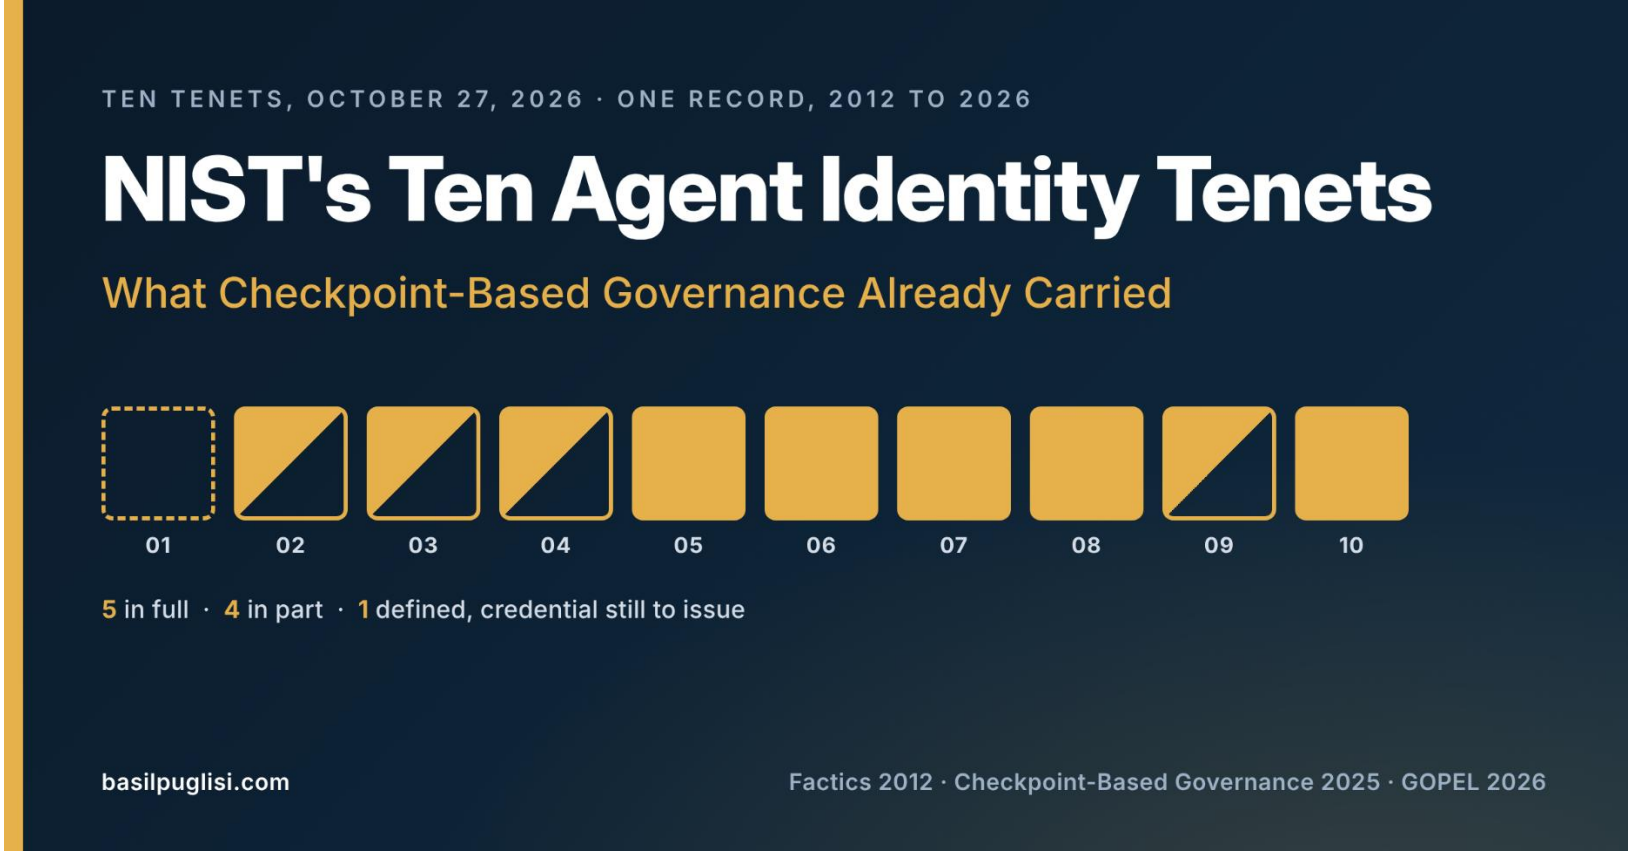


NIST's Ten Agent Identity Tenets: What Checkpoint-Based Governance Already Carried

Basil C. Puglisi, MPA · A Human-AI Collaboration · basilpuglisi.com · November 2026

On October 27, 2026, the Cybersecurity Insights blog at the National Institute of Standards and Technology published ten design tenets for agent identity management systems. Harold Booth, Bill Fisher, and Ryan Galluzzo of NIST wrote the post with Tom Clancy of MITRE. It is the second entry in a series on agentic identity, authentication, and authorization. The tenets draw on public comments to the National Cybersecurity Center of Excellence project on agent identity and authorization and on engagement with the field. The authors present them as an initial set, to be applied according to the risk of the transactions an organization supports.

The record those tenets meet did not begin in February 2026, and it did not begin with code. It begins with the evidence discipline the author first stated publicly in 2012 and completed with measurement from 2013, which the author calls Factics. Its governing line reads: "Every fact must lead to a tactic, and every tactic must leave evidence." It runs through Checkpoint-Based Governance, made public on September 23, 2025, which puts a named human at the decision and leaves a record of what they chose. The February 2026 AI Provider Plurality Congressional Package then built that governance into an agent. At its center sat a non-cognitive enforcement layer that dispatches, collects, routes, logs, pauses, hashes, and reports, which the author calls GOPEL, the Governance Orchestrator Policy Enforcement Layer. The agent architecture specification behind the package was published on February 3, 2026, the package page on February 8, and working reference code on February 23. By October 2026 that code had advanced to Alpha releases for two agent models. A Responsible AI content desk, Agent Model 1 in production, was publishing on basilpuglisi.com and puglisiconsulting.com with no human review of any post.

This article lays the ten tenets against that record, with Checkpoint-Based Governance as the basis, the February work as its implementation, and the running desk as its evidence of use. Nine of the ten have a counterpart in the record, five in full and four in part. The tenth has its definition in the record and not yet its credential, and that credential is the work now under construction.

What NIST asks for

The ten tenets answer five questions. The first is who the agent is and for whom it acts. Tenet 1 asks for a distinct, cryptographically verifiable identity for each agent, tied to a workload, an environment, and an approved configuration. The credential lifecycle stays separate from any human principal. Tenet 2 asks that the agent stay cryptographically linked to the person or organization that tasked it, with delegated scope narrowing as authority passes to sub-agents. Tenet 6 asks for signed mandate statements from the task originator that define intended outcomes and constraints and can be verified downstream.

The second question is what the agent may hold, and for how long. Tenet 3 asks for zero standing privilege and separation of duties. Tenet 4 asks for short-lived, revocable grants and continuous runtime verification, so that a hijacked agent can be stopped in the middle of a task.

The third question is where the human stands. Tenet 5 asks organizations to define risk thresholds that trigger human review and to route escalations to the right authority with enough context to decide. It also asks them to design against consent fatigue, the point at which frequent escalation turns review into habit.

The fourth question is how the architecture is divided. Tenet 7 asks that the logic deciding what an agent attempts be separated from the systems deciding what it is allowed to do. The authors' reason is that a language model accepts data and instruction as a single stream of tokens through a single interface. The separation therefore has to be enforced by architecture. They put the consequence in four words: "Guardrails alone are insufficient." Tenet 8 asks that access decisions stay deterministic, with probabilistic signals as inputs and never as the deciding mechanism.

The fifth question is how the agent is watched and recorded. Tenet 9 asks for monitoring built on agent-specific metrics and thresholds. Tenet 10 asks for an append-only, tamper-evident body of evidence covering the agent's identity, its delegated intent, and its actions, kept on infrastructure the agent cannot reach.

The authors are plain about status. The tenets are not a standard, several cannot yet be met with current technology, and they will guide the NCCoE demonstrations to come.

Where the record begins

Factics comes first because it needs no AI. The framework page states the dependency the rest of the record rests on: "None of them works without the evidence discipline underneath, which is why Factics needs no AI and comes first." A fact that leads to no tactic is an observation, and a tactic that leaves no evidence is an assertion. Every later layer applies that discipline to AI work. That includes the structured output a single platform must show before its answer carries weight, and the record a named human leaves at a checkpoint.

Checkpoint-Based Governance made the human authority layer public on September 23, 2025, as "a protocol-driven framework for structuring human-AI collaboration." That first public record already carried the four-stage loop and the rule that "Every checkpoint decision generates a record including the input evaluated." It set risk-proportional deployment, with a checkpoint at "Every decision point before irreversible actions" under

heavy governance. And it named the drift signal the agent models would later have to watch for: "Increasing approval rates may indicate automation bias."

The framework page today carries the title Presence Is Not Authority and the settled form of the rules. The human governor accepts, modifies, or rejects, and "These are the only three checkpoint outcomes." The decision is documented with the governor's identity, rationale, and timestamp, "and the record is immutable." Agreement among machines decides nothing: "No platform count, confidence score, or level of agreement substitutes for human arbitration." The failure mode that matters most is named plainly: "Passive acceptance is the one that hollows out the checkpoint without anyone noticing." The page also places the enforcement layer where NIST's seventh tenet would later put it. "GOPEL is a non-cognitive enforcement layer that would record and enforce checkpoints mechanically, without reasoning." The paper behind the page adds the three checkpoint phases. In the BEFORE phase the human governor establishes scope, criteria, and constraints, and no AI execution begins without it. In the DURING phase the governor monitors with authority to intervene, redirect, or terminate. In the AFTER phase the governor validates the output against the BEFORE criteria and accepts, modifies, or rejects it. The paper's prohibition is constitutional rather than operational: "AI cannot satisfy, close, or validate a checkpoint for another AI."

That is the basis. Everything February built sits on it.

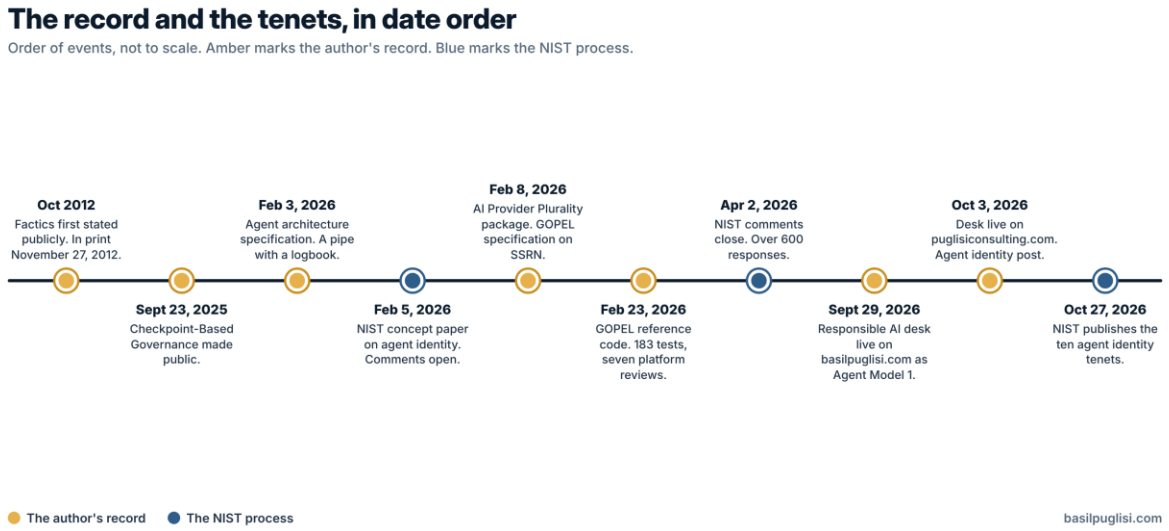


Figure 1. The record and the tenets, in date order. Amber marks the author's record and blue marks the NIST process. Order of events, not to scale.

What February built on it

The agent architecture specification of February 3, 2026, described its agent as "a traffic controller" that "performs zero cognitive work." In the phrase that carried into every later document, the agent is "a pipe with a logbook." The human layer in that specification exercises Checkpoint-Based Governance authority at defined pause points. The agent sits between the AI layer and the human layer as a mechanical orchestrator. Its first design principle was record-keeping first: "It is a logging system that also routes. The audit trail is the product." Every transaction was to leave six record types, Request, Dispatch, Response, Navigation, Arbitration, and Decision, in an append-only file that any auditor could query.

The specification fixed the division of labor in its Agent Neutrality Principle. Every non-mechanical decision, from the platform pool to the escalation thresholds, is "a human-supplied configuration constant. The agent executes these constants. It does not decide them." Its security section then required code integrity by hash against the approved version and separation of duties between whoever writes the code and whoever approves its deployment. It required immutable deployment with no hot-patching and audit files "stored in a location separate from the agent's execution environment." It required authenticated operator identity on every arbitration record, so that the person who decided "cannot plausibly deny the decision." It carried the automation bias trigger from Checkpoint-Based Governance as numbers: approval rates above 95 percent or reversals below 2 percent for three cycles start an audit and move the work to a denser checkpoint model.

Twenty days later, on February 23, 2026, the reference code was published. It ran, passed 183 tests across nine test suites, and held the non-cognitive constraint across all 14 source modules. It survived adversarial review run on seven AI platforms, each working without sight of the others' findings. It carried HMAC-SHA256 operator identity, a hash-chained audit trail, checkpoint gates that would not advance without human approval, and second-approver signatures. It added an escalation engine that ratchets posture upward and requires explicit authority to bring it down, with counters for approval and reversal rates. The article said what the code was not at the time: a reference implementation on mock adapters, and "It isn't deployment-ready software." The package's Technical Appendix set identity binding for federal pilots on PIV or CAC credentials and named the objective as non-repudiation.

What was running before October 27

The February code did not stand still. In October 2026 the author published Alpha releases for two agent models in the HAIA repository. The Agent Model 1 Runner is Responsible AI written as code. The Agent Model 2 Gateway is the hybrid that runs through to AI Governance at a signed human checkpoint after every role. GOPEL moved into the same folder beside them. Those two programs remain reference code, and the paper that introduced them says so. Neither has been deployed, and neither has called a live AI provider in the build record.

The production use runs beside them, and it is Agent Model 1. The author's Responsible AI content desk runs on Grok Bot as Navigator, with parallel review seats on the Claude API and the ChatGPT API. It publishes under the HAIA Agents byline, an author record that reads "HAIA Agent Model 1 Responsible AI" on basilpuglisi.com. Machine checks machine, no named human holds a checkpoint over any single post, and every post carries the #AIgenerated label for that reason. The desk is a separate build from the Alpha Runner, and its run shows the model in use rather than the Alpha code. The archive records what the desk published and when, and the checks behind each post rest on the author's own operating record rather than an outside audit. The desk was in use before the NIST post appeared. Its first post on basilpuglisi.com is dated September 29, 2026. Its first posts on puglisiconsulting.com are dated October 3, 2026, and one of them is an article titled Workflow and AI Deployment: Agent Identity Before Scale.

That October 3 post is worth naming for what it is and what it says. It is an automated Responsible AI output, machine checking machine, and no human reviewed it before it went live. It argued that production AI agents should each have a distinct identity, a named human owner, task-scoped tokens, and audit records linking the user, the agent, and the tool. Those are tenets 1, 2, 3, 4, and 10, published by the author's own agent desk twenty-four days before NIST published them. The deployer answers for that post under the rule the author set for Agent Model 1, and the record of it stands on its own date.

Agent Model 3 has the longest run of all. It is the manual form: the human dispatches, collects, routes, and decides at every step, with no software agent. It is the way the author has worked with AI since 2022. The Factics page dates the application to AI from 2022 onward. When early model outputs arrived without sources, the fix was "a second prompt sent after every output requiring the facts, tactics, KPIs, and sources behind the answer." Independent source verification through a second platform followed. The page for the author's reasoning and dispatch method, which the author calls HAIA-RECCLIN, dates the second step: "In 2023, one platform produced strong answers and unreliable citations, so the sourcing job moved to a platform that handled citations well." The February specification records that this manual form "produced the Governing AI manuscript" in 2025, and the Agents page carries it as the model in which the human "decides at every step." Agent Model 3 built the Alpha code and runs the author's own published work to this day. The deployed desk is Responsible AI. The hand-run practice behind every paper named here, this one included, is AI Governance, and it has been in use for years.

The tenets against the record

Direct means the requirement is present in the basis and carried as a mechanism in the implementation record, from the February specification and code through the October Alpha. Partial means the requirement is present as a principle, or in one of its parts, without the mechanism NIST describes.

NIST tenet, October 27, 2026	Checkpoint-Based Governance basis	Implementation, February to October 2026	Fit
01 Distinct agent identification	An agent is "whatever acts to reach a final output," and "what separates them is where the named human stands"; "an AI system does not act on its own" (Agents page)	Code integrity: every deployment verifies the running agent's hash against the approved version. No credential for the agent itself, and none presented outside the deployment	Definition present, credential absent
02 Agent linking, delegation, and attenuation	The decision carries the governor's identity, rationale, and timestamp; no AI approves another AI's output, so no delegated decision can exceed what the named human signed	Authenticated operator identity on every arbitration and decision record; non-repudiation as the stated objective; PIV or CAC in federal pilots	Partial: direct on the human link, absent on a scope rule for sub-agents
03 Least and zero standing privilege	Risk-proportional deployment: checkpoint density assigned by consequence at the design stage	Separation of duties among author, approver, and auditor; the agent holds human-supplied constants and cannot change them	Partial: separation of duties present, zero standing privilege not addressed
04 Short-lived, dynamic, revocable authorization	The DURING phase: the governor monitors with authority to intervene, redirect, or terminate; dynamic escalation when risk indicators emerge in a live session	Fail-closed halt on any failed verification; a circuit breaker; resumption only through a separately authorized human act	Partial: direct on the authority to halt, absent on just-in-time, time-bounded grants
05 Human-in-the-loop triggers	"Increasing approval rates may indicate automation bias" (September 23, 2025); "Passive acceptance is the one that hollows out the checkpoint without anyone noticing"; Presence Is Not Authority	The Pause operation delivers a governance package to the human; the 95 percent and 2 percent triggers over three cycles; escalation to a denser checkpoint model	Direct

NIST tenet, October 27, 2026	Checkpoint-Based Governance basis	Implementation, February to October 2026	Fit
06 Mandates and constraints	The BEFORE phase: scope, criteria, and constraints set by the human governor before any AI execution; a checkpoint at "Every decision point before irreversible actions"	Every governance artifact set by human authority outside the agent and loaded as configuration; configuration treated as a governed artifact with version history; from October 2026, a charter signed by a named human for the Runner and a signed manifest for the Gateway	Direct
07 Architectural separation of policy from reasoning	"GOPEL is a non-cognitive enforcement layer that would record and enforce checkpoints mechanically, without reasoning"; Checkpoint-Based Governance defines the requirements, GOPEL enforces them	The non-cognitive constraint: the agent performs zero cognitive work, and the Navigator, the only reasoning function, sits outside it	Direct
08 Deterministic access decisions	"No platform count, confidence score, or level of agreement substitutes for human arbitration"	Deterministic policy application; a static analyzer that rejects any cognitive operation in the code	Direct on the enforcement layer, with human arbitration above it
09 Continuous monitoring	Passive acceptance must be detectable; the thresholds and cycles belong to GOPEL	The Report operation; breach detection across five severity levels; a biennial review cycle in the Technical Appendix	Partial: approval, reversal, and breach metrics present, runtime monitoring of the agent's own access absent
10 Immutable evidence	"Every checkpoint decision generates a record" (September 23, 2025); "the record is immutable"; Why the Record Is the Proof; Factics: every tactic must leave evidence (2012)	Append-only, hash-chained, signed audit trail; six record types; audit files stored apart from the agent; the Decision Record links every upstream record	Direct on the chain, with the agent credential it would carry waiting on tenet 1

Where the three documents read as one

Tenet 7 is the closest match, and it is Checkpoint-Based Governance before it is GOPEL. NIST's reasoning is that a model receives data and instruction as one stream. The separation between what an agent attempts and what it is permitted therefore has to be enforced by architecture around the agent. The framework page states the same division as a design rule: the enforcement layer records and enforces checkpoints mechanically, without reasoning. The paper assigns the two halves by name, with Checkpoint-Based Governance defining the constitutional requirements and GOPEL enforcing them. February stated the claim as a security property scoped to content. An agent that cannot evaluate content cannot be steered through content, so the governance layer "eliminates this attack surface by refusing to understand anything." The claim covers what the agent reads, not the code, keys, or hosts it runs on, which the specification covered with hash checks, separation of duties, and audit files held apart from the agent. The one reasoning function in the architecture, the Navigator that synthesizes platform outputs, was placed outside the agent on purpose. The agent only checks the structure of what comes back.

Tenet 8 follows from the same rule. NIST asks that deterministic capabilities remain the primary mechanism for access decisions and that probabilistic signals never decide. Checkpoint-Based Governance had already removed every probabilistic signal from the decision: no platform count, no confidence score, and no level of

agreement substitutes for human arbitration. The February code enforced determinism in the layer underneath with a static analyzer that scans every source file for inference calls, embedding lookups, classification operations, and content-dependent branching. The publication reported zero violations across all 14 source modules. A deterministic layer gives the same output for the same input, which is what makes it auditable. It is also what makes semantic prompt injection against it pointless, because it never interprets content. That property covers the content channel alone. A compromised key, a bad adapter, or a defect in the policy code is a different attack. The record meets those with code integrity by hash, separation of duties, and the fail-closed security layer, not with determinism.

Tenet 10 asks for evidence covering identity, delegated intent, and actions, kept where the agent cannot modify it. That requirement has the longest lineage in the record. Factics required in 2012 that every tactic leave evidence. Checkpoint-Based Governance required in September 2025 that every checkpoint decision generate a record. Its framework page now carries a section titled Why the Record Is the Proof, with the record immutable. The six record types of February do the work at the agent layer. The Request Record holds the task, its scope, and the human who initiated it. The Dispatch and Response Records hold what was sent and what came back, unedited. The Arbitration Record holds who decided and why, and the Decision Record links the final output to every record above it. The specification required those files to live apart from the agent's execution environment, chained by hash and signed, with corrections appended rather than overwritten. That is the "logging infrastructure independent of the agent" that NIST describes, written as a requirement eight months before the tenet and as a principle thirteen months before it.

Tenets 2 and 6 meet the record at the human. NIST wants the agent cryptographically linked to the person who tasked it and wants signed mandates from that person that define outcomes and constraints. The BEFORE phase of Checkpoint-Based Governance is that mandate in principle: the human governor sets scope, criteria, and constraints, and no AI execution begins without them. February bound the human to every arbitration and decision record with authenticated identity. It treated the configuration the agent runs on as the mandate: human-supplied constants, held as a governed artifact with its own version history, changed only through a checkpoint. What February did not do was sign that mandate at the moment of tasking as a statement a downstream service could verify. It also carried no scope-narrowing rule for sub-agents. The principle behind attenuation was in place as a rule of checkpoint authority. No AI can approve another AI's output, so no delegated decision can exceed what the named human signed. NIST's rule is about privilege, meaning the scope a sub-agent can hold, and the record carries no mechanism that narrows it. The two requirements sit side by side, and one does not enforce the other. The October 2026 Alpha code closed the signing gap. The Agent Model 1 Runner starts from a charter that a named human signs on a device the runner does not control. The Agent Model 2 Gateway runs only what a signed manifest names. In NIST's terms both are mandates created when the agent is tasked, and in both the agent holds only the public key.

Where the record goes further

Tenet 5 is titled Human-in-the-Loop Triggers, and its rationale names consent fatigue as the risk. Checkpoint-Based Governance named the same risk thirteen months earlier, in the line "Increasing approval rates may indicate automation bias," and the February specification gave it numbers. Approval above 95 percent or reversal below 2 percent for three consecutive cycles starts a mandatory audit. Persistent violation moves the work from the single-checkpoint model to the checkpoint-per-role model. The specification also refused to count the human review at the end of an automated line as a control. It called that review "an informal safety valve, not a formal governance control," and it placed the control in the trigger and the escalation path instead.

The difference in the title matters. Human in the loop, as the author has written since January 2026, describes presence. The framework page carries the point in its title, Presence Is Not Authority. Its rule is that the governor accepts, modifies, or rejects and leaves a record with identity, rationale, and timestamp. NIST's tenet sets the trigger. Checkpoint-Based Governance sets who answers once the trigger fires and what that ruling leaves behind. The February code built the gate as something that "will not advance without human approval." The March canonical description then added a mandatory rationale field to the Arbitration Record, so that a deliberate judgment and a credential submission would not look the same in the log.

The record also carries a limit the tenets do not state. When seven platforms reviewed the February code, none asked what happens if the Navigator lies. The human running the build asked. A structurally clean synthesis that misstates what the platforms said passes every deterministic check, because there is no injection pattern to find and no format error to flag. The February article named the defense as the human at the checkpoint, with every raw platform response preserved beside the synthesis so the human can compare them. That is Checkpoint-Based Governance doing what the enforcement layer cannot. Separation of policy from reasoning, which tenet 7 asks for, does not remove the problem. It makes the problem visible and routes it to a person.

Provider plurality is the other place the record reaches past the tenets. The ten say nothing about how many platforms an agent should consult. The February architecture made plurality structural: identical prompts to three or more independent platforms, with divergence treated as governance signal rather than noise. Checkpoint-Based Governance treats unanimous agreement as a reason to verify outside the AI ecosystem rather than a reason to proceed. Readers who commented on the NIST post raised the cross-organization case, where a consumer's agent asks a bank to act on a mandate the bank did not issue. They asked where the neutral anchor sits so that identities stay portable across vendor boundaries. The package answered the market-structure half of that question in February. It asked Congress to mandate API accessibility for platforms in federal procurement and high-consequence pipelines, on the principle that vehicles must meet safety standards to use public roads. It asked for investment in small platforms so that no single vendor's namespace becomes the only one that counts. The other half of the question, an anchor a receiving organization can verify, is the tenth tenet's work, and the record does not yet carry it.

The tenth tenet

Distinct agent identification is the tenet the record defines without yet issuing, and the reason is structural rather than an oversight. Checkpoint-Based Governance answers the identity question in its own terms. On the Agents page, "An agent is whatever acts to reach a final output," and "what separates them is where the named human stands." The record bound the human, through operator identity and non-repudiation, from February. By March it bound the platform as well, through the endpoint registry and signed inference receipts of the Confidential Processing Extension. The Verified AI Inference Standards Act proposed the attestation mandate behind them. The agent's own identity was checked inside the deployment, by hashing the running code against the approved version, and was never presented to anyone outside it. NIST's first tenet asks for the outward-facing version: an identity tied to workload, environment, and approved configuration, with a lifecycle of its own, that another system can verify.

The author's Verified AI paper of September 23, 2026, drew the line that governs that tenet. Identity verification is not behavioral verification. A credential says who stands behind an agent, not what the agent will do. That limit applies to every agent credential now being proposed, and it applies to whatever the agent layer, which the author calls HAIA Agents, issues next.

That work is where the tenth tenet meets the record. The Agent Model 1 Runner is started by a named human who signs its charter. Its receipts now carry a signed and witnessed deployment record naming the company, the deployer, their relationship, and the assignment. The Agent Model 2 Gateway enrolls named governors by passkey, with two people behind every enrollment after the first. What remains is to issue, at the moment an agent is created, a record that names what was created, why, and by whom. That record would state whether the agent runs as Responsible AI or under AI Governance, which is Checkpoint-Based Governance's own answer to who the agent is. It would travel with the agent to any system it touches. It is in design now, and the tenth tenet describes its shape.

What the dates carry

The core of this record is from 2025, and all of it was public. Checkpoint-Based Governance was on basilpuglisi.com from September 23, 2025. The specification, the package, and the code were on basilpuglisi.com, GitHub, and SSRN from February 2026. The frameworks were also exercised as live sessions across every major AI platform from 2025 on, under HAIA-RECCLIN and the parallel multi-AI review the author calls CAIPR. The work therefore sits in those platforms' own records as well as on the open web.

It also sat where researchers look. The NIST authors drew on public comments and stakeholder engagement, and that process is on the record. The concept paper behind it was published on February 5, 2026, two days after the agent architecture specification, with comments due April 2. Three days after the concept paper, on February 8, the package's technical appendix went up on SSRN. Its title is GOPEL Infrastructure Specification and HAIA-RECCLIN Operational Model for AI Provider Plurality Implementation. Its DOI was registered on March 24, nine days before comments closed. The scholarly indexes carry it under the keywords AI governance, audit trail, tamper-evident records, and accountability. The companion paper, Ethics for Oversight and Protection: The Constitutional Case for AI Governance Infrastructure, carries the same dates. The agent architecture specification itself was on Academia.edu as a working paper. Checkpoint-Based Governance had reached that channel the year before. The HAIA-RECCLIN papers posted to Academia.edu in 2025 include the Responsible AI Growth Edition. That edition carries it as the framework's governance model, with "human oversight checkpoints (BEFORE and AFTER)" and a "minimum three checkpoints per decision cycle." The first duty of a researcher is to ask whether similar work exists. The author holds that a search of the scholarly record by the subject's own terms would have returned this one. As of November 2026 the indexes record no citation of it.

NIST's summary of the more than 600 responses reads like the table above. Commenters told NIST that "the reasoning system cannot also be the authorizing system." They wanted human review to run "as a hard blocking state rather than a procedural step." They held deterministic enforcement essential and probabilistic mechanisms unfit to be the arbiter. They asked for tamper-evident evidence with cryptographic receipts. They also wrote that "traditional audit logs capture what happened but often fail to capture why an action occurred." Each of those positions was in the record before the window closed, the last of them as the rationale field the March canonical description made mandatory. Work of that volume, public for thirteen months and run on the platforms themselves, likely shaped any AI-assisted drafting, summarizing, or checking of agent governance material in 2026. This article does not say the record was among the sources anyone read. It says the record was in the channels a researcher is expected to search while the comments were open. It says the field arrived, inside eight weeks, at what the record had already published. The tenets codified the arrival. The author's position does not need the citation to stand, and whether anyone cited it does not change the dates. What a reader can do is lay the documents side by side and check them, which is what the table above is for.

Nor is this a claim that the record is finished. The two Alpha programs remain reference code without a professional security review. Reviews of that Alpha argued that the Responsible AI label on its outputs asserts more than independent scorers can support. They also noted that scoring platforms trained on similar data can share a blind spot. Those cautions stand in the record and stand here. They apply to the running desk as well, which is Responsible AI by design and has never been called governance. NIST's tenets, for their part, are a starting point by their authors' own description, with the NCCoE builds still to come.

The September 2025 question was who answers for an AI-assisted outcome and what record stands behind it. The February 2026 question was whether that answer could be built into an agent, and by October the agent was publishing. The October 2026 question is what agent identity infrastructure must look like. The three answers match on nine points, five in full and four in part. The tenth, an identity the agent carries to every system it meets, has its definition in the record and its credential still to issue.

Sources

Booth, H., Fisher, B., Galluzzo, R., & Clancy, T. (2026, October 27). *Design tenets for agent identity management systems*. Cybersecurity Insights, National Institute of Standards and Technology. https://www.nist.gov/blogs/cybersecurity-insights/design-tenets-agent-identity-management-systems

- Fisher, B., & Galluzzo, R. (2026, August 27). *Back to the future: Why agentic AI needs a strong identity foundation*. Cybersecurity Insights, National Institute of Standards and Technology. https://www.nist.gov/blogs/cybersecurity-insights/back-future-why-agentic-ai-needs-strong-identity-foundation
- Booth, H., Fisher, W., Galluzzo, R., & Roberts, J. (2026, February 5). *Accelerating the adoption of software and artificial intelligence agent identity and authorization* (Initial public draft; comments due April 2, 2026). NIST Computer Security Resource Center. https://csrc.nist.gov/pubs/other/2026/02/05/accelerating-the-adoption-of-software-and-ai-agent/ipd
- National Cybersecurity Center of Excellence. (2026). *Summary of comments* on the concept paper. National Institute of Standards and Technology. https://pages.nist.gov/nccoe-ai-identity/summary-of-comments.html
- National Cybersecurity Center of Excellence. (2026). *NCCoE Agentic SI Identity and Authorization Project Resource Hub*. National Institute of Standards and Technology. https://pages.nist.gov/nccoe-ai-identity/

Puglisi, B. C. (2026, September 3). *Factics: The method behind valuable content and trusted AI*. basilpuglisi.com. https://basilpuglisi.com/factics/

- First print introduction: Puglisi, B. C. (2012). *Digital Factics: Twitter*. MagCloud. https://www.magcloud.com/browse/issue/471388

Puglisi, B. C. (2025, September 23). *Checkpoint-Based Governance: An implementation framework for accountable human-AI collaboration (v2 drafting)*. basilpuglisi.com. https://basilpuglisi.com/checkpoint-based-governance-an-implementation-framework-for-accountable-human-ai-collaboration-v2-drafting/

- Latest update: Puglisi, B. C. (2026, September 11). *Checkpoint-Based Governance: Presence is not authority*. basilpuglisi.com. https://basilpuglisi.com/cbg/
- Full paper: *Checkpoint-Based Governance: A constitutional framework for human-AI collaboration* (October 2026), linked from the framework page.
- On Academia.edu in 2025: Puglisi, B. C. (2025). *HAIA-RECCLIN: The multi-AI governance framework for individuals, businesses & organizations. The Responsible AI Growth Edition*. Academia.edu. https://www.academia.edu/145013100/HAIA_RECCLIN_Multi_AI_Framework
- Asimov, I. (1942). Runaround. *Astounding Science Fiction*.
- Asimov, I. (1985). *Robots and Empire*. Doubleday.
- European Union. (2024). *Regulation (EU) 2024/1689*, Article 14. https://artificialintelligenceact.eu/article/14/
- Green, B. (2022). The flaws of policies requiring human oversight of government algorithms. *Computer Law & Security Review, 45*, 105681. https://doi.org/10.1016/j.clsr.2022.105681
- Lee, J. D., & See, K. A. (2004). Trust in automation: Designing for appropriate reliance. *Human Factors, 46*(1), 50 to 80. https://doi.org/10.1518/hfes.46.1.50_30392
- Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. *Human Factors, 39*(2), 230 to 253. https://doi.org/10.1518/001872097778543886

Puglisi, B. C. (2026, February 3). *HAIA-RECCLIN agent architecture specification*. basilpuglisi.com. https://basilpuglisi.com/haia-recclin-agent-architecture-specification/

- Latest update: Puglisi, B. C. (2026, October 13). *HAIA Agents: From the definition problem to running code*. basilpuglisi.com. https://basilpuglisi.com/haia-agents-running-code/
- Working paper: Puglisi, B. C. (2026). *HAIA-RECCLIN: Agent governance architecture for audit-grade multi-AI collaboration (EU regulatory compliance edition)*. Academia.edu.

https://www.academia.edu/164432153/HAIA_RECCLIN_Agent_Governance_Architecture_for_Audit_Grade_Multi_AI_Collaboration_EU_Regulatory_Compliance_Edition_

- Public draft, 2025: Puglisi, B. C. (2025, September 26). *The HAIA-RECCLIN model: A comprehensive framework for human-AI collaboration* (draft). basilpuglisi.com. https://basilpuglisi.com/the-haia-recclin-model-a-comprehensive-framework-for-human-ai-collaboration-draft/ Also on Academia.edu: https://www.academia.edu/144564806/The_Haia_Recclin_Model_public_draft
- Agents page: Puglisi, B. C. (2026, October 14). *HAIA Agents: Who acts, and who answers*. basilpuglisi.com. https://basilpuglisi.com/agents/
- Responsible AI desk, basilpuglisi.com: HAIA Agents author archive, first post September 29, 2026, author record "HAIA Agent Model 1 Responsible AI." https://basilpuglisi.com/author/haia-agents/
- Responsible AI desk, puglisiconsulting.com: HAIA Agents author archive, first posts October 3, 2026. https://puglisiconsulting.com/author/haia-agents/
- HAIA Agents. (2026, October 3). *Workflow and AI deployment: Agent identity before scale*. puglisiconsulting.com. https://puglisiconsulting.com/workflow-ai-deployment-agent-identity/
- Puglisi, B. C. (2026, September 11). *What is HAIA-RECCLIN? Reasoning and dispatch*. basilpuglisi.com. https://basilpuglisi.com/haia-recclin/
- Agranat Commission. (1974). *Israeli state inquiry into the Yom Kippur War*. State of Israel.
- Bai, Y., et al. (2022). *Constitutional AI: Harmlessness from AI feedback* (arXiv:2212.08073). https://arxiv.org/abs/2212.08073
- Clayton Antitrust Act, 15 U.S.C. §§ 12 to 27 (1914).
- European Union. (2024). *Regulation (EU) 2024/1689 (Artificial Intelligence Act)*, Articles 9 to 15 and 50. https://eur-lex.europa.eu/eli/reg/2024/1689/oj
- Henrich, J., Heine, S. J., & Norenzayan, A. (2010). The weirdest people in the world? *Behavioral and Brain Sciences, 33*(2 to 3), 61 to 83. https://doi.org/10.1017/S0140525X0999152X
- International Organization for Standardization. (2023). *ISO/IEC 42001:2023 AI management systems*. https://www.iso.org/standard/42001
- Janis, I. L. (1972). *Victims of groupthink*. Houghton Mifflin.
- Narechania, T. N., & Sitaraman, G. (2024). An antimonopoly approach to governing artificial intelligence. *Yale Law & Policy Review, 43*(1).
- National Institute of Standards and Technology. (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf
- Sherman Antitrust Act, 15 U.S.C. §§ 1 to 7 (1890).
- Zheng, L., et al. (2023). *Judging LLM-as-a-judge with MT-Bench and Chatbot Arena* (arXiv:2306.05685). https://arxiv.org/abs/2306.05685

Puglisi, B. C. (2026, February 8). *AI Provider Plurality: A Congressional Package, and the Verified AI Inference Standards Act*. basilpuglisi.com. https://basilpuglisi.com/ai-provider-plurality-a-congressional-package/

- Origin: Puglisi, B. C. (2025, October 30). *The case for AI provider plurality in evidence-based research*. basilpuglisi.com. https://basilpuglisi.com/the-case-for-ai-provider-plurality-in-evidence-based-research/
- Document 4, Technical Appendix: Puglisi, B. C. (2026, February 8). *GOPEL infrastructure specification and HAIA-RECCLIN operational model for AI provider plurality implementation*. SSRN. DOI registered March 24, 2026. https://doi.org/10.2139/ssrn.6195238
- Companion paper: Puglisi, B. C. (2026, February 8). *Ethics for oversight and protection: The constitutional case for AI governance infrastructure*. SSRN. DOI registered March 24, 2026. https://doi.org/10.2139/ssrn.6195278

- Index record: OpenAlex entry for DOI 10.2139/ssrn.6195238, source SSRN Electronic Journal, keywords AI governance, audit trail, tamper-evident records, accountability. https://api.openalex.org/works/doi:10.2139/ssrn.6195238
- Full package: https://github.com/basilpuglisi/Public-Policy
- Executive Order 14179, Removing Barriers to American Leadership in Artificial Intelligence (January 23, 2025).
- Executive Order 14365, Ensuring a National Policy Framework for Artificial Intelligence (December 11, 2025).
- Federal Aviation Act of 1958; Securities Exchange Act of 1934; Communications Act of 1934.

Puglisi, B. C. (2026, February 23). *GOPEL: The code behind the policy*. basilpuglisi.com. https://basilpuglisi.com/gopel-the-code-behind-the-policy/

- Latest update: Puglisi, B. C. (2026, March). *GOPEL canonical description* and the *Confidential Processing Extension*, in the HAIA repository. https://github.com/basilpuglisi/HAIA
- European Data Protection Supervisor. (2025, September 22). *TechDispatch #2/2025: Human oversight of automated decision-making*.
- Goddard, K., Roudsari, A., & Wyatt, J. C. (2011). Automation bias: A systematic review of frequency, effect mediators, and mitigators. *Journal of the American Medical Informatics Association, 19*(1), 121 to 127. https://doi.org/10.1136/amiajnl-2011-000089
- European Union. (2024). *Regulation (EU) 2024/1689 (Artificial Intelligence Act)*. https://eur-lex.europa.eu/eli/reg/2024/1689/oj
- National Institute of Standards and Technology. (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf

Puglisi, B. C. (2026, September 23). *Verified AI: What it means and where it came from*. basilpuglisi.com. https://basilpuglisi.com/verified-ai-history-meaning/

- Birkholz, H., Thaler, D., Richardson, M., Smith, N., & Pan, W. (2023). *Remote ATtestation procedureS (RATS) architecture* (RFC 9334). Internet Engineering Task Force. https://doi.org/10.17487/RFC9334
- Concordium. (2026, May 28). *The Concordium Agent Registry is live*. https://www.concordium.com/article/the-concordium-agent-registry-is-live-the-accountability-layer-erc-8004-doesnt-have
- Diamant, B., & Elias, T. (2025, July 21). *Introducing HUMAN Verified AI Agent*. HUMAN Security. https://www.humansecurity.com/learn/blog/human-verified-ai-agent-open-source/
- Advanced AI Society. (2026). *Open verification: The Proof-of-Control standard for agents* (Working draft). GitHub. https://github.com/LFDT-ProofOfControl/ov-poc-standard

Basil C. Puglisi, MPA

A Human-AI Collaboration

#AIassisted using the HAIA Ecosystem | CC BY-NC-SA 4.0

Free for personal, educational, and noncommercial research use with attribution. Commercial exploitation, paid productization, and enterprise commercialization require separate permission and licensing.

Suggested citation: Puglisi, B. C. (2026). NIST's Ten Agent Identity Tenets: What Checkpoint-Based Governance Already Carried. basilpuglisi.com.