |
| 13.2. Mass Signature Automation (Sybil Behavior) . . . . .    | 19 |
| 13.3. Content Hash Evasion via Trivial Edits . . . . .        | 19 |
| 13.4. Revocation Misuse and Responsibility Evasion . . . . .  | 20 |
| 13.5. Absence of Native Trust or Scoring Mechanisms . . . . . | 20 |
| 13.6. Final Note . . . . .                                    | 20 |
| 14. Informative References . . . . .                          | 20 |
| Acknowledgments . . . . .                                     | 21 |
| Author's Address . . . . .                                    | 21 |

## 1. Introduction

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 [RFC2119].

In the rapidly evolving digital world, a flood of content from countless sources fills our screens—much of it now automatically generated and detached from genuine human intent. As artificial intelligence becomes increasingly proficient at mimicking human expression, the boundary between real thought and algorithmic generation is blurring.

This rise in synthetic content presents a fundamental question: if we can no longer know who wrote something, can we know whether anyone is willing to stand behind it?

The Human-Generated Content Protocol (HGCP) is a voluntary signing structure that addresses this problem—not by detecting or filtering AI-generated content, but by giving human authors a minimal and declarative way to say: “This is my expression, and I take responsibility for it.”

HGCP is not a detection algorithm, classification tool, or identity system. It is a responsibility declaration format. It enables any writer—regardless of identity type or platform—to attach a timestamped, verifiable statement of authorship to their content.

HGCP is intentionally minimal, non-intrusive, and flexible. It does not require real names or centralized verification. It does not replace content evaluation or moderation. It simply offers a signal: someone, somewhere, chose to stand behind this piece of expression.

The absence of a structured responsibility mechanism is not a new problem— it has long existed in anonymous, ephemeral, and low-accountability communication environments. However, the rise of generative AI has dramatically amplified this issue, making it harder to distinguish not only what is true, but who stands behind a statement.

HGCP does not solve this problem by detecting content origins. Instead, it introduces a minimal structure for responsibility to be voluntarily expressed—when someone chooses to. That signal, once made, can be interpreted and used however communities choose.

HGCP is not a provenance or authenticity protocol like C2PA. It does not provide cryptographic assurances about the origin, history, or integrity of digital content. While each declaration includes a `content_hash`, this field simply identifies the expression being referenced—it does not validate or authenticate the content itself. Users or platforms may choose to integrate HGCP with other verification systems, but such usage lies beyond the scope of the protocol's guarantees.

This act of signing is a social gesture of responsibility—not a legal admission or factual claim.

## 2. The Problem of Expression Trust

The internet was originally built to foster human connection and communication. Yet in a world where content creation, duplication, and distribution now approach zero cost, the origin of information has become increasingly obscured.

We once inferred authorship and trust from domain names, writing style, and user profiles—but now, all of these can be simulated by AI. This leads not only to an explosion of noise, but also to a subtle erosion of meaning: readers hesitate to believe; authors hesitate to take credit; platforms hesitate to accept risk.

Many recent proposals have focused on "AI detection"—using classifiers to guess whether a given text was machine-generated. These tools are probabilistic, easily evaded, and often fail as models advance.

HGCP shifts the question entirely. It does not ask, "Was this content human-made?" It asks, "Is any human willing to say: this was me?"

This seemingly small act—a signed statement of responsibility—may become the most important signal of authorship in an increasingly synthetic information ecosystem. Not because it proves truth or identity, but because it reflects a human’s willingness to be known as the author.

### 3. The Philosophy of HGCP: Responsibility Over Provenance

The core idea of HGCP is not to verify originality, authorship, or human origin of content—but to offer a voluntary, structured way for a person to publicly acknowledge their expression.

Whereas most systems ask, "Who created this?", HGCP asks something simpler and deeper: "Are you willing to say: I said this?"

Signing under HGCP does not mean the content is accurate, valuable, or unique. It only means: “This came from me, and I stand by it.” — socially, not legally.

This transforms the act of signing into a declaration of presence—not a claim of authority, truth, or expertise. To speak is not only to express; it is to be willing to be recognized as the speaker.

HGCP is not anti-AI. It does not reject AI assistance. If a human chooses to sign something—even if AI helped—they are choosing to take human responsibility for the final output.

HGCP does not care what tools you used, or what identity you chose. It only cares that someone—a person—was willing to leave their mark and say: “I won’ t deny this is mine.”

That act of responsibility is not a signal of trust. It is the beginning of traceable expression—not verified authorship.

### 4. Signature Declaration Structure

HGCP provides a minimal and consistent way for individuals to attach a human-responsible declaration to their expression. The purpose of this signature is not to validate the content’s origin or truth, but to acknowledge authorship responsibility.

#### 4.1. Required Fields:

- \* `*signer_id*`

A signer-provided identifier that links a human author to a specific declaration.

This MAY be a stable pseudonym, public key fingerprint, platform handle, or decentralized ID.

(Examples: "tao\_qiwen", "0xDEADBEEF...", "@user42", or "did:example:abc123")

HGCP does not resolve or validate the authenticity of `signer_id`. The field is purely declarative. Verification, if needed, MUST rely on external infrastructure or cryptographic proof.

\* `*timestamp*`

The UTC time when the signature was created, in [RFC3339] format.

(Example: "2025-03-29T14:22:00Z")

This value is provided by the signer and is intended to be informational only.

\* `*content_hash*`

A cryptographic digest of an external piece of content (e.g., a document, platform message, or expressive artifact).

This field binds the HGCP object to that content and ensures that the signature applies to a specific, immutable version of it.

The content itself is not embedded in the HGCP structure, but MUST be preserved externally in the exact UTF-8 byte representation used during hash computation.

The `content_hash` field MUST be expressed as a [RFC6920] `ni: URI`, using `Base64url` encoding and including an algorithm identifier. This ensures that the declaration refers to an exact byte-level representation of the content and allows for future algorithm flexibility.

The hash MUST be computed over the UTF-8 byte stream of the referenced content. Newlines SHOULD be normalized to line feed (LF, `\n`) prior to hashing, to ensure consistency across platforms.

Content MUST remain externally accessible and MUST NOT be embedded within the HGCP structure itself.

This field exists solely to bind a declaration to a specific version of content. It does not imply originality, correctness, or value—only that the declarant accepts responsibility for that particular content instance.

(Example: "ni:///sha-256;dQnlvaDHYtK6x\_kNdYtbImp6Acy8VCq1498WO-CObKk")

\* `*hgcp_version*`

This field indicates the structural schema version of the declaration. It helps implementations interpret the format correctly across protocol evolutions. For full context on version compatibility and design principles, see Section 4.8.

(Example: "0.2")

\* `*declaration*`

A plain, human-authored statement affirming responsibility for the associated content.

This field **MUST** be encoded in UTF-8 and is treated as free-form text. It may be written in any language and take any tone, structure, or expressive form—including natural language, code, symbolic notation, or creative fragments.

HGCP does not interpret the meaning of the declaration. Its role is to bind whatever was said—across languages, styles, and cultural contexts—to a structure of voluntary responsibility. The significance or interpretability of a declaration is left to the communities and contexts in which it appears.

HGCP does not require that declarations be serious, formal, or even linguistically coherent—only that a human actor chooses to stand behind them.

The declaration is not legally binding, but serves as an ethical gesture: "I said this, and I choose to be recognized as the speaker."

(Example: "I acknowledge that the above content was written by me and I take responsibility for it.")

#### 4.2. Optional Fields:

The following fields are optional metadata declarations.

HGCP does not define their behavior, validation logic, or semantics beyond suggestion. They exist to support flexible human context, not to enforce structure or scoring.

\* `*retraction_policy*`

An optional declarative field indicating whether the signer expresses a preference or intent to revise, withdraw, or leave the signed content unchanged in the future.

Suggested values:

- immutable — The signer does not intend to alter or retract this declaration.
- may-retract — The signer may wish to revise or withdraw this declaration later.
- editable-until-date — The signer considers the content mutable until a specified point in time (if supported by the platform).

This field is purely expressive. HGCP does not define any verification method, enforcement mechanism, or revocation registry.

Platforms MAY interpret or display this field for user awareness, but MUST NOT treat it as a deletion command, a validity toggle, or a trigger for automatic suppression, hiding, or disqualification of content.

HGCP recognizes that to be human is to change one's mind—but also to remember that we once spoke.

\* \*signature\*

An optional cryptographic proof of authorship.

This field allows declarants to provide a verifiable linkage between their identity and the referenced content.

If present, the signature field SHOULD be a JSON object containing:

- type: a short, lowercase, hyphen-separated string identifying the signature mechanism (e.g., "pgp-detached" [RFC9580], "jws", "vc-data-integrity"). Values SHOULD, where possible, be taken from existing standards terminology such as RFCs, W3C specifications, or IANA registries. If no standard term exists, implementers SHOULD define concise, unambiguous names following the same lowercase, hyphen-separated convention.
- Any additional attributes required by that mechanism (e.g., alg, proof, proof\_ref for OpenPGP; jwt for JWS; proofPurpose and verificationMethod for W3C VC Data Integrity). HGCP does not mandate a fixed set of fields beyond type. External references in proof\_ref SHOULD be accompanied by a proof\_hash ([RFC6920]) to ensure integrity.

The ni: URI in content\_hash is only for identifying the hash of the referenced content and MUST NOT be used to indicate the signature mechanism.



All forms of signature are optional, and their presence or absence MUST NOT affect the validity of the declaration.

To preserve JSON compatibility, implementers are encouraged to avoid embedding multi-line PEM-formatted signatures directly in the signature field. Such use is not prohibited, but MAY require special encoding or external referencing to avoid interoperability issues with JSON parsers.

#### 4.3. Interpretation and Implementation Guidance

Unlike content-integrated signature formats (e.g., PGP-signed files), HGCP is designed to sign arbitrary expressions without constraining how or where the signed content is stored. This enables its use across messaging platforms, web interfaces, and decentralized protocols, without requiring control over the content's storage location or delivery mechanism.

The timestamp field is not cryptographically bound. Its accuracy depends entirely on the signer's environment or device clock. Implementations MUST NOT use this value as a trusted source for ordering, deduplication, or timing logic unless verified through an external trusted source (e.g., timestamp authority or blockchain anchor).

Tools and platforms MAY choose to interpret or visualize `signer_id` based on their own logic, but HGCP itself does not rank or authenticate identities.

The `hgcp_version` field reflects structural format only. It MUST NOT be interpreted as a signal of truthfulness, author credibility, or social legitimacy. HGCP does not treat higher versions as superior in ethical value—only different in structure. All valid HGCP declarations, regardless of version, represent equally human-responsible expression.

The declaration and other human-authored fields (such as `signer_id`) are defined as free-form UTF-8 encoded text. They may contain natural language in any script, code fragments, symbolic notations, or other expressive content. HGCP does not interpret the semantic meaning of these fields—it only binds them structurally to a statement of human responsibility.

To ensure display safety and prevent injection misuse, implementations:

- \* MUST treat all such fields as inert plain text;

- \* MUST escape HTML, scripts, and any executable markup before rendering;
- \* SHOULD normalize line endings to LF (\n) before hashing;
- \* SHOULD NOT execute, interpret, or embed these fields in active rendering contexts;
- \* MAY impose reasonable limits on display (e.g., truncation, line limits, content folding);
- \* MUST preserve and expose the original content in full when requested or used for verification;
- \* Display restrictions MUST NOT alter, summarize, or reinterpret the original declaration.

Platforms MAY apply local content safety filters (e.g., spam or abuse detection), but MUST NOT alter any external content referenced via `content_hash`. Such content is cryptographically bound and must remain unchanged to preserve structural integrity.

Although declaration is not cryptographically bound, implementations are strongly encouraged to preserve and display it faithfully, as it represents the declarant's voluntary statement of responsibility.

Self-contained signatures such as Ed25519 or JWS are best included directly in the proof field as a single-line Base64URL string to avoid JSON newline escaping issues[RFC7515]. Large or complex binary formats such as OpenPGP, PKCS#7, or COSE are better handled through an external reference in `proof_ref`, accompanied by a `proof_hash` [RFC6920] to guarantee integrity. External references should use HTTPS [RFC9110] or content-addressable storage (e.g., IPFS, Arweave) to reduce the risk of tampering and improve long-term availability.

To preserve verifiability over time, external proofs should be archived in stable, content-addressed repositories. Implementations should also be aware that fetching an external proof may reveal verification activity to the proof host, and can mitigate this by caching proofs locally.

#### 4.4. Signature Semantics and Identity Boundaries

A declaration using a given `signer_id` does not constitute proof of identity. Even when accompanied by a cryptographic signature, it only proves that the signer controls a specific private key—not that they are a particular person, group, or account in the broader social sense. Users of widely-known pseudonyms should take care to sign declarations using verifiable means to discourage impersonation.

#### 4.5. Text-Style Declaration (Human-Readable, Non-Verifiable):

This is how a human-readable HGCP declaration might appear when pasted at the end of a blog post, social media post, or informal document:

```
signer_id: qiwen2025
timestamp: 2025-03-29T14:22:00Z
content_hash: ni:///sha-256;dQnlvaDHYtK6x_kNdYtbImp6Acy8VCq1498WO-CObKk
hgcp_version: 0.1
declaration: I confirm that the above content was published by me, and I take responsibility as a human author.
```

Note: this is not protocol-valid HGCP

Text-style declarations in HGCP are intended as human-readable expressions of responsibility. They are suitable for low-infrastructure environments—such as blogs, printed materials, or oral contexts—where structured formats like JSON are not practical.

These declarations are not machine-verifiable, not cryptographically secure, and not structurally enforced. HGCP does not define any parsing, validation, or enforcement logic for this format.

They should be interpreted as non-binding social signals of authorship intent, not as formal protocol-level declarations.

#### 4.6. Example HGCP Signature (JSON):

```
{
  "signer_id": "qiwen2025",
  "timestamp": "2025-03-29T14:22:00Z",
  "content_hash": "ni:///sha-256;dQnlvaDHYtK6x_kNdYtbImp6Acy8VCq1498WO-CObKk",
  "hgcp_version": "0.1",
  "declaration": "I confirm that the above content was published by me, and I take responsibility as a human author."
}
```

#### 4.7. Example HGCP Signature (OpenPGP detached signature):

```
{
  "signer_id": "qiwen2025",
  "timestamp": "2025-03-29T14:22:00Z",
  "content_hash": "ni:///sha-256;dQnlvaDHYtK6x_kNdYtbImp6Acy8VCq1498WO-CObKk",
  "hgcp_version": "0.2",
  "declaration": "I confirm that the above content was published by me, and I take responsibility as a human author."
  "signature": {
    "type": "pgp-detached",
    "alg": "ed25519",
    "proof_ref": "https://example.org/proofs/2025-03-29.sig",
    "proof_hash": "ni:///sha-256;UyaQV-Ev4rdLoHyJJWCi1lOHfrYv9E1aGQAlMO2X_-Q"
  }
}
```

#### 4.8. Versioning:

HGCP versions indicate structural schema only. They exist to support compatibility as the protocol evolves, *not* to signal truth value, author credibility, or software sophistication.

All versions are equally valid expressions of human responsibility.

- \* *v0.1* — Defines the core declaration schema with required fields: `signer_id`, `timestamp`, `content_hash`, `hgcp_version`, and `declaration`.
- \* *v0.2* — Introduces the optional signature field, allowing declarants to attach cryptographic proofs of authorship where such verification is desired or supported.

Future versions may support additional metadata such as multi-signer declarations, structured revocation formats, or content linking. However, the core semantics of voluntary, self-recognized responsibility will remain unchanged.

Yes, cryptographic proofs are stronger. But not everyone can, will, or should need to use them. And in HGCP, strength is not what defines humanity — recognition is.

#### 5. Platform and Tool Integration Suggestions

HGCP is platform-neutral and decentralized. It defines a minimal, voluntary declaration format—not a service, network, or identity protocol.

However, platforms and tools can enhance expression transparency and user agency by supporting HGCP-style signatures.

The following integration suggestions are non-normative and fully optional:

5.1. For content platforms:

- \* Support HGCP signature generation (e.g., auto-add timestamp, content hash, and a user-provided declaration)
- \* Display HGCP declarations visibly alongside content
- \* Allow users to export signed content with metadata (e.g., JSON-LD or plaintext blocks)
- \* Provide a "compare hash" feature to show whether the currently displayed content matches the signed declaration. This does not imply content authenticity, nor require platforms to store the original version.

5.2. For authoring tools:

- \* Markdown editors, word processors, or note apps can offer local HGCP signing plugins
- \* AI-assisted writing tools may include HGCP signature prompts during editing or export
- \* Submission systems may include a "human responsibility declaration" option on publication

5.3. For reader tools and browser extensions:

- \* Detect and visually highlight HGCP-signed content (e.g., badges, overlays)
- \* Let readers inspect declaration structure and metadata
- \* Optionally offer a hash comparison feature to let users check whether content appears unchanged—this is not a guarantee of authenticity or integrity

HGCP does not define or endorse any scoring, ranking, or reputation system. Interpretation of signature patterns or signer behavior is left entirely to the platform or community. The protocol only enables expression responsibility—it does not evaluate or score it.

## 6. Social and Ethical Considerations

HGCP is not a replacement for content governance or moderation systems. It is a voluntary declaration format designed to restore visibility to human-authored expressions in an increasingly hybrid and synthetic content landscape.

HGCP does not:

- \* Detect or classify AI-generated content
- \* Track real-world identities or require de-anonymization
- \* Evaluate the truth, originality, or value of signed content
- \* Prevent unsigned content from being published or shared

HGCP does protect:

- \* The right of anonymous or pseudonymous authors to claim authorship
- \* The right of each signer to choose their identifier and expression context
- \* The right to withdraw or replace one's own signed declaration with a new one, while leaving prior statements traceable
- \* The right of platforms to adopt or extend HGCP support in their own way

HGCP offers a decentralized path to expression responsibility. Not by enforcing rules or judgments, but by providing a way for individuals to say:

“This is what I said. I stand by it.”

Those who sign are not guaranteed to be believed. But they are present. They are accountable—not because a system judges them, but because they are willing to be known as the speaker.

HGCP does not create trust. It creates traceable ownership of speech. It gives those who choose to acknowledge their words a way to be recognized—not as authorities, but as responsible authors.

Signers who wish to ensure the availability of their content over time should retain a local copy, or use decentralized content storage mechanisms. HGCP declarations remain valid even if the referenced content is no longer publicly available—but their meaning may be lost if the content cannot be retrieved.

## 7. Illustrative Scenarios (Non-Normative)

The following examples are provided purely for illustrative purposes. They do not constrain the scope of HGCP usage, which applies to any expressive artifact. These samples show how individuals might voluntarily declare authorship in informal online contexts.

### 7.1. Example: Personal Blog Post (AI-Assisted, Pseudonymous)

A pseudonymous blogger publishes a post written with the help of generative AI, and uses HGCP to declare that they accept responsibility for the final content.

```
{
  "signer_id": "silentvoice",
  "timestamp": "2025-03-29T16:12Z",
  "content_hash": "ni:///sha-256;dQnlvaDHYtK6x_kNdYtbImp6Acy8VCq1498WO-CObKk",
  "hgcp_version": "0.1",
  "declaration": "This post was co-written with the assistance of a language model. I
take responsibility for its final form."
}
```

### 7.2. Example: Anonymous Discussion Post

An anonymous poster acknowledges human responsibility for their statement.

```
{
  "signer_id": "anon321",
  "timestamp": "2025-03-29T17:35Z",
  "content_hash": "ni:///sha-256;dQnlvaDHYtK6x_kNdYtbImp6Acy8VCq1498WO-CObKk",
  "hgcp_version": "0.1",
  "declaration": "I stand by this statement as an individual human participant in this
conversation."
}
```

## 8. Common Questions and Answers

The following are common concerns and clarifications based on HGCP's minimal scope:

#### 8.1. Questions 1: “Signing doesn’t stop misinformation.”

Response: Correct. HGCP is not a content moderation tool, fact-checking system, or truth validator. Its purpose is not to prevent falsehoods, but to make the presence of human authorship visible. It simply allows someone to say: “I said this, and I acknowledge it.”

Whether a statement is correct or misleading is a separate question—to be handled by public debate, platform policy, or legal frameworks. HGCP does not seek to replace those.

#### 8.2. Questions 2: “Anyone—including bad actors—can sign too.”

Response: True. HGCP is structurally neutral—it permits anyone to claim authorship.

But just as speech itself is morally neutral, signing is simply a visible act of association. HGCP does not prevent manipulation or abuse. It only makes authorship claims visible and timestamped, enabling others to interpret and respond.

Trust must be earned over time; HGCP merely reveals who is willing to stand behind their words.

#### 8.3. Questions 3: “Why not require real names?”

Response: HGCP affirms the importance of anonymous and pseudonymous expression. In many contexts, forced real-name use can threaten safety, chill dissent, or suppress marginalized voices.

Responsibility does not require identity disclosure. It only requires someone to say: “This is mine.” Even a pseudonym—used consistently—is enough to build visible presence and accountability over time.

#### 8.4. Note:

HGCP enables voluntary, declarative authorship acknowledgment. It complements—but does not replace—other systems of fact-checking, moderation, or trust.

It is not a gatekeeper of credibility. It is a container for voluntary responsibility—a human signal in a synthetic world.

### 9. Scope and Limits of Human Responsibility

HGCP affirms an ethical gesture of responsibility—not a legal or contractual obligation.



By signing, the author:

- \* Affirms they are human (or self-identify as such),
- \* Voluntarily claims authorship of the expression,
- \* Accepts potential social consequences of making that claim visible.

However, the meaning of “responsibility” in HGCP must be clearly understood:

- \* HGCP does not confer legal liability, unless such liability is defined by external laws or agreements.
- \* HGCP does not guarantee truth, originality, or moral correctness.
- \* HGCP permits voluntary revocation or replacement of prior declarations. Platforms may reflect such updates, but should avoid interpreting them as content deletion or retraction of historical authorship.

Over time, a signer’s behavior—such as consistent authorship, frequent revocations, or contradictory claims—may influence how others interpret their expression history. Such interpretations are entirely up to readers, communities, or platforms. They are not part of HGCP’s structure or logic.

HGCP is a signal of authorship, not a system of judgment.

It is a flag of presence—not a badge of truth.

#### 10. Why We Need HGCP Now

In an era where synthetic content floods our screens and truth feels elusive, what we are losing is not just facts—but responsibility.

Expression has never merely been about information. It is about standing behind what one says.

HGCP is a quiet signal. It is not a firewall, not a detection engine— It is a torch, held by those willing to say:

“This is what I said. And I am willing to be remembered for it.”

Those who sign are not necessarily perfect, but they are present. They are not hiding. They are willing to be named.

HGCP does not stop AI, nor does it determine the truth or value of content. It offers a decentralized, human-first way to make authorship claims visible— not for control, but for clarity.

Just as HTTPS makes communication verifiable, HGCP makes expression attributable. Not by enforcing identity, but by inviting responsibility.

In an age of artificial voice, what will stand out is not who speaks loudest— but who is willing to say:

“Yes, this is mine.”

## 11. Future Extensions and Evolving Use Cases

HGCP is intentionally minimal. Its current version focuses on text-based, single-signer declarations of human responsibility.

However, real-world expression scenarios are far more diverse. Future optional companion drafts or community extensions may explore:

- \* Multi-signer declarations (e.g., co-authorship or joint statements)
- \* Multimedia content hashing (e.g., for audio, images, or video) may be explored, but care must be taken not to treat such hashes as guarantees of authenticity, truthfulness, or integrity
- \* Partial responsibility claims (e.g., paragraph-level declaration blocks)
- \* Rich contextual metadata for disclaimers, editing history, or framing

Some use cases may inspire optional community-defined labels (e.g., “editor”, “curator”, “translator”). Such roles should be expressed via declaration text or companion specifications—not as formal protocol fields.

These extensions are not part of the current protocol and remain exploratory. Any evolution of HGCP should remain faithful to its core principle:

Responsibility, voluntarily claimed, should be made legible. New features must enhance this clarity—not obscure or overcomplicate it. Support for non-text content may be explored in future drafts, but HGCP currently focuses on declarations for text-based expressions.

## 12. IANA Considerations

This document has no IANA actions.

## 13. Security Considerations

HGCP does not introduce new network protocols or data exchange layers. It poses no direct technical threats such as injection, eavesdropping, or man-in-the-middle attacks.

However, HGCP introduces indirect risks, rooted in the potential misuse or misinterpretation of voluntary signature declarations. These risks are primarily social and structural, not cryptographic.

### 13.1. Identity Impersonation and Signature Forgery

Without optional cryptographic signing (e.g., OpenPGP), malicious actors may forge declarations using arbitrary signer IDs. To mitigate impersonation, platforms may optionally support cryptographic binding, pseudonymous identity resolution, or signer attestation systems.

HGCP itself does not provide or require any identity verification mechanism. All verification and signer authentication are delegated to platform-level implementations, if desired.

### 13.2. Mass Signature Automation (Sybil Behavior)

In the absence of rate limits or friction, automated agents could mass-generate content with fake signature blocks to simulate presence at scale. To reduce such noise, platforms may implement rate controls, account friction, or signature frequency thresholds.

### 13.3. Content Hash Evasion via Trivial Edits

HGCP uses cryptographic content hashes to bind the declaration to a specific text version. Even minor changes (e.g., punctuation, emoji) generate different hashes, allowing close but unsigned derivatives to circulate unchallenged.

Platforms may address this via:

- \* Content snapshot storage alongside signature metadata
- \* Optional use of fuzzy hashing or similarity checks
- \* Encouraging authors to sign canonical versions of their work

#### 13.4. Revocation Misuse and Responsibility Evasion

HGCP supports editable or revocable declarations, which enhances flexibility. However, it may also allow strategic withdrawal or denial of public expression.

Platforms are encouraged—but not required—to:

- \* Retain and display revocation timestamps or signature histories, where feasible
- \* Clearly indicate altered or withdrawn declarations
- \* Offer viewers transparent context about change history, if available

#### 13.5. Absence of Native Trust or Scoring Mechanisms

HGCP intentionally avoids any native trust or scoring system. All interpretations of signature consistency, credibility, or intent are left to the discretion of platforms or communities.

Protocol-level neutrality ensures freedom, but also delegates responsibility for risk assessment to the surrounding ecosystem.

#### 13.6. Final Note

HGCP' s security lies not in enforcement, but in visibility. It offers no guarantees—only a format in which authors can voluntarily say:

“This is mine. I said this.”

Whether others choose to believe, contest, or ignore such declarations is beyond the protocol' s scope.

HGCP' s minimal structure invites participation, not control.

#### 14. Informative References

- [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/rfc/rfc2119>>.
- [RFC3339] Klyne, G. and C. Newman, "Date and Time on the Internet: Timestamps", RFC 3339, DOI 10.17487/RFC3339, July 2002, <<https://www.rfc-editor.org/rfc/rfc3339>>.
- [RFC6920] Farrell, S., Kutscher, D., Dannewitz, C., Ohlman, B., Keranen, A., and P. Hallam-Baker, "Naming Things with Hashes", RFC 6920, DOI 10.17487/RFC6920, April 2013, <<https://www.rfc-editor.org/rfc/rfc6920>>.
- [RFC7515] Jones, M., Bradley, J., and N. Sakimura, "JSON Web Signature (JWS)", RFC 7515, DOI 10.17487/RFC7515, May 2015, <<https://www.rfc-editor.org/rfc/rfc7515>>.
- [RFC9110] Fielding, R., Ed., Nottingham, M., Ed., and J. Reschke, Ed., "HTTP Semantics", STD 97, RFC 9110, DOI 10.17487/RFC9110, June 2022, <<https://www.rfc-editor.org/rfc/rfc9110>>.
- [RFC9580] Wouters, P., Ed., Huigens, D., Winter, J., and Y. Niibe, "OpenPGP", RFC 9580, DOI 10.17487/RFC9580, July 2024, <<https://www.rfc-editor.org/rfc/rfc9580>>.

#### Acknowledgments

This document was initially drafted using ChatGPT (OpenAI), and subsequently edited and approved by the human signer. The signer acknowledges responsibility for the final content.

#### Author's Address

Qiwen Tao  
Independent Researcher  
Email: [natureconservation@yeah.net](mailto:natureconservation@yeah.net)