

CARCS

COMPLIANCE ACCOUNTABILITY RECORD & CASE STUDY

A Structured Documentation Protocol for Human-Governed AI Work

Version 1.4 · May 2026

Basil C. Puglisi, MPA

A Human-AI Collaboration

basilpuglisi.com

Working paper released for feedback and collaboration · Published under Open Access

#AIassisted using the HAIA Ecosystem

The Documentation Gap in Governed AI Work

Organizations and practitioners are deploying AI at scale, and most governance discussion still centers on what AI should do, what it should not do, and which policy guardrails should contain it. What receives far less attention is the practical question that arrives after the work is finished: what record exists of how the work was done, who decided what, and which evidence shaped the decision?

This question is no longer theoretical. Auditors, regulators, researchers, and internal review teams are asking for record keeping, traceable human oversight, and documented methodology. The EU AI Act raises the pressure in regulated settings, while enterprise compliance programs and professional publishing raise it everywhere else. Any practitioner who has tried to reconstruct a complex multi-platform project two weeks later already knows the operational truth. Preserved conversation histories are not the same thing as a governed record.

The problem is not lack of trace, because AI work leaves plenty of trace. The problem is that those traces are scattered across platforms, trapped in provider-specific formats, and organized around conversation flow rather than around the questions an audit actually asks. An audit asks who decided, on what evidence, at which point, and why. Raw platform histories can answer those questions, but only after someone reads through everything and painstakingly reconstructs the decision logic from fragments spread across sessions and providers. That is not documentation. That is archaeology.

CARCS exists to close that gap.

What CARCS Is

CARCS stands for Compliance Accountability Record and Case Study. It is a structured documentation protocol that transforms the raw evidence of AI work sessions into a portable, organized record suitable for internal review, external audit, academic citation, or publication as a case study.

CARCS is human-triggered. A practitioner runs a prompt at the end of a session or project, the synthesis platform draws on available session evidence, and the result is a ten-section document that captures what happened, who decided, what the evidence was, and what comes next. The practitioner reviews and approves the document before it is finalized. Nothing in the record is finalized until a named human has signed off on it.

CARCS is not a framework and does not govern how AI work is conducted. It documents how AI work was conducted, and the distinction matters because CARCS can apply to any AI workflow, not only those operating under HAIA methodology. Any practitioner or organization doing structured AI work who needs a record of that work can use CARCS. The prompt in the appendix to this paper is deliberately designed without framework-specific terminology so it functions across any context.

Within the HAIA ecosystem, CARCS serves as the human-produced documentation layer for Agent Model 3 operations, where the practitioner orchestrates AI work manually across multiple platforms without automated logging infrastructure. Outside the HAIA ecosystem, CARCS applies wherever AI work needs to be documented, organized, and made legible to anyone who was not present for the work itself.

Why Structured Documentation Matters

Structured documentation matters first at the evidentiary level. When an organization is asked to show that a human reviewed and approved an AI output before action was taken, a conversation history is thin evidence. It shows what was said, but it does not reliably show what was decided. It shows what the model produced, but it does not always show whether the human applied judgment. A structured record that names the arbiter, classifies the decision, and records the rationale belongs to a different evidentiary category. It is not just activity preserved on a screen but governance made legible.

It matters just as much at the operational level. AI work often stretches across multiple sessions, several platforms, and more than one week. Each new session on a platform without durable memory forces the practitioner to reconstruct context, decisions, and current project state before useful work can resume. A CARCS record produced at the end of each session changes that pattern. The next session opens with a governed briefing rather than a memory exercise, which reduces restart time, lowers reconstruction error, and keeps the work compounding forward.

It also matters at the institutional level. Each documented decision, each preserved dissent, and each named act of human arbitration becomes part of the organization's empirical record of how it actually works with AI. Over time that record says far more about governance maturity than a policy statement ever can. Without that evidence, claims of oversight remain claims. With it, they begin to look like practice.

How CARCS Works

CARCS produces a ten-section draft record from a three-part prompt suite, subject to practitioner completion where evidence or provenance is incomplete. The sections are organized to move from context to evidence to implications to continuity, and the structure is designed so that each section answers a specific question an auditor or reviewer would ask.

Two sections carry particular weight for practitioner completion. Section 5 is only as strong as the explicit decision trace in the underlying record, and Section 10 depends on external provenance data the synthesis platform cannot generate independently. Both sections are flagged for practitioner review when evidence is absent or ambiguous.

The opening sections establish context: what project this is, what session type it represents, what the evidence status of the record is, and what working posture existed at session open. The evidence status declaration is particularly important. Not all synthesis

platforms have equal access to session history, and CARCS requires the record to declare upfront whether it is drawing on complete project memory, a single session window, partial cross-session memory, or manually assembled evidence.

The middle sections capture the substance of the work: what prompts were dispatched and to which platforms, what the platforms produced including convergence and dissent, and what the human arbiter decided at each point. The dissent requirement deserves emphasis. AI platforms operating in parallel frequently disagree, and those disagreements are often the most valuable signal in the entire session. A synthesis that smooths divergent outputs into a unified summary eliminates precisely the information that governance exists to preserve. CARCS requires dissent to be isolated and documented, not averaged away.

The human arbiter record classifies every decision using a four-type taxonomy drawn from operational practice: corrective override, where the platform made an error the human caught; creative supersession, where the human produced something better than the platform's best output independently; checkpoint confirmation, where the human actively reviewed and approved rather than passively accepted; and deferred decision, where the human identified an item but chose not to resolve it in this session. This taxonomy matters because it distinguishes between different kinds of human engagement. Governance that cannot tell the difference between a human who caught an error and a human who approved without reading is not governance.

The quality of Section 5 depends directly on decision tracing during the session itself. The recommended practice is to tag governance decisions inline using [T0-OVERRIDE] at the moment they occur. This tag is a searchable anchor that the synthesis platform extracts when generating Section 5, producing a far more accurate record than reconstruction from memory after the fact. A practitioner who tags decisions during a session produces a classified Section 5 without additional effort at record generation time. A practitioner who tags nothing produces a Section 5 flagged for manual completion. Tagging is not mandatory, but its absence is the single most common cause of incomplete Section 5 records.

The governance observations section carries the human arbiter's direct commentary on what the session showed about the methodology in practice, written in third person with the arbiter named explicitly in the prose. Named attribution is mandatory, because accountability without a name attached to it is not accountability.

The closing sections handle continuity and implications: what the session reveals about the methodologies being used, what the next practitioner needs to know before the next session begins, what items remain open, and what the underlying evidence base for the record is. The raw evidence index connects the CARCS narrative back to the actual platform session records, applying SHA-256 integrity standards where tooling permits and Practitioner Disclosure where it does not.

Practitioners who want to move from Attestation Grade to Hash Verified before GOPEL infrastructure is in production can do so today with a basic Python script. Exporting the session as a text or JSON file and running a SHA-256 hash against that file using Python's built-in hashlib library takes under a minute and produces a verifiable fingerprint that can

be recorded in Section 10. Any subsequent modification to the exported file produces a hash mismatch, which is the chain-of-custody mechanism the Raw Evidence Index is designed to establish. No specialized tooling is required. A locally computed hash stored in operator-controlled storage is self-attested integrity, not externally verified chain of custody. External verification requires anchoring the hash to a third-party timestamp service or equivalent infrastructure, which GOPEL is designed to provide.

Legal Position and Current Exposure

Any document positioning itself as a compliance artifact must address the question of legal standing, and the honest answer in April 2026 is that the legal terrain is shifting and no guarantee is possible.

In *United States v. Heppner*, No. 25-cr-00503-JSR (S.D.N.Y. Feb. 17, 2026), Judge Jed Rakoff declined to extend attorney-client privilege or work product protection to documents Bradley Heppner generated using Claude, Anthropic's consumer AI platform, placing significant weight on Anthropic's privacy policy terms the user had accepted and on the absence of counsel direction. The ruling is the first of its kind and most legal commentary, including Akin Gump and the Harvard Law Review, names Claude specifically because the consumer-versus-enterprise distinction is central to how future cases may receive different treatment. Commentary from Akin Gump ("SDNY Rules Communications With a Public Generative AI Platform Are Not Protected by Attorney-Client Privilege or Work Product Doctrine," February 2026) and the Harvard Law Review ("United States v. Heppner," March 2026) has characterized the ruling as fact-specific rather than categorical, suggesting that different configurations, particularly closed enterprise systems or counsel-directed work, may receive different treatment.

The author holds that the court decided *Heppner* without augmented intelligence theory before it, without applying the expectation of privacy doctrine from *Katz v. United States* (1967) to the gap between what platform terms technically permitted and what a reasonable practitioner expected, and without examining whether terms of adhesion on platforms that have become functionally essential cognitive infrastructure can constitute knowing waiver. These are arguments that were not made, not arguments the court considered and rejected. The author views these as valid grounds for future challenge as case law develops.

The author further holds that infrastructure changes could possibly shift this risk environment materially. A closed, cryptographically attested, operator-controlled channel for AI inference, of the kind proposed in the Governance Orchestrator Policy Enforcement Layer (GOPEL) specification and the Verified AI Inference Standards Act (VAISA) submitted to the 119th Congress, may address the primary confidentiality exposure that public commercial platform terms currently create. This is the author's position on a developing situation, offered with conviction and without guarantee.

Practitioners in regulated or legally sensitive contexts should consult legal counsel before treating any CARCS record as a privileged compliance artifact under current infrastructure conditions.

There is also a less-discussed risk that practitioners should weigh before choosing a compliance-grade label. A CARCS record that names the arbiter, classifies decisions, preserves dissent, and documents deferred items is also a roadmap for any challenger seeking to attack those decisions. A Checkpoint Confirmation with thin rationale is worse than no record because it documents shallow review under a label claiming depth. Preserved dissent is discoverable and can be used against the practitioner as evidence that the other platforms warned of a risk and the arbiter approved anyway. This is not a reason to avoid producing records. It is a reason to choose the Evidentiary Use label carefully and to ensure that every Checkpoint Confirmation in Section 5 carries rationale that reflects the actual depth of review.

Versioning and File Management

Every CARCS document receives a unique versioned filename following the convention HAIA_CARCS_[ProjectIdentifier]v[Major][Minor].md. Version increments are mandatory on every file change without exception. The scale is: a .1 increment for governance or structural changes, meaning any change to content, argument, schema, or sections; and a .01 increment for detail changes, meaning prose, grammar, spelling, voice edits, and formatting corrections. Prior versions are never overwritten. Each revision produces a new file so the complete change history is visible in the filename sequence.

When an error is discovered after approval, the correction procedure is to produce a new versioned file labeled ERRATA in the Document Control block, describe what was wrong and what the correction is, and mark the prior version as superseded with a pointer to the corrected file. A compliance record with known errors and no errata trail is worse than no record.

EU AI Act Context

The EU AI Act's high-risk AI system provisions are currently scheduled to take full effect on August 2, 2026, though active European Commission proposals under the Digital Omnibus package may extend the high-risk timeline by up to 16 months. Practitioners in EU-regulated contexts should confirm the operative dates with legal counsel before treating August 2026 as a hard deadline. Where the provisions do apply, practitioners deploying AI governance should assess CARCS records against Articles 12 and 14, which address record-keeping and human oversight respectively. CARCS is designed to produce evidence directly relevant to both. Article 11, which addresses technical documentation, requires organizational documentation beyond what CARCS alone provides, and practitioners in EU-regulated contexts should supplement CARCS accordingly.

The HAIA Context

CARCS originated inside the HAIA ecosystem, a body of open-source human-AI governance work developed by Basil C. Puglisi, MPA and published at github.com/basilpuglisi/HAIA under Creative Commons. HAIA stands for Human Artificial Intelligence Assistant, and the ecosystem it names is a layered architecture of governance frameworks, measurement instruments, and operational protocols built around a single governing principle: human judgment must remain at the center of every consequential AI-assisted decision.

The HAIA ecosystem includes several interconnected frameworks. Factics, developed in 2012, is the foundational methodology pairing facts with tactics and measurable outcomes. HAIA-RECCLIN structures how AI platforms respond to human requests, assigning seven functional roles and requiring ten-field structured output. HAIA-CAIPR (Cross AI Platform Review) governs parallel multi-AI orchestration. GOPEL (Governance Orchestrator Policy Enforcement Layer) automates the mechanical operations of CAIPR through seven deterministic steps with zero cognitive work by design, producing hash-chained, append-only audit trails. CBG (Checkpoint-Based Governance) is the constitutional authority layer. HEQ and AIS (Human Enhancement Quotient and Augmented Intelligence Score) measure what human-AI collaboration produces that neither party could produce independently.

The three HAIA operating models define how much automation runs between the human and the AI output. Model 1 runs a full pipeline with a single human checkpoint at the end. Model 2 pauses for human review after each functional role. Model 3, which CARCS was built to serve, runs with no agent and no automated infrastructure. The practitioner orchestrates every step manually, and the gap is structure. CARCS is the structured synthesis layer that closes it.

This paper and the prompt suite in the appendix were designed deliberately for use beyond the HAIA ecosystem. The ten-section schema, the evidentiary grading system, the dissent mandate, and the accountability requirements all function independently of whether the practitioner knows what RECCLIN or CAIPR mean. Practitioners operating within the HAIA ecosystem will find CARCS significantly more powerful because the infrastructure already exists to populate it. They are not adopting a documentation add-on. They are activating the documentation layer that their existing governance practice was already generating the evidence to support.

CARCS and the Governance Infrastructure

A CARCS record documents what happened and who decided. The HAIA ecosystem components that surround CARCS determine what kind of evidence that record carries.

Each component in the ecosystem addresses a distinct governance function. When that component is active in the deployment being documented, the CARCS record inherits the evidentiary properties that component produces. When it is not active, the record carries the evidentiary properties of whatever is active, down to the practitioner's own attestation

at minimum. The record is honest about which infrastructure produced the evidence it contains.

CBG (Checkpoint-Based Governance) provides the constitutional authority the record traces. CBG declares who holds binding decision power, at which checkpoints, under what constraints, and through what accountability channels (moral, employment, civil, criminal). CBG also declares the trust assumptions the governance structure rests on: the Constitutional Wall (any reduction in substantive human engagement at a checkpoint converts AI Governance to Responsible AI regardless of physical presence), the Asimov harm boundary (no human governor may direct an AI-assisted outcome that injures a human being, allows harm through inaction, or harms humanity), and the passive acceptance detection thresholds (95% automated approval rate, 2% reversal floor). When CBG is active, the CARCS record carries constitutional accountability. When it is not, the record documents Responsible AI mode and declares the absence of named human authority.

GOPEL (Governance Orchestrator Policy Enforcement Layer) provides the lifecycle evidence the record indexes. GOPEL's six audit record types map to the stages of governed AI work: Request Record (task initiated), Dispatch Record (prompt sent to platforms), Response Record (platform output received), Navigation Record (synthesis performed), Arbitration Record (human decided), and Decision Record (authorized output linked to every upstream record). SHA-256 hash chains bind every record to every other in a tamper-evident sequence. When GOPEL is active, Section 10 of the CARCS record carries Hash Verified integrity. When it is not, Section 10 carries Attestation Grade integrity, and the practitioner may optionally compute a local SHA-256 hash as described in this paper's Hash Verification section.

GOPEL CPE (Confidential Processing Extension) provides data boundary classification. CPE's four-profile system (Profile 0 through Profile 3) classifies how data exposure is governed during AI platform inference. When CPE is active, the CARCS record documents which data protection profile was applied and what evidence exists that the profile held during execution.

GOPEL CICE (Checkpoint Information Chain Extension) provides delegation chain and identity binding. CICE closes the semantic governance gap at the human checkpoint by chaining context across Navigator interactions. The Delegation Record requires a mandatory throughput-versus-risk rationale, specifies which checkpoint categories receive full review versus delegated review, defines triage criteria, sets review cadence and sampling rates, names escalation triggers, and carries an expiry date. When CICE is active, the CARCS record documents the complete chain from platform response to human decision, including who was authorized to decide at each point and whether delegation occurred explicitly or by default.

GOPEL Post-Quantum Cryptographic Agility Amendment provides the cryptographic trust assumption declaration. The amendment extends GOPEL's hash chain specifications to accommodate evolving cryptographic threats while preserving the seven deterministic operations and the non-cognitive constraint. When the Post-Quantum Amendment is

active, the CARCS record documents which cryptographic primitives are in use and what mathematical assumptions they rest on.

HAIA-Overwatch provides continuous monitoring from outside the trust boundary. Overwatch observes the GOPEL pipeline from a separate trust boundary, verifies infrastructure integrity and exchange alignment on every transaction, and escalates to human governance authority only when discrepancies are detected. It touches nothing, blocks nothing, modifies nothing, and watches everything. When Overwatch is active, the CARCS record documents that a monitoring layer existed outside the system being governed and whether it detected discrepancies during the period covered by the record.

Operating Model Selection determines checkpoint density and evidence generation scope. Model 1 (Agent Responsible AI) runs a full pipeline with a single endpoint checkpoint. Model 2 (Agent AI Governance) pauses after each functional role for human review. Model 3 (Manual Human AI Governance) runs with no automated agent. The model selection is itself a CBG decision, documented in the audit file with risk classification rationale. The CARCS record documents which model was active because the model determines how much of the governance infrastructure produced evidence and how much relied on the practitioner's direct attestation.

The configuration declaration in the CARCS prompt suite (see Appendix A, Prompt One) requires the practitioner to state which components are active before record generation begins. This declaration is not decorative, because it sets the evidentiary ceiling of the record. A CARCS record produced under Model 3 with no GOPEL infrastructure carries Attestation Grade integrity. A CARCS record produced under Model 2 with GOPEL, CICE, CPE, and Overwatch active carries Hash Verified integrity with data boundary classification, delegation chain evidence, and continuous monitoring confirmation. Both are valid records, but they are not equivalent evidence.

APPENDIX A

The CARCS Prompt Suite

The following prompt suite can be run at the end of any AI session or project. It is designed to work on any AI platform without requiring prior knowledge of the HAIA ecosystem. Run it in the platform that has the most complete view of the work. All three prompts are required for a complete record.

Session Tips Private Evidence Pulls: Upload session exports to the project context before running prompts. Records built from private file stores are more defensible than those relying on public platform memory alone. **Platform Memory Note:** Identify your platform's memory type before running: Full Evidence, Session Limited, Memory Partial, or Practitioner Reconstructed. **Live Session Tip:** Tag governance decisions inline with [T0-OVERRIDE] during the session. This creates a searchable anchor for Section 5. **Voice Gate Tip:** Write Section 6 observations in third person during the session rather than relying on AI voice conversion afterward.

Pre-Flight Checklist

Complete before running any prompt. If your answers to questions 1 and 2 are No, the resulting record will be Attestation Grade with significant gaps. Do not classify it as Compliance-Oriented.

- Have you tagged governance decisions using [T0-OVERRIDE] or explicit language such as “approved,” “rejected,” or “going with X”?
- Have you exported or can you access session evidence from all platforms used?
- Human arbiter name for Section 6: _____ (Write your full name before running Prompt One. If left blank, Section 6 will be flagged [PRACTITIONER NAME REQUIRED] throughout.)
- What is the intended use of this record? (Compliance / Research / Case Study / Continuity)
- If the synthesis platform declares Full Evidence, verify by asking it to name at least three prior sessions by date or topic. If it cannot, downgrade to Memory Partial.

Prompt One: Session Identity, Working Context, and Methodology

CARCS PROMPT ONE

GLOBAL RULE: This record is written in third person throughout.****

No first-person singular anywhere in the document.

Your task is to draft the first three sections of a CARCS record from the evidence actually available in your active context, accessible memory, or practitioner-supplied materials. Do not claim access to evidence you cannot see. If evidence is missing, partial, or ambiguous, state that explicitly and flag the field for practitioner completion.

Before generating, state your memory type as the first line of your output:

- Full Evidence: Complete access to all project sessions
- Session Limited: Access only to this chat window
- Memory Partial: Cross-session memory without full conversation histories
- Practitioner Reconstructed: Evidence was manually pasted into this session

This declaration sets the evidence status class for the record.

GOVERNANCE INFRASTRUCTURE DECLARATION

Before generating any section, declare the active governance configuration for this session. This declaration sets the evidentiary ceiling of the record.

For each component, state ACTIVE (with version) or NOT ACTIVE:

- CBG (Checkpoint-Based Governance): [version or NOT ACTIVE]
- GOPEL (Governance Orchestrator Policy Enforcement Layer): [version or NOT ACTIVE]
- GOPEL CPE (Confidential Processing Extension): [version or NOT ACTIVE]
- GOPEL CICE (Checkpoint Information Chain Extension): [version or NOT ACTIVE]
- GOPEL Post-Quantum Amendment: [version or NOT ACTIVE]
- HAIA-Overwatch: [version or NOT ACTIVE]
- Operating Model: [Model 1 / Model 2 / Model 3]

If operating outside the HAIA ecosystem, state: NON-HAIA DEPLOYMENT.

The record proceeds with Attestation Grade integrity by default.

This declaration populates the record infrastructure baseline.

CARCS PROMPT ONE

Section 10 integrity grade must be consistent with this declaration:
Hash Verified requires GOPEL active. Attestation Grade is the ceiling without GOPEL. Continuous monitoring evidence requires Overwatch active. Delegation chain evidence requires CICE active. Data boundary classification requires CPE active. Constitutional accountability requires CBG active.

EVIDENCE CEILING RULE: Summarize only what is visible in the evidence

available to you. Do not infer unseen sessions, missing decisions, hidden files, or unstated governance actions. Flag anything unconfirmed as [PRACTITIONER REVIEW REQUIRED] rather than presenting it as confirmed fact.

SECTION 1: SESSION IDENTITY

- Project or case study identifier
- CARCS version number (v1.0 for a new record)
- Date of this session
- Session type: New project, continuation, or project close
- Evidence status class
- If continuation: which prior CARCS was loaded at session open

SECTION 2: WORKING CONTEXT AT SESSION OPEN

- What instructions, working agreements, or operating rules were active
 - Which platforms had those instructions loaded and whether confirmed
 - If continuation: what was confirmed as carried over
 - If no context confirmation was performed, state that explicitly
- Absence of confirmation is an observable condition, not a failure.

SECTION 3: METHODOLOGY

- Prompts dispatched (full text or accurate summary)
- Platforms used and how work was organized across them
- What function each platform served for each task
- Prompt type: structured (source-level diversity) or open (perspective diversity)
- Whether any ambiguity was resolved before work began

CARCS PROMPT ONE

Stop after Section 3. Wait for review before proceeding to Prompt Two.

— *End of prompt. Review output before proceeding.* —

Prompt Two: Platform Outputs, Human Arbiter Record, and Governance Observations

CARCS PROMPT TWO

GLOBAL RULE: This record is written in** third person throughout.**

No first-person singular anywhere in the document.

CONTEXT GUARD: Confirm you have access to Sections 1-3 from Prompt One.

If not visible, state [CONTEXT LOST: RUN PROMPT ONE FIRST] and halt.

SECTION 4: PLATFORM OUTPUTS

- Role self-assignments where applicable
- Convergence signals: where platforms agreed
- Divergence and dissent: where platforms disagreed (with attribution)

DISSENT MANDATE: Do not smooth divergent outputs into a unified summary if organic consensus did not occur. Document all outputs that contradict the majority view. If no clear majority, document all verbatim and flag for human review. State platform count for convergence assessments. For Session Limited records: state dissent analysis is structurally unavailable and recommend multi-platform follow-up for compliance use.

SECTION 5: HUMAN ARBITER RECORD

Document only explicit decisions visible in available evidence.

Search for [T0-OVERRIDE] tags as strong decision anchors.

Fallback: scan for phrases like 'approved,' 'rejected,' 'going with X.'

Do not infer decisions from conversation flow or output shape.

Flag unconfirmed decisions as [PRACTITIONER REVIEW REQUIRED].

Classify each decision as one of:

- Corrective Override: Platform error caught and corrected by human
- Creative Supersession: Human produced better output independently
- Checkpoint Confirmation: Human explicitly reviewed and approved (Silence, continuation, or lack of objection does NOT count)

CARCS PROMPT TWO

- Deferred Decision: Item identified but deferred to Section 9

For each decision record: classification, what was decided, rationale, alternatives considered, and [EXPLICIT] or [INFERRED] anchor tag.

If no decisions found, add: ZERO DECISION SIGNALS DETECTED.

This flag indicates a detection limit, not absence of human governance.

SECTION 6: GOVERNANCE OBSERVATIONS

Write in third person. Name the human arbiter explicitly in every observation involving a human decision. Named attribution is mandatory.

Correct: 'Basil C. Puglisi directed the synthesis platform to...'

Incorrect: 'I directed the platform...'

Incorrect: 'The human arbiter directed...' (name required)

If name unavailable, flag as [PRACTITIONER NAME REQUIRED].

Preserve substance exactly; rewrite only for third-person voice shift.

Stop after Section 6. Wait for review before proceeding to Prompt Three.

— *End of prompt. Review output before proceeding.* —

Prompt Three: Methodology Implications, Continuity, Open Items, and Raw Evidence Index

CARCS PROMPT THREE

GLOBAL RULE: This record is written in third person throughout.****

No first-person singular anywhere in the document.

CONTEXT GUARD: Confirm you have access to Sections 1-6 from Prompts One and Two. If not visible, state [CONTEXT LOST: RUN PROMPTS ONE AND TWO FIRST] and halt.

SECTION 7: METHODOLOGY IMPLICATIONS

- Which methodologies are evidenced by this session
- Which need updating based on this session
- Any finding that changes how a protocol should be documented
- Any new concept or principle not yet formally named

CARCS PROMPT THREE

SECTION 8: CONTINUITY RECORD

Write for the next session, not as a backward summary:

- Which instructions should be confirmed active before the first task
- Which platforms have governance memory loaded at session close
- Which decisions are open and awaiting disposition
- Project state: complete / in progress / not started

SECTION 9: OPEN ITEMS

List every unresolved question, deferred decision, and flagged item.

For each: description, priority (Critical/High/Standard/Low),

and suggested resolution path. Nothing omitted because it seems minor.

SECTION 10: RAW EVIDENCE INDEX

PROVENANCE RULE: Do not invent filenames, URLs, hashes, or access paths. Mark any unavailable field [MISSING: PRACTITIONER COMPLETION REQUIRED]. A missing field limits evidentiary strength.

For each platform session: platform name, session date, export filename or URL (where available), access type, SHA-256 hash (where tooling permits).

Hash Verified and Attestation Grade are NOT equivalent evidentiary states. Declare which applies. For Practitioner Reconstructed records, include a Practitioner Disclosure listing what was included, omitted, and why.

CARCS COMPLETION STATUS

- All ten sections generated across three prompts
- Memory type and evidence status class declared
- Section 4: dissent documented, or single-platform caveat stated
- Section 5: classified decision present or flagged for completion
- Section 6: third person, arbiter named explicitly
- Section 8: written for next session, not backward summary
- Section 9: populated or states 'No open items identified'
- Section 10: integrity grade declared

CARCS PROMPT THREE

PRACTITIONER INSTRUCTION: Review complete document before finalizing.

Select the narrowest defensible Evidentiary Use label (see table in paper).

Add a Limitations paragraph enumerating specific evidentiary gaps.

Apply provenance labels [AI SYNTHESIZED] / [PRACTITIONER CONFIRMED] where the distinction matters.

Save as: HAIA_CARCS_[ProjectIdentifier]_v1_0.md

Never overwrite a prior version. Each revision produces a new file.

LEGAL NOTICE: This record may be discoverable in legal or regulatory proceedings. Consult legal counsel before treating as privileged.

VERSIONING: .1 increment for governance/structural changes.

.01 increment for prose, grammar, or formatting corrections.

Errors found after approval: produce new versioned file labeled ERRATA.

This document is #Alassisted.

— *End of prompt. Review output before proceeding.* —

Evidentiary Use Decision Table

After generating the complete record, select the narrowest defensible label and add it to the Document Control block. Add a Limitations paragraph enumerating specific evidentiary gaps.

Evidence Status Class	Sections 5 + 10 Complete	Sections 5 + 10 Partial or Flagged
Full Evidence	Compliance-Oriented Record with Declared Limitations	Compliance-Oriented Draft Pending Practitioner Completion
Session Limited	Research / Case Study Draft	Continuity and Internal Review Record
Memory Partial	Research / Case Study Draft	Continuity and Internal Review Record
Practitioner Reconstructed	Research / Case Study Draft	Session Save Point Only

The model produces the inputs. The practitioner makes the classification. When in doubt, select the narrower label.

APPENDIX B

Sources

The following sources are cited in the body of this paper. All legal commentary sources are listed in the order most directly relevant to the Heppner analysis. Author works are listed chronologically.

Primary Legal Sources

Katz v. United States, 389 U.S. 347 (1967).

United States v. Heppner, No. 25-cr-00503-JSR (S.D.N.Y. Feb. 17, 2026) (written memorandum opinion).

Legal Commentary

Akin Gump Strauss Hauer & Feld LLP. (2026, February). SDNY rules communications with a public generative AI platform are not protected by attorney-client privilege or work product doctrine. <https://www.akingump.com/en/insights/alerts/sdny-rules-communications-with-a-public-generative-ai-platform-are-not-protected-by-attorney-client-privilege-or-work-product-doctrine>

Debevoise & Plimpton LLP. (2026, February 11). SDNY rules AI-generated documents are not protected by privilege. <https://www.debevoisedatablog.com/2026/02/11/district-court-rules-ai-generated-documents-are-not-protected-by-privilege/>

Harvard Law Review. (2026, March). United States v. Heppner. Harvard Law Review Blog. <https://harvardlawreview.org/blog/2026/03/united-states-v-heppner/>

New York State Bar Association. (2026, March 10). Loose AI prompts sink ships: How Heppner shook the legal community. NYSBA News. <https://nysba.org/loose-ai-prompts-sink-ships-how-heppner-shook-the-legal-community/>

O'Melveny & Myers LLP. (2026, February 24). S.D.N.Y. first-of-its-kind ruling: AI-generated documents are not privileged. <https://www.omm.com/insights/alerts-publications/sdny-first-of-its-kind-ruling-ai-generated-documents-are-not-privileged/>

Paul, Weiss, Rifkind, Wharton & Garrison LLP. (2026, February 20). SDNY court considers whether AI-generated documents are subject to privilege protections. <https://www.paulweiss.com/insights/client-memos/sdny-court-considers-whether-ai-generated-documents-are-subject-to-privilege-protections>

Proskauer Rose LLP. (2026, February 18). SDNY addresses privilege and work product implications of using unsecured public AI tools. <https://www.proskauer.com/alert/sdny-addresses-privilege-and-work-product-implications-of-using-unsecured-public-ai-tools>

Regulatory Sources

European Parliament & Council of the European Union. (2024). Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32024R1689>

IAPP. (2026, February 3). European Commission misses deadline for AI Act guidance on high-risk systems. International Association of Privacy Professionals. <https://iapp.org/news/a/european-commission-misses-deadline-for-ai-act-guidance-on-high-risk-systems>

Author's Published Works

Puglisi, B. C. (2012). Digital Factics (1st ed.). Digital Ethos Press.

Puglisi, B. C. (2025). Digital Factics (2nd ed.). Digital Ethos Press.

Puglisi, B. C. (2025). Governing AI: When capability exceeds control (ISBN 9798349677687). Digital Ethos Press.

Puglisi, B. C. (2026a). Checkpoint-based governance: A constitution for human-AI collaboration (v5.0) [Governance specification]. basilpuglisi.com.

Puglisi, B. C. (2026b). GOPEL: Governance orchestrator policy enforcement layer (v1.5) [Technical specification]. basilpuglisi.com.

Puglisi, B. C. (2026c). HAIA: Human artificial intelligence assistant ecosystem [Open-source framework, Creative Commons]. GitHub. <https://github.com/basilpuglisi/HAIA>

Puglisi, B. C. (2026d). HAIA-RECCLIN: Human-AI governance methodology (3rd ed.) [Methodology specification]. basilpuglisi.com.

Puglisi, B. C. (2026e). The human enhancement quotient: Measuring cognitive amplification through AI collaboration. Social Science Research Network. <https://ssrn.com/abstract=6583419>

Puglisi, B. C. (2026f). AI provider plurality and the Verified AI Inference Standards Act (VAISA): Legislative framework for the 119th Congress. Social Science Research Network. <https://ssrn.com/abstract=6195238>

All URLs confirmed active as of April 2026. Readers seeking independent verification of United States v. Heppner may search the PACER federal court records system using docket No. 25-cr-00503-JSR.

CARCS v1.4 · Basil C. Puglisi, MPA · basilpuglisi.com · May 2026

#AIassisted using the HAIA Ecosystem

Open-source governance work: github.com/basilpuglisi/HAIA