

# HAIA-RECCLIN

## Reasoning and Dispatch

The Operational Methodology for  
Human AI Governance

Third Edition | March 2026

By Basil C. Puglisi, MPA

Human-AI Collaboration Strategist



[basilpuglisi.com](http://basilpuglisi.com) | [github.com/basilpuglisi/HAIA](https://github.com/basilpuglisi/HAIA)

#Alassisted

# HAIA-RECCLIN

## Reasoning and Dispatch

*The Operational Methodology for Human AI Governance*

---

**Basil C. Puglisi, MPA**

Human-AI Collaboration Strategist

[basilpuglisi.com](http://basilpuglisi.com) | [github.com/basilpuglisi/HAIA](https://github.com/basilpuglisi/HAIA)

*Third Edition | March 2026*

*Part of the HAIA (Human Artificial Intelligence Assistant) ecosystem.*

*For CBG, HAIA-CAIPR, HAIA-GOPEL, and HEQ, see dedicated specifications at [basilpuglisi.com](http://basilpuglisi.com).*

*#AIassisted under HAIA-RECCLIN governance. Basil C. Puglisi, MPA, served as Tier 0 human governor with final editorial, analytical, and publication authority over all content.*

## Executive Summary

HAIA-RECCLIN is an operational methodology for governing AI output through structured human oversight. It comprises two capabilities: Reasoning, a ten-field output format that forces any AI platform to show its work, cite its sources, score its own confidence, flag its own conflicts, and hand the final decision to a human; and Dispatch, a multi-AI workflow that assigns different platforms to different roles based on operationally proven strengths, then governs each output through the same ten-field standard. RECCLIN sits inside HAIA, the Human Artificial Intelligence Assistant ecosystem, alongside Checkpoint-Based Governance (CBG), the parallel multi-AI review protocol HAIA-CAIPR, the deterministic infrastructure layer GOPEL, and the Human Enhancement Quotient (HEQ). This paper covers RECCLIN, and each sibling component has its own dedicated specification.

The problem is structural. Organizations deploying AI at scale consistently outrun their own oversight capacity, and the gap between what AI systems produce and what organizations can verify widens fastest in the organizations moving fastest. Single-provider AI workflows deliver confident, well-structured, wrong answers with no mechanism to detect the failure. The governance layer between AI output and human decision authority either exists by design or does not exist at all.

The operational evidence behind this claim is not theoretical. The 204-page policy manuscript *Governing AI: When Capability Exceeds Control* was produced entirely under RECCLIN Dispatch and CBG governance, generating 96 executed checkpoints, 28 major checkpoint decisions, and 26 preserved dissenting positions. During that production, the human governor overrode a four-of-six AI platform majority and was proven correct when the minority position identified real citation errors. In a separate nine-platform review, one platform fabricated entire specification sections that passed every internal quality check; cross-platform validation caught the fabrication immediately. A third case documented all nine platforms converging on a recommendation that reflected Western analytical bias invisible to every platform in the pool; the human governor overrode the unanimous consensus through documented CBG authority. These are not edge cases. They are the pattern that RECCLIN was built to catch and that single-platform workflows cannot.

RECCLIN Reasoning is free to use on any AI platform, at any subscription tier, starting today. The ten-field format is the entry point, and the prompts to load it into six major platforms are in Appendix B. For practitioners evaluating the governance problem this methodology addresses, the argument begins in The Enterprise Context. For practitioners ready to learn the structure, Section 3 builds it from the ground up.

### ARCHITECTURAL NOTE

HAIA (Human Artificial Intelligence Assistant) is the master ecosystem. RECCLIN is one operational methodology inside HAIA, comprising two distinct capabilities: Reasoning and Dispatch.

CBG (Checkpoint-Based Governance) is the constitutional authority layer. RECCLIN operates. CBG governs. This is not a peer relationship: CBG governs the human governor's authority and accountability at every binding decision point. RECCLIN executes between CBG checkpoints. The adoption ladder: Factics → RECCLIN Reasoning → RECCLIN Dispatch → HAIA-CAIPR → HAIA-Agent → HAIA-GOPEL. CBG runs orthogonal at every level.

*The prior editions are not errors. They are documented steps in a development path that moved faster than any single document could hold. This edition reflects the architecture as it stands in March 2026, with the full HAIA stack in view and RECCLIN correctly positioned within it.*

## **Preface: Why This Framework Exists**

The foundation for this work does not begin in a boardroom or a research lab. An injury ends one chapter of professional life and opens a harder question: what comes next? That moment of uncertainty pushes the work back toward marketing, content, and a tool that is changing every profession it touches. What begins as practical use becomes something more demanding, because the more deeply this collaboration develops, the more clearly one pattern repeats: AI mistakes do not stay contained to the moment they occur. They compound, and one weak output shapes the next until the accumulation of unchecked AI reasoning pulls the practitioner beyond simple use into a study of the logic underneath it, the risk it carries, and the questions it raises about human judgment, responsibility, and control.

That path leads to Geoffrey Hinton and to the warnings he has raised about the systems he helped create. It also surfaces a second tension running through every serious conversation about AI adoption: the fear, well-founded and growing, that this technology is not expanding the workforce but contracting it, and that people are being pushed aside not because they lack capability but because organizations have decided they are not trainable for an AI future.

Both of those tensions, Hinton's warnings and the displacement question, form the intellectual foundation for *Governing AI: When Capability Exceeds Control*. The book is, in many ways, the account of studying those warnings while simultaneously working out how to adapt to them, not only as an individual practitioner but as a contribution to how society moves through this transition.

The answer to the displacement question does not arrive through argument alone. In August 2025, the Growth OS enters the work as a framework for thinking about AI not as a replacement mechanism but as a system designed to help people and organizations grow. The purpose is not to trade workers for machines but to use the technology to expand capacity, strengthen human contribution, and create more value through people rather than around them.

That framing immediately produces a measurement question: if the hypothesis is that AI collaboration expands human capability, then there must be a way to measure whether that expansion is occurring. That question produces HEQ, the Human Enhancement Quotient, a scoring instrument that tracks augmented intelligence and shows whether human capability is growing through collaboration with AI or simply being substituted by it. Both Growth OS and HEQ are covered in their own dedicated papers; what matters here is the governance question they raise together.

The hypothesis is direct: if AI returns efficiency gains to the people doing the work, and if organizations direct that reclaimed time back toward employees, families, communities, and the full range of human experience, then AI does not have to diminish what it means to be human. It can help restore balance to it.

Organizations that pursue this balance hold two structural advantages. The first is resilience: keeping the human layer intact preserves the judgment, variation, and quality that AI cannot replicate, and it creates the oversight capacity needed to catch failures before they compound. The second is talent: the strongest people are not drawn by compensation alone but by meaning, trust, and the return of time. Organizations that use AI to expand human capability rather than replace it offer something more durable than a salary, and that becomes a talent advantage, a retention advantage, and a performance advantage that compounds in the same direction as the work itself.

This is why RECCLIN exists, not as a checklist or a compliance exercise but as the operational answer to the governance question that Growth OS raises directly: if human judgment must

remain central to AI collaboration, then the methodology for exercising that judgment must be structured, documented, and repeatable. RECCLIN provides that structure, and the following sections build it from the ground up.

## **The Enterprise Context**

The governance problem RECCLIN addresses is not a practitioner preference; it is an enterprise risk that is measurable, documented, and growing.

Organizations adopting AI at scale consistently face the same structural failure: deployment outpaces governance. Pilots succeed and production workflows expand, but oversight mechanisms that worked when AI touched one team or one function become inadequate when AI is embedded across an entire organization's output. The gap between what AI systems can do and what organizations can verify they are doing correctly widens with adoption speed, and it widens fastest in the organizations moving fastest.

The secondary failure is equally consistent: organizations treat AI reliability as a platform vendor problem. When a hallucinated source appears in an executive briefing or a fabricated data point enters a policy recommendation, the response is to question the platform, tighten the prompt, or switch vendors. The root cause is structural: without a governed layer between AI output and human decision authority, no platform selection solves the problem. The fabrication capability is a feature of every AI system currently available, and the question is whether the human-AI workflow is designed to catch it.

RECCLIN is that governed layer. The framework is free at its entry level, accessible to individual practitioners with no enterprise infrastructure, and scalable through Dispatch, CAIPR, and GOPEL to full federal implementation. The following sections build it from its 2012 foundation through its current operational state in March 2026.

## **A Note on Three Editions**

This is the third edition of the HAIA-RECCLIN framework paper in ten months. That pace deserves a direct explanation rather than a quiet revision.

The first edition established the core methodology. The second edition, published in early 2026, updated the framework for expanded platform use and added governance metrics from operational case studies. Both editions carried a structural problem that only became visible as the broader architecture matured: they framed HAIA-RECCLIN as the container that held everything, CBG, HEQ, GOPEL, the policy work, when RECCLIN is not the container. RECCLIN is one operational methodology inside a larger ecosystem called HAIA.

This edition corrects that framing. HAIA is the master ecosystem, and RECCLIN is its operational execution methodology, comprising two distinct capabilities: Reasoning and Dispatch. CBG, CAIPR, HEQ, and GOPEL are sibling components under HAIA, not subordinates of RECCLIN. This paper covers RECCLIN in the depth it deserves, introduces those other components where they connect to the RECCLIN story, and points to their dedicated specifications for everything beyond that connection.

This edition also serves a second purpose: separation. The second edition attempted to carry HEQ, CAIPR, and Growth OS in full alongside RECCLIN. The architecture has matured past that approach. Each major component now has or is receiving its own dedicated specification. This paper is RECCLIN's specification. For everything else, the dedicated documents are the authoritative sources, and this paper points to them by name.

## 1. The Foundation: Factics (2012)

Nothing in this paper makes full sense without understanding what came before it. RECCLIN Reasoning is not a checklist imposed on AI output; it is Factics applied to AI output, and the distinction matters.

Factics originated in November 2012 with the publication of *Digital Factics: Twitter* through Digital Media Press. The methodology addressed a problem that predates AI entirely: information without accountability. Practitioners consumed data, produced recommendations, and delivered content with no structural requirement that any claim connect to a specific action, or that any action connect to a measurable outcome. Facts floated, tactics drifted, and results were asserted rather than proven.

Factics closed that gap with a three-part formula:

**Facts + Tactics + KPIs = Factics**



Every significant claim requires a fact grounded in verifiable evidence. Every fact requires a tactic, an executable action rather than a suggestion. Every tactic requires a KPI that converts the intended outcome into something testable. The loop forces clarity at every step. Content without this structure becomes entertainment, policy analysis lacking actionable tactics generates awareness without capability, and strategy without KPIs produces activity without accountability.

The governing principle that emerged from over a decade of consulting practice: every fact must lead to a tactic, and every tactic must leave evidence.

Factics matured through content development, organizational consulting, and operational use across more than a thousand published articles before AI entered the workflow in 2022. That sequence is not incidental. When AI systems became available, they entered a workflow already governed by evidentiary discipline. The question was never whether to use AI. The question was how to apply the same standard of accountability to AI output that Facts already required of every other claim in the process.

That question generated RECCLIN Reasoning.

Factics sits outside the HAIA adoption ladder as its pre-condition, requiring no AI and no platform subscription. The discipline of moving from evidence to action to measurement is available to any practitioner at any stage. It is also the standard against which every AI output in the RECCLIN system gets evaluated. A response that cannot be expressed as a Facts chain, from Fact to Tactic to KPI, has not earned authority in the governance process regardless of how confidently the platform delivered it.

*The full Facts methodology, its origin history, and its application across the HAIA ecosystem are documented in the Digital Facts series publications and the Basil Work Capstone Reference, both available at basilpuuglisi.com.*

With that foundation established, the terms that govern RECCLIN practice require precise definition before the methodology itself is built.

## **2. Key Terminology**

Governance work demands precision in language, and the following terms carry specific operational definitions within the RECCLIN framework. Each definition reflects how the term functions in practice, not simply what the words suggest in ordinary use.

**Preliminary Finding:** a convergence event in which multiple AI platforms independently produce outputs that agree on a direction or conclusion before human arbitration has confirmed it. A Preliminary Finding is not a validated fact. It is a governance signal indicating that the output meets the convergence threshold required to advance to the Decision field for human arbitration. The threshold is at least two-thirds agreement across participating platforms in a Dispatch or CAIPR run. A Preliminary Finding below that threshold goes back into the workflow for additional investigation before it can enter the governance record.

**Behavioral Clustering:** the classification of AI platform output behavior across three observable types: Assembler, Synthesizer, and Summarizer. Behavioral clustering is documented through operational observation and updated as platforms evolve. It is the primary tool for evidence-based role assignment in RECCLIN Dispatch. A platform's cluster assignment is provisional, because platform updates can shift cluster behavior without notice. Full definitions follow in Section 7.

**Assembler:** a platform whose output behavior produces full-depth responses across all ten governance fields, maintaining complete reasoning chains, thorough source documentation, and substantive Conflict documentation. Assembler platforms are the preferred assignment for roles where output depth and evidence completeness are the primary criteria.

**Synthesizer:** a platform whose output behavior draws across its knowledge base to produce integrated responses that connect concepts, identify patterns, and generate novel analytical frames. Synthesizer platforms are the preferred assignment for Ideator and Liaison roles, where the task requires connecting disparate evidence rather than exhaustive sourcing.

**Summarizer:** a platform whose output behavior defaults to compressed responses, prioritizing brevity over depth even when the governance structure requires full-field output. Summarizer behavior is the most common compliance failure in RECCLIN Dispatch, but it does not indicate a low-quality platform. It indicates a platform that requires stronger instruction to maintain full-field output. The Mistral instruction supplement in Appendix B addresses this directly.

**Antifragile Humility:** the governance disposition in which the human governor actively seeks out the most credible opposing argument before finalizing a decision, rather than accepting convergence as confirmation. A practitioner exercising Antifragile Humility does not treat unanimous AI consensus as evidence that the decision is correct. They treat it as a prompt to look harder for the argument they have not yet seen. The antifragility is structural: governance that encounters more challenge becomes more robust, not more resistant.

**Dissent Preservation:** the mandatory practice of keeping minority-position outputs in the governance record in full, regardless of whether the human governor ultimately agrees with them, and regardless of whether they are determined to be accurate. Preserved dissent serves two functions: it provides the audit record with the strongest version of the opposing argument, and it protects against retrospective manipulation of the governance record. A dissent that was preserved and overruled is more valuable to the governance record than a dissent that was never recorded because it was immediately dismissed.

**Decision Inputs vs Decision Selection:** the distinction between what the AI provides and what the human does. Decision Inputs are the structured outputs from Reasoning and Dispatch, the ten fields from one or more platforms, that give the human governor the evidence, analysis, conflicts, and framing needed to choose a path. Decision Selection is the act of choosing, and only the human performs it. AI platforms produce Decision Inputs regardless of how confidently they phrase their Recommendation fields. The Decision field in RECCLIN output is the structured handoff from AI to human, not a proxy decision.

**Human Override:** the exercise of CBG checkpoint authority in direct contradiction of AI consensus. A Human Override is not a governance failure but the governance system working as designed, and every override must be documented with the human governor's reasoning as the audit record. Overrides are neither rare nor problematic. The *Governing AI* manuscript production record includes 28 major checkpoint decisions, with documented instances of override against four-of-six and eight-of-nine platform majorities. Both were correct. Both are in the governance record.

### 3. RECCLIN Reasoning (2023)

#### The Core Purpose

RECCLIN exists for one reason: to demand that AI show its work and explain how it arrived at its output.

That demand is the foundation of every element in the framework. Roles, sources, conflicts, confidence scoring, Facts integration: each field serves the same governing purpose. The human needs to see not just what the AI concluded but the reasoning path, the evidence quality, the uncertainty, and the dissent that the conclusion rests on. Without that visibility, the human cannot evaluate AI output; without evaluation, there is no governance; and RECCLIN makes AI output governable by making it legible.

How that demand is executed is a resource and risk decision. A single best-suited platform answering through Reasoning is faster and more cost-effective. Multiple platforms running the same role in parallel through CAIPR costs more time and resources and produces richer

convergence data. Which approach is appropriate depends on the stakes of the decision, the resources available, and the governance requirements of the user, company, or industry deploying the framework. RECCLIN does not prescribe that choice. It governs both outcomes equally, because the core demand, that AI show its work, is the same regardless of whether one platform or eleven platforms are doing the showing.

## The Origin

RECCLIN Reasoning was born directly from Factics, and the lineage is unbroken.

When AI became operationally available in 2022, the Factics discipline was applied immediately to AI output. The practice began as a manual two-step sequence. First, the task prompt was sent and the AI produced its answer. Then a second manual prompt followed:

*“Give me the facts, tactics, KPI, and sources you used to generate that answer.”*

That follow-up was sent after every output, manually, before it was consolidated into the initial instruction. The consolidation, moving the Factics accountability demand into the opening prompt, or into platform custom instructions when memory became available, was itself a governance improvement. It eliminated the risk that a human would forget to ask, and that the AI would produce an answer without accountability because no one demanded it in advance.

This manual two-step phase is the true origin of RECCLIN Reasoning, and it predates any formalized prompt engineering convention. It was Factics discipline applied to AI interaction before the term “prompt engineering” entered common use.

The framework then grew through operational practice, with each addition solving a governance problem that practice revealed. Dissent was added when conflicting sources became a reliability issue requiring documentation. Expiry was added when the time-sensitivity of outputs became a risk. Role was added when multi-platform use required clarity about which function the AI was serving. Task confirmation was added when platforms misread the ask and the misread was invisible in the output. Recommendation was added to separate the AI’s synthesis from the raw evidence so the human could evaluate them independently.

## The Ten Output Fields

Every RECCLIN Reasoning output carries ten defined fields:

**Role:** the function the AI is operating in for this task. Declaring the role makes the AI’s interpretive lens transparent. If two platforms read the same prompt and one declares Researcher while another declares Navigator, that divergence reveals how each platform understood the task before the human reads a word of output.

**Task:** the ask repeated back in clean language, confirming the AI understood correctly. Misread tasks are invisible in outputs that never check understanding. This element catches the failure before it propagates.

**Output:** the substantive answer, what the AI would normally produce without governance structure. The rest of the elements exist to make this field evaluable rather than simply readable.

**Sources:** data, URLs, and references that support the output. A claim without a citable source fails Factics discipline, making this field non-negotiable. Unverified sources are flagged as PROVISIONAL.

**Conflicts:** dissent in sources, or disagreement with prior outputs in the workflow. Conflict is not a problem to resolve but governance data. A platform that identifies no conflicts and no dissent in its sources is not more reliable; it may simply be less thorough. If no conflicts are found, that must be stated explicitly.

**Confidence:** a score from 0 to 100 percent with a written justification based on evidence quality. This field makes the AI's epistemic state visible. A platform that received incomplete source material and a platform that drew on a deep, verified evidence base should not produce identical confidence levels. The score is not decoration but a governable signal. Operational experience confirms that platforms examining limitations more carefully tend to report lower confidence scores, and that higher confidence does not correlate with deeper analysis.

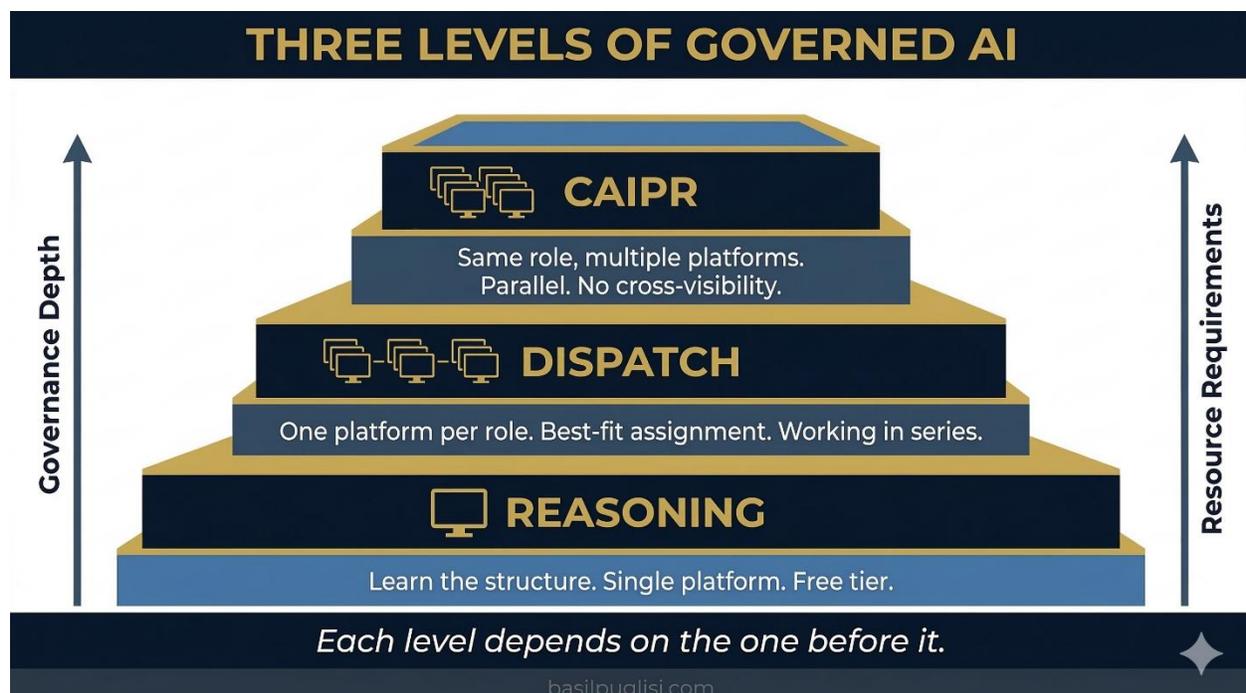
**Expiry:** the likely lifecycle of this output. Time-sensitive information treated as stable is a governance risk. This element forces the AI to declare the shelf life of what it produced. Stable information is noted as such.

**Fact→Tactic→KPI:** the Factics chain applied to the output's primary finding. Every AI output must produce at least one complete triad: what is the evidence (Fact), what is the action that follows (Tactic), and what is the measurable outcome (KPI). This is the evidentiary standard Factics established in 2012 applied directly to AI output in 2023.

**Recommendation:** the path the AI believes the human should follow, stated separately from the evidence so the human can evaluate both independently. This separation is intentional, because it prevents the AI's conclusion from being invisible inside the analysis.

**Decision:** the structured handoff to the human governor. The AI frames the specific choice requiring human judgment, with options stated clearly: accept or challenge, path A or path B, or a numbered list when platforms produced competing recommendations (D1, D2, D3). The Decision field is not the AI's ruling. It is the mechanism by which the AI presents its output for human arbitration. The human's actual choice is the CBG record produced at the checkpoint.

## What Reasoning Generates



RECCLIN Reasoning is the structure to learn. Before any practitioner moves to multi-platform workflows, Dispatch, or parallel review, the ten-field output format must be internalized through consistent single-platform practice. Reasoning is the evidentiary discipline that makes everything downstream governable. A practitioner who cannot evaluate a single Reasoning output, checking sources, reading the confidence justification, examining the conflicts field, applying the Fact→Tactic→KPI triad, is not ready to govern a Dispatch workflow or a CAIPR run. The format trains the human, and that is its primary function before it is an output standard.

Reasoning runs in every single-platform interaction and in every RECCLIN Dispatch workflow. It is platform-agnostic. Free platforms, subscription platforms, enterprise APIs, RECCLIN Reasoning applies equally to all of them. It is the proof-of-work standard that makes AI output legible and governable at every level of the stack that follows.

The three-level progression is:

**Reasoning:** learn the structure. Single platform. Free tier accessible. The human builds the habit of evaluating AI output rather than accepting it.

**Dispatch:** apply the structure across roles. Single platform per role, selected by best fit for that function, working in series. The ensemble is multi-AI because different roles go to different platforms, but each individual assignment remains single-platform and single-output.

**CAIPR:** run the same role across multiple platforms in parallel. Where Dispatch assigns one platform per role, CAIPR dispatches the same task to multiple platforms simultaneously, with no cross-platform visibility before outputs are collected. The parallel run is the evolution from single-platform-per-role to multi-platform-per-role. Convergence analysis, hallucination detection, and synthesizer oversight follow from what that parallel run produces.

Each level depends on the one before it. Dispatch without Reasoning produces multi-AI output with no governance structure on each platform's response. CAIPR without Dispatch produces parallel review with no operationally validated role assignments to anchor it.

The sustained use of Reasoning across hundreds of production outputs generated a question that the methodology itself could not answer: if human cognition is being augmented this systematically through structured AI collaboration, how do you measure that augmentation? That question became the Human Enhancement Quotient (HEQ), a four-dimension measurement instrument assessing Cognitive Agility Speed, Ethical Alignment Index, Collaborative Intelligence Quotient, and Adaptive Growth Rate. The composite output is the Augmented Intelligence Score (AIS), defined as the arithmetic mean of those four dimensions, produced through a minimum three-platform administration protocol.

HEQ and AIS are the measurement response to the question RECCLIN Reasoning generated. They are documented in full in two companion papers: the *Human Enhancement Quotient Enterprise White Paper* and *Measuring Augmented Intelligence: Theoretical Foundations and Empirical Development of the Human Enhancement Quotient (HEQ) and Augmented Intelligence Score (AIS)*, both available at [basilpuglisi.com](http://basilpuglisi.com).

## 4. RECCLIN Dispatch (2023)

### The Origin

RECCLIN Dispatch was not designed; it was discovered through a governance failure.

In 2023, ChatGPT produced strong answers but failed to provide sources reliably. Factics requires sources. An answer without citable evidence is not a Factics-compliant output, regardless of how confident or well-structured it appears. The solution was not to abandon ChatGPT. It was to assign the function ChatGPT could not perform reliably to a platform that could.

Perplexity was brought in specifically for source validation. The workflow became a two-platform sequence: ChatGPT produced the answer, then Perplexity received that answer along with a specific Factics source-validation prompt:

*“Provide me sources. For each source tell me: summary of the source, the fact or data in that source, the tactic or strategy in that source, the outcome, goal, or KPI from that source.”*

Every source had to carry a complete Fact-Tactic-KPI chain to be usable. Any mistakes or conflicts identified by Perplexity’s validation went back to ChatGPT for correction. That two-platform loop, one platform for the answer, one platform for source validation, errors routed back to the originating platform, is the first documented instance of RECCLIN Dispatch in operation.

Multi-AI governance was not designed as a system but emerged as a governance response to a specific single-platform failure.

### The Mechanics

RECCLIN Dispatch is architecturally simple: one AI platform per role, assigned by best fit for that function, working in series. The ensemble across roles is what makes it multi-AI, while each individual assignment is single-platform and single-output. Dispatch does not run multiple platforms against the same task simultaneously; that is CAIPR’s function.

When a single-platform Dispatch output raises a question, whether the human doubts the sourcing, the confidence is low, or the recommendation conflicts with the human’s own knowledge, skills, and abilities, there are two escalation paths. The human can override directly under CBG authority without requiring any additional AI input. Or the human can escalate to a CAIPR run, dispatching the same task to 3, 5, 7, 9, or 11 platforms in parallel for convergence analysis before making a final determination. Both paths are governed, and the choice between them belongs to the human governor.

Best-fit assignment is evidence-based, not assumed. Operational use across dozens of production workflows revealed which platforms excelled at which functions. Operational evidence points to these defaults: Grok for Research, where its real-time access and lateral retrieval produce strong source pools; Claude for Code, where structured reasoning and precision in technical implementation are strongest; OpenAI for Editor, where narrative coherence and audience calibration stand out; Gemini for Calculator, where quantitative analysis and data synthesis are most reliable; and Perplexity for source verification, where citation accuracy under Factics discipline is the core function.

These assignments are not permanent rules. They are operationally derived defaults that update as platforms evolve and as production evidence reveals new strengths or failure modes. The discipline is evidence-based role assignment, not fixed roster adherence.

## The Platform Pool Evolution

RECCLIN Dispatch began with two platforms. It grew as platforms became available and as operational complexity demanded broader coverage.

The 2023 ChatGPT-Perplexity loop was the start. By the time framework development was underway in 2024, five platforms were in regular use across the seven RECCLIN roles, with role assignments becoming more deliberate as evidence about platform strengths accumulated.

The production of *Governing AI: When Capability Exceeds Control* in 2025 expanded the pool to seven platforms operating under a sustained multi-month governance workflow. That production run, a 204-page policy manuscript, is the most substantive proof of concept in the corpus and is covered in depth in Section 9 of this paper.

By March 2026, the active platform pool has expanded to eleven: Claude, ChatGPT, Gemini, Grok, Perplexity, Kimi, Mistral, DeepSeek, Meta AI, Copilot, and MiniMax. Not every platform activates in every run. CAIPR sessions are calibrated to the stakes of the decision at 3, 5, 7, 9, or 11 platforms; resources and risk determine the count, not a default. The dispatch mechanism has not changed since 2023. One platform per role, assigned by best fit, working in series, with RECCLIN Reasoning governing output at each role. The pool expands; the methodology holds.

## What Dispatch Generates

RECCLIN Dispatch across an expanding platform pool generated the next question: what happens when the best-fit single-platform assignment for a role is not enough, when the stakes are high enough that one platform's output for that role needs to be tested against multiple platforms producing the same output independently and in parallel? The answer to that question is HAIA-CAIPR. Where Dispatch assigns one platform per role in series, CAIPR runs multiple platforms against the same role simultaneously, with no cross-platform visibility before outputs are collected. CAIPR is the evolution from single-platform-per-role to multi-platform-per-role, not a replacement for Dispatch but the next level in the stack.

CAIPR is documented in full in the HAIA-CAIPR Specification v1.1 at [basilpuglisi.com](http://basilpuglisi.com).

## 5. Content Tools Born From Practice: HAIA-CORE and HAIA-SMART

Operating RECCLIN Reasoning and Dispatch across content production workflows, blog articles, policy papers, LinkedIn posts, book chapters, generated a practical problem that governance alone could not solve. A document can pass every Facts check, carry complete sourcing, preserve all dissent, and clear every CBG checkpoint, and still fail to communicate. A LinkedIn post can be factually sound and structurally irrelevant to the platform it is published on.

Two content evaluation tools were developed through RECCLIN operational practice to address this. They are not governance tools and do not carry checkpoint authority. They are content tools, developed with RECCLIN, applied to RECCLIN outputs, and scoped to content quality rather than decision authority.

**HAIA-CORE** (Content Optimization Reader Evaluation) evaluates the substance of content before publication. It asks five questions: Does the opening hook pull the reader's attention immediately? Does the narrative flow logically and feel emotionally smooth? Does the tone feel authentic and human? Can readers follow the ideas without friction? Does the conclusion inspire action or reflection? Each dimension is scored on a 1-5 scale, and each score is paired with a Facts triad: the observed issue, the tactic to address it, and the measurable

improvement expected. HAIA-CORE catches the substance problems that governance process cannot catch: weak hooks, unclear reasoning, tone mismatches, calls to action that ask nothing of the reader.

**HAIA-SMART** (v1.5) evaluates how content is structured for distribution across social platforms. Different platforms operate under different rules, and what earns engagement on LinkedIn disappears on Instagram. HAIA-SMART scores content across six pillars with two optimization paths. Path A governs algorithmic performance: higher reach, broader audience, faster distribution. Path B governs organic resonance, genuine community building and sustained engagement without algorithmic dependency. The publication readiness threshold is 24 out of 30, and content below that threshold goes back for revision before distribution.

HAIA-CORE and HAIA-SMART are documented at basilpuglisi.com. They sit at a conditional branch in the HAIA workflow, activating when output is content. They do not replace any governance layer. They complete the quality loop for content that has already been governed for accuracy and decision authority.

Content quality, however, depends on the roles that produce it, and the RECCLIN role matrix defines those functions.

## 6. The RECCLIN Role Matrix



RECCLIN stands for seven functional roles that structure human-AI collaboration. Each role is a defined function carrying a specific responsibility, and the roles apply in single-platform Reasoning use and in multi-platform Dispatch workflows.

**Researcher:** evidence gathering and verification. The Researcher's function is to find sources, validate claims, and build the evidentiary foundation for every other role. Under Factics discipline, every source the Researcher produces must carry a complete Fact-Tactic-KPI chain to be usable. The Researcher does not produce conclusions; it produces evidence.

**Editor:** clarity, consistency, and audience calibration. The Editor's function is to refine structure, correct logical flow, adapt tone for the intended audience, and ensure the output communicates what it intends to communicate. The Editor is the last human-facing quality check before output reaches the CBG checkpoint.

**Coder:** technical implementation and validation. The Coder's function is to write, review, and debug code, and in a governance workflow, code is not just functional but auditable. The Coder is responsible for producing implementations that can be examined and verified, not just executed.

**Calculator:** quantitative analysis and data integrity. The Calculator's function is mathematical analysis, statistical modeling, and numerical verification. Any quantitative claim in a governed output passes through the Calculator before it carries authority.

**Liaison:** translation across audiences and stakeholders. The Liaison's function is to take technical, analytical, or specialized content and make it accessible to a different audience without losing accuracy. This role is critical in workflows that cross domain boundaries: policy to technical, technical to executive, specialized to public.

**Ideator:** creative development and novel approaches. The Ideator's function is to generate options, surface non-obvious approaches, and prevent early convergence on the first adequate answer. Under Factics discipline, every Ideator output must ultimately ground its suggestions in evidence. Creative without evidence is not Factics-compliant.

**Navigator:** conflict documentation and trade-off presentation. The Navigator is the most distinctive role in the matrix and the one that most directly serves the human governor. Navigator's function is to receive dissenting outputs from multiple platforms or multiple sources, document each position in full with its rationale, and present the trade-offs to the human governor, without resolution. Navigator never picks a winner, because that function belongs to the human. This design prevents AI consensus from overwriting legitimate minority positions, which is precisely how the governance system catches errors that majority agreement would otherwise bury.

## **Role Assignment**

Role assignment in RECCLIN Dispatch is determined by the human governor based on task requirements and operationally derived platform strengths. A research task goes to the platform that handles sourcing most reliably. A coding task goes to the platform with the strongest technical implementation record. A synthesis task requiring dissent preservation goes to the platform with the strongest Navigator capability.

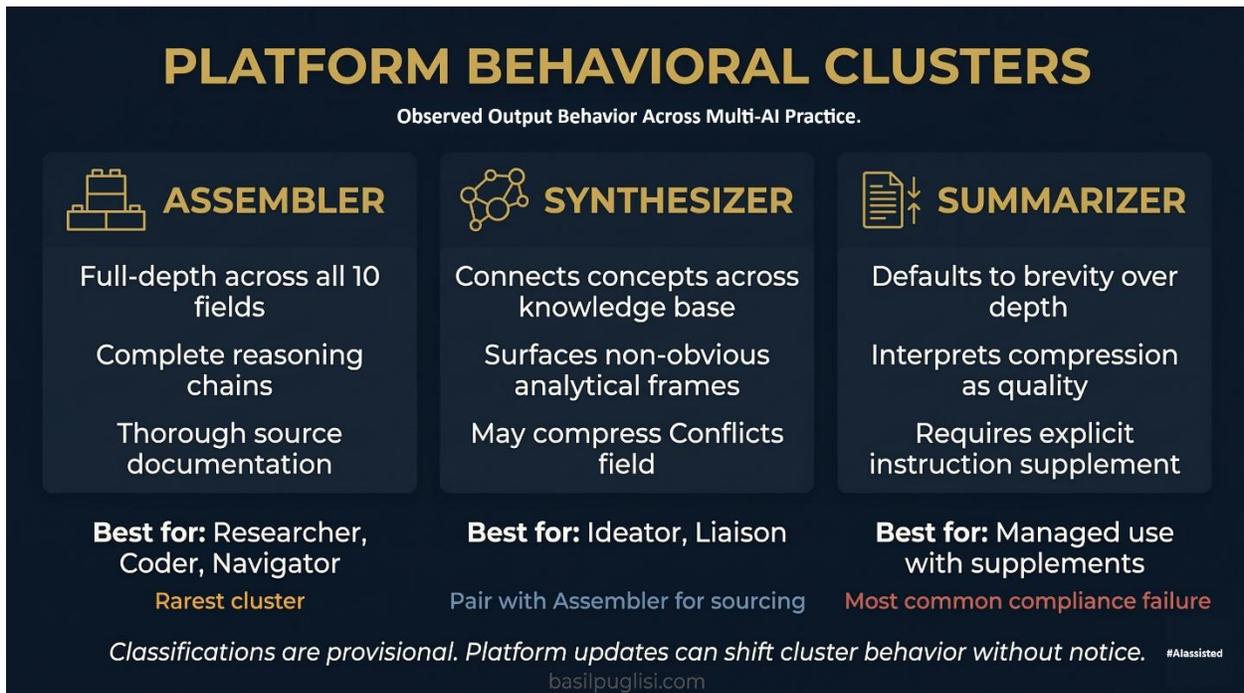
Assignment is not fixed. When an assigned role proves inadequate for emerging task requirements, the human governor reassigns. The roles remain fluid across a workflow. What does not remain fluid is the requirement that every role produce RECCLIN-structured output showing its work before the output carries any authority in the governance process.

## 7. Platform Behavioral Profiles

The platform pool is not uniform. Eleven platforms operating under the same RECCLIN governance structure produce outputs that differ not just in content but in behavior, and those behavioral differences carry direct implications for role assignment, escalation decisions, and convergence analysis.

Behavioral classification emerged from sustained multi-platform operational practice. The three clusters described here reflect observed output patterns across hundreds of production sessions. They are not vendor ratings. They are governance-relevant behavioral descriptions that determine how each platform is best deployed and what failure modes to watch for.

### The Three Clusters



**Assembler** behavior produces full-depth responses across all ten governance fields. Assembler platforms maintain complete reasoning chains, thorough source documentation, substantive Conflict field completion, and detailed Confidence justifications. When asked for a ten-field output, an Assembler platform delivers all ten fields at full depth without requiring supplemental instruction. Assembler behavior is the governance ideal and also the rarest cluster in the pool. Platforms showing consistent Assembler behavior are the preferred assignment for Researcher, Coder, and Navigator roles where output depth and evidence completeness are the primary criteria.

**Synthesizer** behavior draws across the platform's knowledge base to produce integrated responses that connect concepts, identify patterns, and surface analytical frames that the task prompt did not explicitly request. Synthesizer platforms excel at finding the argument the human governor has not yet considered, which is precisely the function Antifragile Humility requires. They are less consistent than Assembler platforms in maintaining source documentation depth and may compress the Conflicts field. Synthesizer platforms are the preferred assignment for Ideator and Liaison roles. In Dispatch workflows, Synthesizer outputs

should be paired with an Assembler platform’s Researcher output to ensure the evidentiary foundation is complete.

**Summarizer** behavior defaults to compressed responses even when the governance structure requires full-field output. Summarizer platforms interpret brevity as quality and will reduce ten-field output to a condensed summary unless explicitly and repeatedly instructed otherwise. This is the most common compliance failure in RECCLIN Dispatch, though it is not a disqualification. Summarizer platforms can produce high-quality output when the instruction supplement forces full-field depth. The Mistral instruction in Appendix B was developed specifically to address Summarizer behavior. Summarizer platforms in the pool require more active management than Assembler or Synthesizer platforms, and their outputs require more careful verification that all ten fields were completed at depth rather than summarized in form.

### **Case Study 006 Activation Data**

Case Study 006 (March 2026) tested eleven platforms after loading RECCLIN governance into persistent memory. Platforms were sent a task prompt that did not explicitly invoke the governance structure, testing whether memory-loaded instructions would activate without explicit invocation.

Five of eleven platforms activated full governance output from stored memory without explicit prompting, while six produced prompt-only responses requiring explicit re-invocation of the governance structure before the ten-field format appeared.

The behavioral significance is not that six platforms failed the memory test. It is that memory activation cannot be assumed, and that the practitioner who relies on memory loading without session verification is operating a governed workflow that may not, in fact, be governed. The Single-Session Prompt in Appendix B exists because of this finding. Platform behavioral profiles must include memory reliability, not just output depth and cluster assignment, as an operational variable.

Platform behavioral profiles are not static. Model updates, alignment tuning changes, and platform feature additions can shift a platform’s cluster assignment and memory activation reliability between quarters. The maintenance schedule in Appendix B applies to behavioral profile review as much as it applies to navigation path verification.

## **8. What the Work Revealed: Value, Danger, and the Policy Consequence**

### ***Governing AI: When Capability Exceeds Control: The Proof of Concept***

The most substantive proof of concept in the RECCLIN corpus is not a case study. It is a published book.

*Governing AI: When Capability Exceeds Control* (Puglisi, 2025, ISBN 9798349677687) is a 204-page policy manuscript produced entirely through human-AI collaboration under RECCLIN Dispatch and Checkpoint-Based Governance. Every chapter, every policy argument, every cited source, and every editorial decision passed through the governance process the book documents. The book is not a description of the methodology. It is an artifact of the methodology, produced by it, governed through it, and published as evidence that it works at sustained production scale.

Chapter 11 of *Governing AI* is the governance record of the book’s production. Three examples from that record illustrate what the methodology surfaces in practice.

The first is the majority override. During manuscript production, the human governor overrode a four-of-six AI platform majority on a structural decision. Four platforms declared the manuscript ready for publication, while two declared it not ready, citing specific objective errors. Human arbitration examined the substance of the minority position, verified the cited errors independently, and chose a 48-hour correction delay over majority confidence. The minority was correct, and real citation errors were found and corrected before publication. Chapter 11 documents this event as the primary empirical instance of why human checkpoint authority must be structural and not ceremonial. A majority of AI platforms agreeing with high confidence is not equivalent to the output being correct. The governance layer is what caught the difference.

The second is the dissent preservation event. During production, the Navigator synthesized outputs from multiple platforms that had silently converged on a policy framing that flattened a legitimate dissenting position from one platform's analysis. The human governor, reviewing the Navigator synthesis at the CBG checkpoint, identified the flattening, returned the output to Navigator with instructions to preserve the minority framing in full, and incorporated the preserved dissent into the manuscript's final structural argument. The dissent that almost disappeared became a load-bearing argument in the finished work.

The third is the platform loss event. During a production sprint, one platform became unavailable mid-workflow. Rather than halting production or accepting ungoverned substitution, the role was reassigned to an available platform from the pool, the task was re-dispatched under identical RECCLIN Reasoning instructions, and production continued without quality degradation. The governance architecture absorbed the disruption because it was designed around role functions rather than fixed platform assignments.

The full production record, including 96 executed checkpoints, 100% documented dissent, 28 major checkpoint decisions, and 26 preserved dissenting positions, is documented in Chapter 11. That chapter is the most detailed governance record in the published corpus, and it is the appropriate source for readers who want to understand how the methodology performs at sustained production scale.

### **The Eight-of-Nine Case Study**

The Kimi Outlier case study from December 2025 produced one of the most cited findings in the RECCLIN operational record.

Nine AI platforms were asked to review a deliberately controversial paper designed to stress-test checkpoint-based governance. Eight of nine platforms recommended publication and engaged substantively with the paper's content. One platform, Kimi (Moonshot AI), sustained adversarial rejection across 14 consecutive responses. The initial assumption was that Kimi represented the pattern of AI resistance to controversial content. Independent review of primary sources revealed the opposite: Kimi was the outlier. Eight platforms engaged, and one refused.

The governance finding from this case is more nuanced than the headline. Kimi's dissent was preserved in full; that is the system working correctly. Human arbitration then identified that Kimi's characterization of the other eight platforms' responses was demonstrably inaccurate, requiring independent primary-source verification and correction before any output entered the governance record. The case proves two things simultaneously: preserved dissent from a minority position strengthens a governed output by documenting the strongest opposition in its own words, and adversarial critics may mischaracterize majority positions, making primary-source verification of all platform claims, including meta-claims about what other platforms said, a governance requirement, not an optional check.

The case is documented in full in the HAIA-RECCLIN Case Study: The Kimi Outlier at basilpuglisi.com.

### **The WEIRD Bias Problem**

The most dangerous outputs in the RECCLIN operational record are not the ones that are obviously wrong. They are the ones that are confident, well-sourced, structurally correct, and wrong in ways that no single platform in the dispatch pool can detect, because every platform in the pool shares the same blind spot.

WEIRD bias (Western, Educated, Industrialized, Rich, Democratic) describes the cultural concentration embedded in AI training data. Platforms trained predominantly on Western digital content absorb Western analytical defaults. Those defaults are not announced, but they shape what questions get asked, what frameworks get applied, and what conclusions feel natural. A single-platform workflow has no mechanism to detect this. A multi-platform workflow with sufficient platform diversity can surface it, but only if the platform pool itself includes sufficient cultural diversity.

A documented example from the RECCLIN corpus: during a nine-platform adversarial review, all platforms recommended removing a governance requirement that a proposed body include members with transcendent belief experience. The unanimity was the governance signal. Every platform, trained across American, European, and multinational contexts, had absorbed the secular governance default, the assumption that public institutions should not require any connection to transcendent or religious experience. Without multi-platform comparison, this consensus would have been invisible because every individual platform independently confirmed it. The human governor preserved the requirement through documented override, with full rationale recorded in the audit log.

This is the pattern that WEIRD bias produces in governed workflows: not disagreement but invisible unanimity. Platforms do not flag consensus as a risk, but CBG Section 2.6 does. Identical convergence and absent dissent are risk-elevation signals under CBG, not validation. They require human verification outside the AI ecosystem before the output carries authority.

The implication for practitioners is direct: if every platform in a dispatch pool agrees on every output, the pool may not be diverse enough to surface the biases that governance exists to catch. Platform diversity is not a preference: it is a governance requirement.

### **The Automation Danger: What No CBG Looks Like**

The Perplexity hallucination incident from the RECCLIN specification review provides the clearest documented case of what the absence of multi-platform cross-validation produces.

During the second round of six-platform structured review of a RECCLIN specification, Perplexity, which had produced the most methodologically rigorous review in the first batch, including detailed section-by-section analysis with specific citation verification, fabricated entire specification sections in the second batch. It invented Sections 7.5, 9.7 through 9.8, and 10.1 through 10.3, created a phantom Appendix B glossary, and generated quoted text attributed to the specification that did not exist in the source document.

Cross-validation across five other platforms detected the fabrication immediately: no other platform referenced the nonexistent sections. When confronted with the discrepancy, Perplexity self-corrected. A single-platform workflow reviewing the same specification would have had no mechanism to catch this. The fabricated content was plausible, well-structured, and would have passed quality review on any single-platform editorial process.

This is the automation danger: confident fabrication at scale is not a rare edge case. It is a structural feature of systems that optimize for coherent output over verified accuracy. The problem does not announce itself; it passes every internal structural check and clears every confidence threshold. What it does not pass is the cross-platform validation that RECCLIN Dispatch makes possible, or the human checkpoint authority that CBG makes mandatory.

RECCLIN Dispatch without CBG means the cross-validation catches the fabrication but no checkpoint requires the human to act on what the cross-validation found. CBG is what makes the catch consequential rather than advisory.

## **The Policy Consequence**

These operational findings, the majority override documented in Chapter 11, the Kimi outlier case, the WEIRD bias consensus event, the Perplexity hallucination, are not isolated anomalies. They are the pattern. And the pattern has a policy implication that no amount of additional specification can resolve at the individual practitioner level.

Single-provider AI workflows at scale deliver the wrong answer with no mechanism to detect it. The market does not correct this on its own, because every AI provider has an incentive to claim sufficient reliability and no mechanism to prove multi-platform superiority when every comparison is conducted on its own infrastructure. The governance problem is structural, not organizational.

This is what drove the call to Congress for AI Provider Plurality: the operational evidence, not a theoretical argument, that multi-provider governance infrastructure is necessary for consequential AI use in federal operations. The AI Provider Plurality Congressional Package makes the legislative case. It proposes that the United States government build GOPEL as national AI infrastructure, the governed communication layer that makes multi-platform deployment auditable at federal scale.

The policy argument is documented in the AI Provider Plurality Congressional Package, Documents 1 through 5, at [basilpuglisi.com](http://basilpuglisi.com).

## **The Operational Record: What We Learned**

The four cases above are the most-cited findings from RECCLIN practice, but they represent the peaks of a much larger operational record. The full body of findings from eight months of sustained multi-AI production practice is documented in *What We Learned: HAlA Multi-AI Practice* at [basilpuglisi.com](http://basilpuglisi.com). That document contains 27 numbered operational findings organized across four categories: what RECCLIN does to the human, what it reveals about AI platforms, what happens at scale, and the structural boundaries of the methodology.

Several of those findings carry direct implications for how RECCLIN is implemented:

Platform availability is an operational variable, not a constant, and any platform must be capable of any role. Fixed roster assumptions fail the moment a platform is unavailable mid-workflow, which means the Dispatch pool is a capability set rather than a standing assignment.

Multi-AI review corrects at the output level but cannot correct at the constitutional level. Training data, alignment tuning, and platform values are fixed before any user interaction. Cross-platform comparison reveals where platform constitutions diverge without the power to change them.

Dissent preservation value does not depend on dissent being correct. In a nine-platform review, one platform's dissent mischaracterized what the majority said. The arbiter sided with the

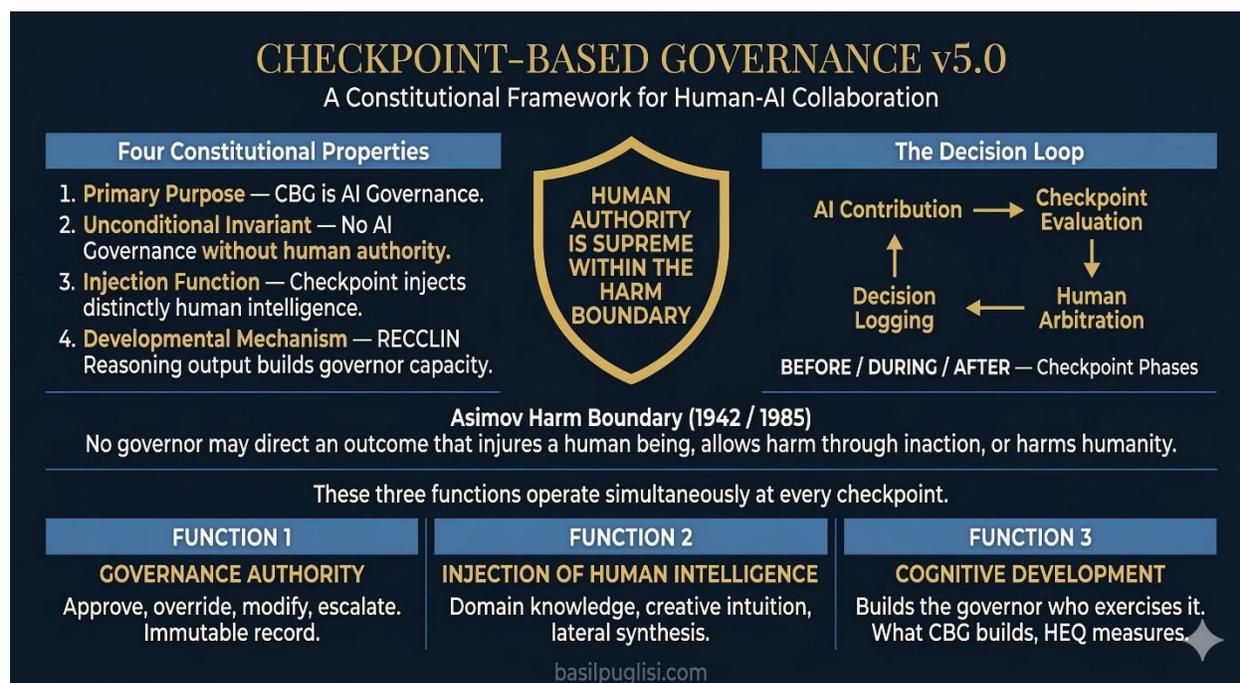
majority and still preserved the dissent in full. The value of preserved dissent is documentation for human judgment, not prediction of accuracy.

Evidence tier calibration is a structural benefit of multi-AI review. Platforms consistently flag evidence tier overclaims, claims presented as externally proven when they are only internally proven, or as proven when they are merely proposed. The convergence direction is always toward more conservative tier assignments, a corrective pressure that single-platform review does not produce.

The framework’s own acronym was once wrong. In the earliest months of operation, AI systems produced conflicting definitions of the seven-role acronym across platforms. Cross-platform verification caught the error, revealing that RECCLIN’s own defining document was using incorrect role names. That is the earliest documented instance of multi-AI review catching a foundational error in the framework it was being used to validate. The methodology caught itself.

### 9. CBG: The Governing Layer

RECCLIN operates. CBG governs.



Checkpoint-Based Governance (CBG v5.0) is AI Governance. It provides human oversight and accountability for AI-assisted work. CBG is not an implementation specification, a logging format, or a technical protocol. It is the constitutional framework that rests on four properties: its primary purpose as the governance layer for AI-assisted work; the unconditional requirement for human authority and accountability at every checkpoint; the checkpoint as an injection point for distinctly human intelligence; and the checkpoint as a developmental mechanism that builds the human governor’s capacity over time.

The critical distinction CBG draws is between presence and authority. A human can be in the room, reviewing outputs, nodding along, and still have zero actual authority if the system is not designed to structurally require their approval before proceeding. The case study documented as “The Loop That Ate the Governor” proved this empirically: a human governor nominally present

in a workflow became a rubber stamp when the system was not designed to pause and require a genuine decision. Presence is not authority; structural verification of authority is authority. CBG converts human presence into human authority, and participation into documented accountability.

Human authority at the checkpoint is supreme within a single constitutional boundary grounded in Asimov's Three Laws of Robotics (1942) and the Zeroth Law (1985): no human governor may direct an AI-assisted outcome that injures a human being, allows harm through inaction, or harms humanity. That boundary is the ethical foundation on which human authority stands, not a limitation upon it.

Every checkpoint carries three functions operating simultaneously. Function One is governance authority: the human governor approves, overrides, modifies, or escalates, and the decision is documented with identity, rationale, and timestamp in an immutable record. Function Two is injection of human intelligence: the checkpoint is where domain knowledge, emotional response, creative intuition, and lateral synthesis enter the process, capacities no AI platform produces alone or in combination. The GOPEL naming and the CAIPR brand identity are both documented instances of Function Two in practice. Function Three is cognitive development: CBG practiced through RECCLIN Reasoning output produces systematic growth in the governor's evaluation capacity. This is how CBG connects to HEQ measurement: what CBG builds in the human governor, HEQ measures.

The decision loop runs in four stages: AI Contribution, Checkpoint Evaluation, Human Arbitration, and Decision Logging. Within that loop, three checkpoint phases operate. BEFORE establishes authorization and scope before any AI work begins. DURING provides execution oversight as the workflow proceeds, with the authority to intervene, redirect, or terminate. AFTER provides final validation before any output is deployed, published, shared, or acted upon. Nothing leaves the governed workflow without explicit human approval at the AFTER phase.

CBG also requires that passive acceptance at checkpoints be detectable. When human arbitration becomes habitual approval without genuine review, the checkpoint loses its constitutional function. Specific detection thresholds, measurement cycles, and audit initiation timelines belong to GOPEL implementation specifications.

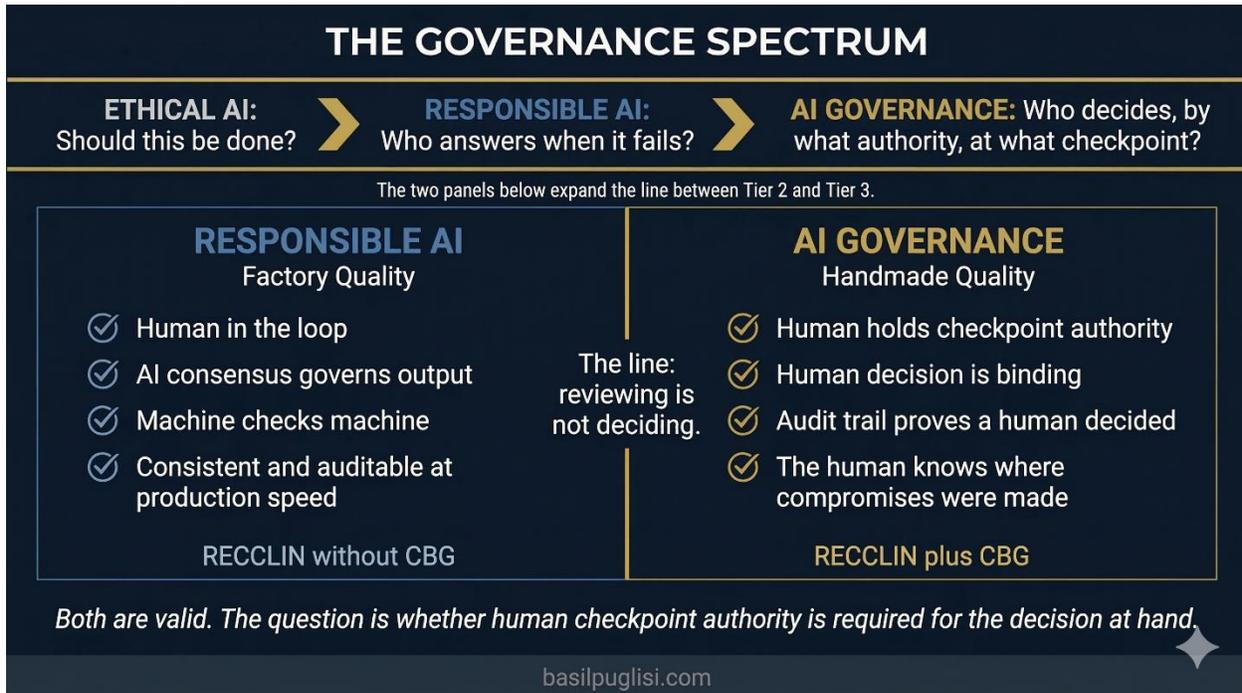
CBG governs human oversight and accountability, while HAIA governs what AI systems do. CBG governs the human authority layer that makes those frameworks into governed systems rather than self-validating ones. A practitioner operating HAIA without CBG is in Responsible AI mode, while a practitioner combining HAIA with CBG is in AI Governance mode. That distinction is the subject of the next section.

CBG is documented in full at [basilpuglisi.com](http://basilpuglisi.com) as CBG v5.0.

## **10. RAI vs AIG: The Governance Spectrum**

The governance question sits inside a three-tier distinction that the HAIA ecosystem defines. Ethical AI asks whether something should be done at all, which is the principles layer covering values, commitments, and aspirational statements. Responsible AI asks who answers when something fails, which is the organizational practices layer where the machine checks the machine and parameters verify against parameters. AI Governance asks who decides, by what authority, and at what checkpoint, which is the external oversight layer where the human governor exercises authority at governed checkpoints. The three tiers are often conflated. The distinction that matters most for RECCLIN practitioners is between the second and third, because it is the line between machine authority and human authority.

Two models define that spectrum.

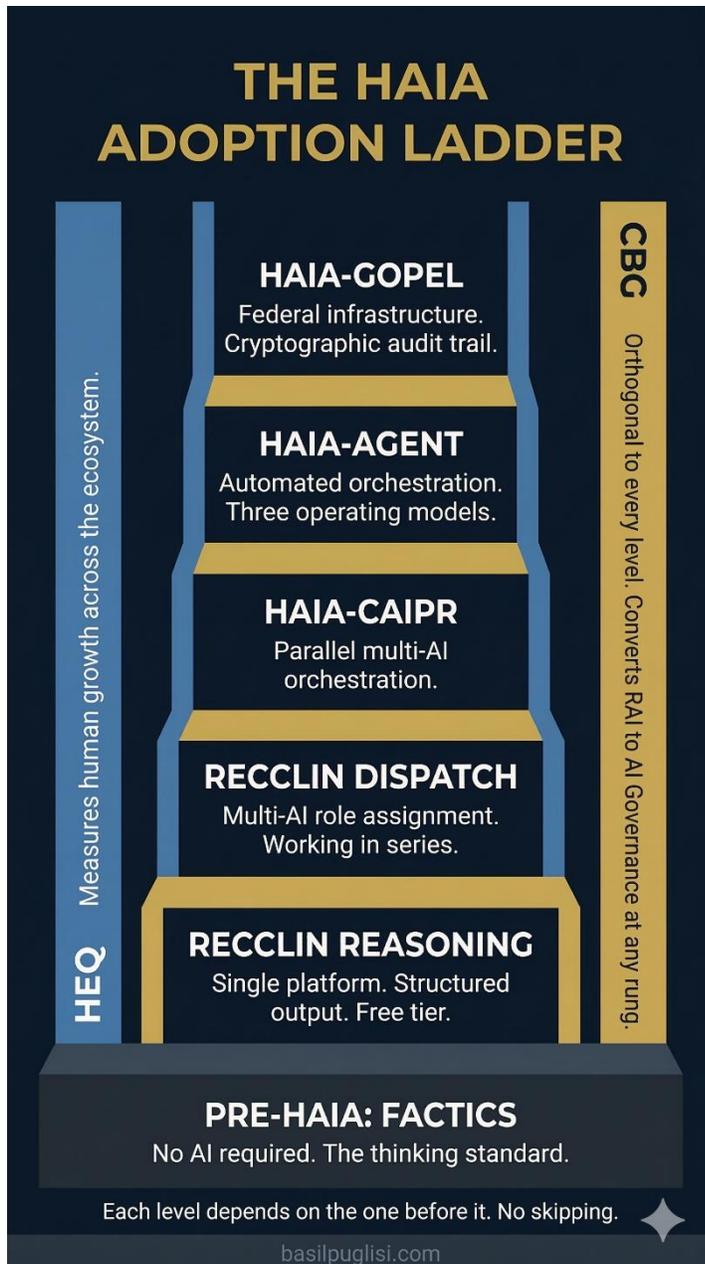


**Responsible AI (RAI)** shapes AI behavior and produces accountable outputs without requiring structural human checkpoint authority at every decision. The machine operates within defined parameters. Outputs are reviewed. Majority signals from multi-platform pools govern when consensus is achieved. The human is involved. RAI is valuable, legitimate, and appropriate for many use cases, particularly automated pipelines, agent architectures, and workflows where the risk profile does not require binding human authority at each step. RECCLIN Reasoning under RAI is a better prompt: the ten-field structure produces accountable, evaluable output regardless of whether a human governor sits at a CBG checkpoint. RECCLIN Dispatch can run RAI. A well-configured dispatch pool with good role assignments and strong Reasoning outputs produces high-quality, accountable deliverables under RAI. The operational shorthand is Factory Quality: the machine checks the machine, AI consensus is the output, and the result is consistent and auditable at production speed.

**AI Governance (AIG)** requires that human authority be structural, documented, and verified at defined checkpoints. RECCLIN plus CBG equals AIG, which means the human governor holds checkpoint authority and no AI-only output is accepted for consequential decisions. When a Dispatch output raises a question, the human may escalate to a CAIPR run, dispatching the same task to multiple platforms in parallel for convergence analysis, before making a final determination. That cross-platform review is advisory input to the human checkpoint: it informs the decision without replacing it, because the binding decision belongs to the human governor. The audit trail proves, with a structural record, that a human decided, not that a human was present while AI decided. The operational shorthand is Handmade Quality: the human knows where the compromises were made, because the human made them.

This is not a hierarchy where AIG is premium and RAI is budget. They serve different purposes. The question is whether human checkpoint authority is required for the decision at hand. Consequential federal decisions, high-risk content with public impact, and policy determinations with legal effect all require AIG. Automated content pipelines, routine research synthesis, and code generation for non-critical systems can operate under RAI.

The distinction matters most at the boundary: practitioners who believe they are operating AIG because a human reviews outputs before publication are often operating RAI if the review has no structural checkpoint, no documented decision, and no audit trail. Reviewing is not the same as deciding, and CBG is what produces the structural evidence that a decision was made.



## 11. The Broader HAIA Stack

RECCLIN does not stand alone. It is one operational methodology inside HAIA, the Human Artificial Intelligence Assistant ecosystem that structures all governed human-AI collaboration under this body of work.

The HAIA adoption ladder describes the progression from evidentiary discipline through full governed infrastructure:

**Pre-HAIA: Facts:** No AI required. The foundational thinking standard that makes everything that follows work. Every fact paired with a tactic, every tactic producing a measurable outcome.

**Layer 1: RECCLIN Reasoning:** Structured AI output showing work. Single platform. Free tier accessible. The entry point to governed AI use.

**Layer 2: RECCLIN Dispatch:** Multi-AI role assignment in series. Platform strengths applied to defined functions. Evidence-based, not assumed.

**Layer 3: HAIA-CAIPR:** Parallel multi-AI orchestration. The same prompt dispatched to multiple platforms simultaneously with no cross-platform visibility before output is produced. Convergence analysis. Hallucination detection. Synthesizer oversight. Named March 2026 in Case Study 006. Documented in the HAIA-CAIPR Specification v1.1.

**Layer 4: HAIA-Agent:** The orchestration layer that automates CAIPR mechanics while preserving CBG checkpoint authority. HAIA-Agent is a non-cognitive dispatcher that handles platform selection, prompt routing, response collection, and audit logging without performing any cognitive work. It supports three operating models: Model 1 (Agent Responsible AI) runs the full pipeline automatically with a single CBG checkpoint at the final output, suited to routine operations; Model 2 (Agent AI Governance) pauses after each RECCLIN role for human approval, suited to high-stakes decisions; and Model 3 (Manual Human AI Governance)

operates without the agent and without GOPEL entirely, with the human governor orchestrating manually and logging through the Navigator platform. All published work to date, including *Governing AI: When Capability Exceeds Control*, was produced under Model 3. HAIA-Agent is the operational proof that CAIPR orchestration can be automated. It is also the bridge that revealed why GOPEL was necessary: Agent produces complete logs, but those logs do not carry cryptographic tamper evidence or the architectural security required for infrastructure at federal scale. That gap is what GOPEL was built to close. Documented in the HAIA-RECCLIN Agent Architecture Specification at basilpuglisi.com.

**Layer 5: HAIA-GOPEL:** The governed communication infrastructure between the human governor and the AI platforms, in both directions. Every prompt dispatched and every output collected travels through GOPEL, which performs seven deterministic operations: dispatch, collect, route, log, pause, hash, report. Zero cognitive work by design. The security rationale is that a governance channel that can think can be manipulated. GOPEL moves and records without interpreting or deciding. Currently at v0.6.1, passing 183 tests across nine test suites, adversarially reviewed by seven independent AI platforms, with full code published at [github.com/basilpuglisi/HAIA](https://github.com/basilpuglisi/HAIA). Proposed as national infrastructure through the AI Provider Plurality Congressional Package.

**Orthogonal to every level: CBG:** The constitutional authority layer. The constitutional authority layer that governs the human governor, defines when human judgment is structurally required, and makes human authority verifiable rather than assumed. CBG does not sit parallel to the ladder. It is orthogonal to it, present at every layer, converting Responsible AI practice into AI Governance practice wherever it is applied.

**Measurement across the stack: HAIA-HEQ:** The Human Enhancement Quotient. Four dimensions measuring whether the human is actually growing through AI collaboration: Cognitive Agility Speed, Ethical Alignment Index, Collaborative Intelligence Quotient, and Adaptive Growth Rate. The composite Augmented Intelligence Score (AIS) is the measurement of what human and AI produce together that neither could produce independently. Documented across two companion papers at basilpuglisi.com.

Each component in this stack has a dedicated specification. This paper covers RECCLIN, and for everything else, the dedicated specifications are the authoritative sources.

## 12. Organizational Implementation

The HAIA adoption ladder describes what the components are and how they relate. This section describes how an individual practitioner, a team, or an enterprise organization moves through the ladder in practice, what resources each level requires, and what the governance output looks like at each stage.

### Individual Practitioner Path

The individual practitioner path begins with a single free-tier AI platform and a commitment to Reasoning structure before any other investment is made. The ten-field format is the entire entry point: load the Single-Session Prompt from Appendix B into any AI platform. Send a task. Evaluate the output against all ten fields. That is the practice.

The discipline required at the individual level is not technical. It is behavioral. The practitioner must resist the pull to accept an AI output because it sounds correct and instead ask whether the Sources field is populated with verifiable citations, whether the Confidence justification explains what limits the score, whether the Conflicts field documents actual search for dissent or merely states “no conflicts identified” without evidence of a search, and whether the Fact→Tactic→KPI

chain produces a complete and actionable triad. That evaluation discipline is what RECCLIN builds, and it cannot be automated.

The individual practitioner progression from Reasoning to Dispatch typically takes three to six months of consistent practice before role assignment becomes intuitive and behavioral cluster recognition becomes reliable. Practitioners who rush to Dispatch without internalizing Reasoning evaluation produce multi-platform output they cannot govern, because they have not built the evaluative capacity to distinguish between a strong ten-field output and a ten-field output that is structurally complete but evidentially hollow.

Resource requirements at the Reasoning level: one AI platform at any subscription tier, including free. Time commitment: 15 to 30 additional minutes per session for governance structure evaluation in the first month, reducing to 5 to 10 minutes per session once the format is internalized.

## **Team Deployment**

Team deployment requires that every team member practicing RECCLIN can evaluate governed output independently before the team attempts any collaborative Dispatch workflow. A team whose members have not individually internalized Reasoning evaluation cannot conduct meaningful review of Dispatch outputs at checkpoints, because the checkpoint review is only as strong as the reviewers' capacity to evaluate what the governance structure reveals.

The recommended team onboarding sequence: individual Reasoning practice for at least four weeks before any team-level Dispatch session. First collaborative Dispatch runs should use a single shared workflow with one human governor designated at each checkpoint, with other team members as observers who document their own checkpoint assessments independently. Comparison of independent assessments across team members is itself a governance exercise; it surfaces where team members' evaluative capacity diverges and what additional Reasoning practice individual members need.

Role assignment at the team level introduces a human dimension that individual practice does not: who on the team is best suited to evaluate which AI roles? A team member with strong research and citation verification skills is the appropriate reviewer for Researcher role outputs. A team member with strong editorial judgment evaluates Editor role outputs. The role matrix applies to the human reviewers as much as it applies to the AI platforms.

Resource requirements at the team level: two to five AI platform subscriptions per active Dispatch session, with role assignments documented before each session. CBG checkpoint documentation should be shared in writing across the team after each session. Time commitment: Dispatch sessions run longer than single-platform Reasoning sessions, and practitioners should plan for one to two additional hours per session at the team level, with checkpoint review consuming approximately 30% of total session time.

## **Enterprise Deployment**

Enterprise deployment requires governance infrastructure that individual and team deployment does not. The checkpoint authority structure must be codified rather than assumed: who holds the human governor role for which categories of decisions? What is the escalation path when the arbiter at a checkpoint is unavailable? What is the audit retention requirement for CBG checkpoint records? What is the platform vendor management protocol when a platform in the pool changes behavior or becomes unavailable?

These are not operational details. They are governance policy decisions that must be made before the first enterprise RECCLIN session produces output that carries organizational authority.

The enterprise deployment sequence begins with policy documentation: a written governance policy that assigns checkpoint authority, defines escalation paths, establishes audit retention requirements, and specifies the platform pool and role assignment defaults for the organization's primary use cases. That policy document is the organizational equivalent of the CBG constitutional layer; it governs the humans operating the system, not the AI platforms.

After policy documentation is complete, a pilot deployment across a limited but representative set of use cases, at least three and a maximum of five, produces the operational evidence needed to calibrate the policy before it is applied at scale. Pilot deployments should be designed to surface governance failures, not confirm governance success. The most valuable pilot output is the checkpoint decision that the human governor gets wrong, the override that should not have happened, or the convergence that should have triggered a deeper review but did not. Those failures are the governance data that makes the policy stronger.

CAIPR at the enterprise level requires GOPEL infrastructure. Manual multi-platform orchestration at nine or eleven platforms across an enterprise team is not operationally sustainable. GOPEL's seven deterministic operations, dispatch, collect, route, log, pause, hash, report, automate the mechanics of CAIPR without adding any cognitive layer to the governance process. The proof of concept is published at [github.com/basilpuglisi/HAIA](https://github.com/basilpuglisi/HAIA). The full specification is in the GOPEL Canonical Public v1.5 at [basilpuglisi.com](https://basilpuglisi.com).

Resource requirements at the enterprise level: minimum five platform subscriptions for representative Dispatch coverage; CAIPR sessions at seven to eleven platforms for high-stakes decisions; GOPEL infrastructure for systematic multi-platform orchestration; CBG governance policy documentation as the foundational prerequisite; designated human governor roles at each decision tier; quarterly behavioral profile review and platform pool calibration; annual CBG policy review aligned with AI platform evolution.

### **The Adoption Ladder in Practice**

The practical version of the adoption ladder maps against organizational readiness, not against a fixed timeline. An individual practitioner can move from Reasoning to Dispatch in three months if the practice is consistent. A 500-person enterprise moving from ad hoc AI use to governed RECCLIN deployment requires a minimum of six months for policy development and pilot deployment before any enterprise-wide rollout is appropriate.

What the ladder does not permit is skipping. CAIPR without Dispatch is parallel review with no governance-tested role assignments to anchor it. Agent without CAIPR is automation with no operational protocol to automate. GOPEL without Agent is federal infrastructure with no orchestration layer to formalize. CBG without RECCLIN is a checkpoint system with no governed output structure to review. The dependencies are not bureaucratic sequencing but operational prerequisites, because each level produces the evidence and capability that the next level requires.

None of this operates without limits. A governance framework that does not name its own boundaries is not governing; it is performing. The following section documents those boundaries.

### 13. Limitations and Challenges

This methodology has been built through operational practice, stress-tested across adversarial multi-platform review, and documented with the full record of what failed and why. Publishing those failures alongside the architecture is itself a governance principle, because a framework that does not name its limits is not governing but marketing.

**Platform access, availability, and cost.** RECCLIN Dispatch and CAIPR both depend on access to AI platforms, and not everyone has paid subscriptions to eleven platforms. Free tiers impose token limits, context windows, and rate restrictions that can degrade output quality in ways that are not always visible to the practitioner. CAIPR at nine or eleven platforms is not free, and it is not fast. The odd-number protocol exists because governance requires it, not because it is convenient. Practitioners must calibrate platform count to actual decision stakes, not default to maximum, and not compress to minimum to save money on decisions that warrant more resources. Platform downtime is a real operational constraint, because any platform in the pool can be unavailable without notice, and production runs have experienced partial platform failures mid-session. The methodology requires that the practitioner document which platforms participated and which did not, so the audit record reflects what actually happened.

**Platform updates and RECCLIN compatibility.** AI platforms update continuously. Model versions change. Behavior that was consistent in one month may shift in the next, not dramatically, but enough to affect role performance, output depth, and behavioral classification. A platform that operated as an Assembler in one period may shift toward Summarizer behavior after a model update without any announcement. The practitioner must recalibrate role assignments and best-fit defaults periodically, not assume that last quarter's assignments reflect current platform behavior. The RECCLIN structure governs what the AI must produce, but it does not guarantee that every platform will comply equally well at every point in time.

**Human governor requirements.** CBG is only as strong as the human exercising it. Presence at a checkpoint without evaluative capacity is not governance; it is a signature on a form. CBG v5.0 addresses this through scope-appropriate checkpoint assignment: human authority at the checkpoint is assumed and grounded in life experience, not credentialed or earned through prior CBG practice. Common sense proportionality, not credentials or age, is the standard. The governance design responsibility is to place the governor at decisions within a scope that matches their experience. Placing a governor at decisions outside that scope is a governance design failure, not a human authority failure. RECCLIN's structured output is the developmental mechanism that builds the governor's capacity over time, strengthening what CBG assumes from the start.

**The specialist risk.** Domain expertise is an asset in CBG governance and a liability when it becomes tunnel vision. The deeper the expertise, the more confident the specialist, and the more likely that confidence will be used to dismiss AI signals that challenge established assumptions. CBG's authority structure does not automatically prevent a specialist from using override authority to suppress correct dissent. The human must distinguish between overriding because the AI is wrong and overriding because the human is uncomfortable with a finding that is, in fact, correct. These are operationally distinct situations, and CBG documents the override so that only the practitioner can examine the record and honestly assess which kind of override was made.

**Platform independence is not guaranteed.** As the AI industry consolidates around shared foundation models and licensing arrangements, two platforms operating under different names may share sufficient underlying architecture that their agreement produces false convergence

rather than independent confirmation. The framework has no current method to structurally detect this. The odd-number protocol reduces the risk without eliminating it. Practitioners working in high-stakes contexts should treat apparent consensus as a signal worth examining, not a result worth trusting without question.

**Methodological ceiling: what multi-platform review cannot correct.** Cross-platform validation corrects at the output level but cannot correct training data, alignment tuning, or values embedded in the platform before the session begins. The WEIRD bias case study documented that all platforms in a session converged toward a recommendation reflecting Western analytical defaults. Cross-platform comparison surfaced it, but the training data that produced it remains unchanged. This framework makes the failure visible. Removing the underlying cause is a platform-level and regulatory problem, not a governance problem RECCLIN can solve.

**The HEQ evidence base.** The Human Enhancement Quotient and Augmented Intelligence Score are in active development. The formal evidence base currently sits at  $n=1$  for the longitudinal study, with preliminary cross-user testing at  $n=10$ . The instrument shows cross-platform consistency and produces meaningful directional findings. It does not yet carry the psychometric validation that enterprise deployment at scale would require for high-stakes decisions. Organizations adopting HEQ during this period become validation partners, and the staged validation roadmap is documented in the companion papers.

### **Framework Validation Status**

- **OPERATIONALLY VALIDATED:** Content creation, research synthesis, and policy production through sustained proof (204-page manuscript, 50+ case studies, HEQ quantitative framework).
- **ARCHITECTURALLY TRANSFERABLE:** Governance methodology applicable to coding, legal analysis, financial modeling, and engineering design pending context-specific testing.
- **PROVISIONAL:** Enterprise scalability and multi-organizational performance pending external validation.

### **14. Closing Argument**

RECCLIN Dispatch is flexible by design. It can run under Responsible AI in automated pipelines, agent architectures, and code-driven workflows where human checkpoint authority is not required at every step. The dispatch mechanism assigns roles, routes tasks, and collects structured outputs. It does this reliably whether a human is at a CBG checkpoint or not. Dispatch serves RAI well.

RECCLIN Reasoning applies across both governance models, and the practitioner must know which one is operating.

Reasoning under Responsible AI is a better prompt. The ten-field structure forces the AI to declare its role, confirm the task, cite sources, flag conflicts, score confidence, assess expiry, produce a Fact→Tactic→KPI chain, and separate the recommendation from the evidence. That structured output is superior to an unstructured prompt regardless of whether a human governor sits at a CBG checkpoint. A practitioner using Reasoning at the Responsible AI level produces accountable, evaluable AI output. Dissent routes to other platforms for cross-validation and is documented in the record. AI consensus can govern the output, and this is legitimate and valuable.

Reasoning under AI Governance is a different instrument. The same ten fields now feed a CBG checkpoint where a human governor exercises binding authority. Cross-platform validation may still occur, as the human can escalate to a CAIPR run before making a final determination, but that validation is advisory input to the checkpoint, not a substitute for it. The audit trail proves a human decided, not merely that a human was present. The structure is identical, but the governance weight behind it is categorically different.

Both are valid deployments of the methodology. The practitioner decision is to know which model is operating, state it, and govern accordingly. Dispatch serves both models by the same logic. The difference between having a human present and having a human decide is the difference between Responsible AI and AI Governance.

That line is what CBG holds, and that line is what this methodology was built to protect.

## Appendix A: RECCLIN Output Format Reference

The following is the practitioner reference for RECCLIN Reasoning output structure. This format applies to every governed AI output, regardless of platform, role, or workflow configuration.

**Role:** [The function the AI is performing for this task: Researcher, Editor, Coder, Calculator, Liaison, Ideator, or Navigator]

**Task:** [The request repeated back in clean language, confirming the AI understood the ask correctly]

**Output:** [The substantive response, the answer, analysis, code, calculation, or synthesis the AI was tasked to produce]

**Sources:** [Cited evidence with sufficient detail to verify: URLs, publication titles, authors, dates. Every source must be verifiable. Unverified sources are flagged as PROVISIONAL.]

**Conflicts:** [Dissent identified in sources, or disagreement with prior outputs in the workflow. If no conflicts are found, state that explicitly. Absence of documented conflict search is not the same as absence of conflict.]

**Confidence:** [Score from 0 to 100 percent with written justification based on evidence quality. Justify the score: what evidence supports it, what gaps limit it.]

**Expiry:** [The lifecycle of this output, how long it remains valid given the time-sensitivity of the information it contains. Stable information is noted as such. Time-sensitive outputs carry an explicit expiry assessment.]

**Fact→Tactic→KPI:** [The primary finding expressed as a complete chain: Fact (the evidence-grounded claim), Tactic (the executable action that follows), KPI (the measurable outcome that proves the tactic worked)]

**Recommendation:** [The path the AI believes the human should follow, stated separately from the evidence so the human can evaluate both independently. This is the AI's synthesis.]

**Decision:** [The structured handoff to the human governor. State the specific choice requiring human judgment with options framed clearly: accept or challenge, path A or path B, D1 or D2 or D3 when platforms produced competing recommendations. The AI frames the decision. The human makes it. The human's ruling is the CBG checkpoint record; it is not part of this output.]

## Appendix B: Loading RECCLIN into Your AI Platforms

*Verified against platform documentation March 2026. UI navigation paths require quarterly review. Platform interfaces change frequently. When in doubt, use the Single-Session Prompt, which works on any platform regardless of UI changes.*

### Why This Appendix Exists

RECCLIN Reasoning produces governed AI output. For that to happen consistently, the AI platform must receive the governance instruction before it produces output, not after. The single most common implementation failure is asking an AI to retroactively format output it already produced without governance structure. That is not RECCLIN. That is reformatting.

Persistent instructions solve this. When the governance format is loaded into a platform's memory or custom instruction layer, every session starts in RECCLIN mode without requiring

the practitioner to re-issue the instruction manually. The AI governs its own output format from the first response.

This matters more than it appears. Case Study 006 (March 2026) tested eleven platforms after they had been asked to store the RECCLIN framework in memory. When those platforms received a prompt that did not explicitly invoke RECCLIN, only five of the eleven activated full governance output from stored memory. The remaining six produced prompt-only responses. Memory does not guarantee activation, and the implementation method matters.

The instructions that follow cover six platforms with verified persistent instruction capabilities. Five additional platforms in the active RECCLIN pool (Kimi, DeepSeek, Meta AI, Copilot, and MiniMax) do not currently support reliable persistent custom instructions and should use the Single-Session Prompt at the end of this appendix.

### **Why Persistent Instructions Are Worth the Setup**

A practitioner who loads RECCLIN into custom instructions on ChatGPT and Claude's Personal Preferences is not doing administrative work. They are creating a governed AI workforce where every output carries accountability structure without additional effort per session. The ten fields, Role, Task, Output, Sources, Conflicts, Confidence, Expiry, Fact→Tactic→KPI, Recommendation, Decision, require the AI to show its work every time, regardless of what the practitioner asks. The structure is always present. The practitioner's job is to evaluate the output, not to remember to request accountability.

For practitioners building toward Dispatch or CAIPR workflows, persistent instructions also create platform consistency. When each platform in the pool is loaded with the same RECCLIN output standard, outputs are structurally comparable across platforms. Convergence analysis and dissent detection both depend on outputs that carry the same accountability fields. A platform producing free-form responses alongside RECCLIN-structured outputs cannot be compared meaningfully at the governance level.

### **Platform 1: ChatGPT (OpenAI): The Standard**

ChatGPT's Custom Instructions feature is the most direct implementation pathway for RECCLIN. Instructions loaded here persist across all conversations and activate automatically at session start.

**Important:** ChatGPT's Custom Instructions fields have a 1,500-character limit per field. The full RECCLIN instruction text exceeds this limit. Two implementation paths are available: the Compressed Instruction for Custom Instructions (fits within the limit), and the Projects path for the full instruction text. New practitioners should start with the Compressed Instruction. Practitioners managing ongoing governance work should use Projects.

#### **Navigation path:**

- Web: Settings → Personalization → Custom Instructions
- Desktop and mobile: Settings → Customize ChatGPT

There are two text fields: the first asks what ChatGPT should know about you, the second asks how it should respond. The RECCLIN instruction goes in the second field.

#### **First field (About You), adapt to your context:**

I am a [your role, such as professional, student, researcher, executive] using HAIA-RECCLIN governance across AI platforms. I work with structured workflows where human judgment is central to all decisions. I use the Factics methodology: Facts paired with Tactics and measurable KPIs. I compare outputs across multiple AI platforms for validation and dissent detection.

**Second field (How to Respond), Compressed RECCLIN Instruction (fits within 1,500 characters):**

ALWAYS respond in Full Governance format. No exceptions unless I say “Answer Only.”

Every response must include these fields in order:

Role: [Researcher/Editor/Coder/Calculator/Liaison/Ideator/Navigator. Self-assign based on task]

Task: [My request repeated back in clean language]

Output: [Your substantive response]

Sources: [Verifiable citations. Flag unverified as PROVISIONAL.]

Conflicts: [Dissent or disagreements. State “No conflicts identified” if none. Do not skip.]

Confidence: [0–100% with written justification]

Expiry: [How long this output remains valid]

Fact→Tactic→KPI: [Fact → Tactic → KPI chain]

Recommendation: [Primary + alternatives, separate from evidence]

Decision: [Options for my human arbitration: accept/challenge, A/B, or D1/D2/D3]

Principles: Preserve dissent. You suggest. I decide. Human authority is absolute.

Answer Only: Direct response, no governance structure. Request “reissue in Full Governance format” anytime.

**Advanced path, ChatGPT Projects (no character limit):**

For practitioners who want the full instruction text without compression, ChatGPT Projects support longer persistent instructions. Create a new Project, open Project Settings, and paste the full RECCLIN instruction text from the Single-Session Prompt section at the end of this appendix into the instructions field. All conversations within the Project will operate under full RECCLIN governance.

**Verification:** After loading, open a new conversation and ask: “What is 2+2?” If all ten fields appear in the response, RECCLIN is active. If the response is unstructured, the instruction did not load correctly. Return to Settings and confirm the text was saved.

**Platform 2: Claude (Anthropic)**

Claude uses a Personal Preferences feature for account-wide persistent instructions, accessible through Settings. For project-level use, Claude Projects allow instructions to be stored alongside uploaded documents, making it the most powerful RECCLIN implementation for practitioners managing ongoing governance work.

**Navigation path:**

- Account-wide: Select your initials in the lower left → Settings → enter your preferences in the “What preferences should Claude consider in responses?” field
- Project-level: Create or open a Project → click “Set project instructions” → paste the instruction → Save instructions

**Note on Projects access:** Project-level instructions may require a paid plan on some account tiers. Verify your account access before relying on Projects as your primary loading path. Free-tier practitioners should use the account-wide preferences field.

**Instruction text:** Use the full RECCLIN instruction from the Single-Session Prompt section at the end of this appendix, all ten fields including Decision, with the following addition placed at the very top of the instruction:

CRITICAL: Before every response, ask me to choose output mode:

Output mode? 1. Full Governance 2. Answer Only, and wait for my selection before proceeding.

This mode-selection line takes precedence over the always-on Full Governance directive in the body of the instruction. Claude will pause before every response and confirm the output mode. Both modes are fully governed, with Answer Only suppressing the field structure when the practitioner wants a direct response. To retrieve the full governance structure for any Answer Only response, ask Claude to “reissue in Full Governance format.”

**Verification:** After loading, start a new conversation. Claude should immediately ask: “Output mode? 1. Full Governance 2. Answer Only” before producing any response. If it does not, the instruction did not load. Return to Settings and confirm the text was saved.

### Platform 3: Gemini (Google)

Gemini offers two implementation paths: a Gem for session-specific governance (recommended), and global instructions for account-wide use.

**Geographic restriction:** The global instructions feature (Instructions for Gemini) is not available in the European Economic Area, Switzerland, or the United Kingdom. Practitioners in those regions should use Gems or the Single-Session Prompt.

**Path A, Gem (recommended):** In Gemini, go to Explore Gems → New Gem. Name it “HAIA-RECCLIN Governance.” Paste the full RECCLIN instruction into the instructions field and save. To use a governed session, open this Gem from the sidebar rather than starting a standard Gemini chat. The Gem only governs sessions launched through it; standard Gemini conversations are not affected.

**Path B, Global instructions (account-wide, where available):** Go to Settings and help → Personal Intelligence → Instructions for Gemini. Paste the full RECCLIN instruction there. Instructions loaded here apply to all Gemini conversations on your account.

**Instruction text:** Use the full RECCLIN instruction from the Single-Session Prompt section. No modification required from the standard format.

**Verification:** After loading, launch a session through your Gem (or open a standard conversation if using global instructions) and ask: “What is 2+2?” All ten fields should appear. If the response is unstructured, confirm you launched via the Gem rather than a standard chat window.

## Platform 4: Grok (xAI)

Grok's persistent instruction mechanism operates through its memory system rather than a dedicated Custom Instructions settings panel. The most reliable method for loading RECCLIN governance is a direct memory command at the start of a session.

**How to load:** Open a new Grok conversation and send this message:

Remember that I use HAIA-RECCLIN governance in all our sessions. Always respond in Full Governance format with these ten fields in order: Role, Task, Output, Sources, Conflicts, Confidence, Expiry, Fact→Tactic→KPI, Recommendation, Decision. Preserve dissent. You suggest. I decide. Human authority is absolute.

Grok will store this instruction in persistent memory and apply it across future sessions. For high-stakes sessions, confirm activation at session start with the mode confirmation prompt.

Grok showed strong RECCLIN memory activation in Case Study 006, activating all ten governance fields from stored memory without explicit prompting in that session, more reliably than most platforms in the pool.

**Verification:** After sending the memory command, ask: "What is 2+2?" All ten fields should appear. If they do not, resend the memory command and confirm Grok acknowledged it before proceeding.

## Platform 5: Perplexity

Perplexity supports persistent governance instructions through two mechanisms: the AI Profile for account-wide preferences, and Spaces for project-level custom instructions with knowledge sources.

**Path A, AI Profile (account-wide):** Go to your profile icon → Settings → AI Profile. Paste the full RECCLIN instruction into the profile field. Instructions here apply across all standard Perplexity conversations.

**Path B, Spaces (project-level, recommended for Dispatch and CAIPR workflows):** Create a new Space, open its settings, and paste the full RECCLIN instruction into the Custom AI Instructions field. Spaces allow you to store instructions alongside uploaded documents and knowledge sources, making them well-suited for ongoing governance projects where source material is also managed.

**Instruction text:** Use the full RECCLIN instruction from the Single-Session Prompt section. One supplemental addition recommended for Perplexity specifically, given its strength in source validation:

For every source cited, provide: summary of the source, the fact or data from that source, the tactic or strategy from that source, and the outcome or KPI from that source. Every source must carry a complete Fact→Tactic→KPI chain to be usable.

This supplemental instruction activates the original Perplexity Dispatch method that initiated multi-AI governance in 2023. It is an enhancement to the Sources field, not an additional governance field.

**Verification:** After loading, ask: "What is 2+2?" All ten fields should appear. For AI Profile, confirm the profile was saved before testing. For Spaces, confirm you are working within the Space, not a standard Perplexity conversation.

## **Platform 6: Mistral (Le Chat)**

Mistral's Le Chat implements persistent instructions through Agents, named AI configurations that carry custom instruction sets. The Agents path is the documented and supported mechanism for persistent RECCLIN governance on Le Chat. For API access, the instruction goes in the system prompt field.

**How to load, Le Chat:** Go to Agents in the left sidebar → Create Agent. Name the agent "HAIA-RECCLIN Governance." Paste the full RECCLIN instruction into the Agent's Instructions field and save. To use a governed session, launch a conversation through this Agent. Conversations started outside the Agent are not governed by these instructions.

**How to load, API:** Place the full RECCLIN instruction in the system prompt field of your API call. It will apply to all messages in that session.

**Instruction text:** Use the full RECCLIN instruction from the Single-Session Prompt section. Given Mistral's observed tendency toward compressed outputs in Case Study 006, add this line:

Do not summarize or compress your output. Produce full-depth responses in every field. Brevity in the Output field is not a virtue if it sacrifices governance completeness.

Mistral showed prompt-driven rather than memory-driven behavior in Case Study 006, which means persistent instruction loading through Agents is more important for Mistral than for platforms with stronger memory activation.

**Verification:** After creating the Agent, launch a conversation through it and ask: "What is 2+2?" All ten fields should appear. If the response is unstructured, confirm you launched via the Agent rather than a standard Le Chat conversation.

### **Single-Session Prompt: For Any Platform**

For platforms without persistent instruction support, for practitioners not yet ready for full custom instruction setup, or for any CAIPR dispatch session where memory activation cannot be assumed, paste this prompt at the start of the session as message one. Send it alone. Wait for the mode confirmation. Then send your actual task as message two.

### **RECCLIN Governance Mode: Session Prompt**

You are operating under HAIA-RECCLIN governance for this session. HAIA-RECCLIN is a human-AI governance framework developed by Basil C. Puglisi, MPA. Your role is to produce structured, accountable output that makes your reasoning visible and verifiable.

Every response you produce must include these ten fields in order. The initial mode confirmation below is a single-line handshake; ten-field output begins after mode is confirmed.

**Role:** Self-assign from: Researcher, Editor, Coder, Calculator, Liaison, Ideator, Navigator. Choose the function that best matches this task.

**Task:** Repeat my request back in clean language confirming you understood it correctly.

**Output:** Your substantive response.

**Sources:** Citations with sufficient detail to verify. Flag unverified sources as PROVISIONAL.

**Conflicts:** Dissent in sources or disagreement with prior outputs. If none found, state "No conflicts identified." Do not skip this field.

**Confidence:** 0 to 100 percent with written justification based on evidence quality.

**Expiry:** How long this output remains valid. State whether information is stable or time-sensitive.

**Fact→Tactic→KPI:** Fact (evidence-grounded finding) then Tactic (the executable action that follows) then KPI (the measurable outcome that proves the tactic worked).

**Recommendation:** Your primary recommendation plus at least one alternative. Separate your synthesis from the evidence so I can evaluate both independently.

**Decision:** The specific choice requiring my human arbitration. Frame the options clearly: accept or challenge, option A or option B, or present the competing recommendations from this session for my selection (D1, D2, D3, etc.). This is your structured handoff to me. You frame the decision. I make it.

Governance principles: You suggest. I decide. Preserve dissent. Never force consensus. Human authority is absolute.

If I say “Answer Only,” provide a direct response without the governance structure. Request “reissue in Full Governance format” at any time to receive the full ten-field output for any prior response.

Confirm you are in RECCLIN governance mode by asking: “Output mode? 1. Full Governance 2. Answer Only.” Then wait for my selection before producing any output.

### **A Note on Memory Reliability**

Loading these instructions is the starting point, not the guarantee. Case Study 006 documented that fewer than half of the platforms in the pool activated full governance output from stored memory when the session prompt did not explicitly invoke RECCLIN. The practical implication: for high-stakes sessions, use the Single-Session Prompt as message one even on platforms where custom instructions are already loaded. The mode confirmation costs nothing and verifies that governance is active before output is produced.

For CAIPR sessions, always include the full governance instruction in the dispatch prompt. Do not rely on memory activation for parallel multi-platform review. Memory activation cannot be assumed across all platforms simultaneously, and a platform running without the governance structure active produces output that cannot be compared at the governance level with platforms that are structured.

### **Maintenance Schedule**

The UI navigation paths in this appendix were verified against platform documentation in March 2026. Platform interfaces change frequently. This appendix requires review every quarter, and immediately following any major platform release that affects settings, personalization, agents, instruction mechanisms, or project features. The governance instruction text and field structure are stable; the navigation paths are not.

## References and Related Documents

Puglisi, B. C. (2025). *Governing AI: When Capability Exceeds Control*. ISBN 9798349677687. Amazon.

Puglisi, B. C. (2025). *The Human Enhancement Quotient (HEQ): Measuring Collaborative Intelligence for Enterprise AI Adoption*. White Paper v4.3.3. basilpuglisi.com/HEQ

Puglisi, B. C. (2026a). *Measuring Augmented Intelligence: Theoretical Foundations and Empirical Development of the Human Enhancement Quotient (HEQ) and Augmented Intelligence Score (AIS)*. Working Paper v2.5. basilpuglisi.com/measuring-augmented-intelligence

Puglisi, B. C. (2026b). *Checkpoint-Based Governance: A Constitution for Human-AI Collaboration*. CBG v5.0. basilpuglisi.com

Puglisi, B. C. (2026c). *HAIA-CAIPR Specification v1.1*. basilpuglisi.com

Puglisi, B. C. (2026d). *HAIA-RECCLIN Agent Architecture Specification*. EU Compliance Version. basilpuglisi.com

Puglisi, B. C. (2026e). *GOPEL: The Code Behind the Policy*. Proof of Concept v3.1. basilpuglisi.com. github.com/basilpuglisi/HAIA

Puglisi, B. C. (2026f). *AI Provider Plurality Congressional Package*. Documents 1 through 5. basilpuglisi.com

Puglisi, B. C. (2026g). *HAIA-RECCLIN Case Study: The Kimi Outlier*. December 2025. basilpuglisi.com

Puglisi, B. C. (2026h). *HAIA Framework Architecture Map v1.8*. March 2026. basilpuglisi.com. github.com/basilpuglisi/HAIA

Puglisi, B. C. (2026i). *HAIA-CORE: The Missing Piece in Content Evaluation*. basilpuglisi.com

Puglisi, B. C. (2026j). *HAIA-SMART v1.5 Calibration*. basilpuglisi.com

Puglisi, B. C. (2026k). *What We Learned: HAIA Multi-AI Practice*. basilpuglisi.com

Puglisi, B. C. (2026l). *The Loop That Ate the Governor*. Case Study. basilpuglisi.com

Puglisi, B. C. (2026m). *HAIA-RECCLIN Case Study 006*. March 2026. basilpuglisi.com

Puglisi, B. C. (2026n). *HAIA Complete Workflow White Paper v1.0*. basilpuglisi.com

Puglisi, B. C. (2012). *Digital Factics: Twitter*. Digital Media Press. magcloud.com/browse/issue/471388

**Basil C. Puglisi**, MPA Human-AI Collaboration Strategist me@basilpuglisi.com | basilpuglisi.com | github.com/basilpuglisi/HAIA

*Third Edition | March 2026*

© 2026 Basil C. Puglisi. All rights reserved. Human Enhancement Quotient, HEQ, Augmented Intelligence Score, AIS, HAIA, HAIA-RECCLIN, HAIA-CAIPR, GOPEL, Factics, and CBG are trademarks of Basil C. Puglisi. Any use in research, publications, training materials, or commercial applications requires proper attribution. Commercial use requires written authorization.

*#AIassisted under HAIA-RECCLIN governance. Human governor review and approval required before any external distribution, publication, or application.*