Verified AI Inference Standards Act

A Congressional Framework for Trustworthy AI Data Processing

Basil C. Puglisi, MPA
Human-AI Collaboration Strategist
basilpuglisi.com

March 2026

Strategic Policy Proposal for Congressional Review
Senate HELP / Senate Commerce / House Energy and Commerce / House Science, Space, and Technology
AI Provider Plurality Congressional Package, Document 5 of 5

Executive Summary

Every time a hospital, insurer, school, financial institution, or federal agency sends sensitive data to an external AI platform, something happens that no current law governs at the technical level and no standardized mechanism can verify. The data arrive at the platform's inference endpoint. They are processed inside an environment the sending organization cannot inspect, cannot audit at the moment of inference, and cannot verify against any technical standard. Contractual promises govern what should happen. Nothing verifiable proves what did.

This window, between when data leave a trusted boundary and when a response returns, is the Invisible Moment. It is the most consequential unprotected interval in modern data governance, and it grows more consequential every day as AI inference becomes the central operation of healthcare, finance, education, defense, and government.

Some providers have begun offering confidential computing features for specific workloads, and cloud platforms broadly support attestation at the infrastructure layer. What the market still lacks is a standardized, customer-callable, per-transaction inference evidence artifact across mainstream managed LLM APIs. The infrastructure exists at the cloud layer. The standardized proof does not exist at the API layer where customers, governance systems, and regulators need it.

The Verified AI Inference Standards Act (VAISA) closes that gap. It does not require new science. The technical foundations exist in the RFC 9334 remote attestation architecture (an IETF Informational document providing the conceptual foundation for attestation workflows), NIST confidential computing standards, FIPS 203/204/205 post-quantum cryptography, and hardware trusted execution environments already deployed by every major cloud provider. What does not exist is the federal requirement that AI platforms expose those protections to the customers who send them sensitive data. VAISA creates that requirement.

This is the transponder problem applied to AI. Air traffic control works because aircraft carry transponders that broadcast identity, altitude, and position. If airlines refused to install transponders, the radar infrastructure would still function, but every aircraft would appear as an unidentified blip. The safety system would be structurally blind, not because the system failed, but because the participants refused to make themselves visible. AI governance faces the same structural problem today. The governance infrastructure exists. The hardware exists. The platforms won't install the transponders. First, existing federal law already mandates verifiable audit mechanisms for protected health data under 45 CFR 164.312(b), for student records under FERPA, for financial data under Gramm-Leach-Bliley and the SEC's cybersecurity disclosure rules, and for government data under the Protecting AI and Cloud Competition in Defense Act of 2025 (if enacted). The technical infrastructure those mandates require does not exist at the AI API layer. Second, every major standards body moving toward AI governance, NIST, the Cloud Security Alliance, the EU AI Act, and the UK Parliament, has identified data-in-use verification as a required capability. The standards are converging. The market is not following. Third, the future of safe multi-AI governance, the ability to detect bias, prevent system corruption, and maintain human authority over AI decisions across multiple platforms, requires that every AI API junction be verifiable. Without attestation at each inference boundary, governance frameworks are blind at the exact layer where risk manifests.

VAISA mandates a four-profile classification system for AI inference, directs NIST and HHS to publish a Verified Confidential Inference Standard within 12 months, aligns enforcement across OCR, FTC, ONC, and CISA, establishes a private right of action for individuals harmed by violations, provides safe harbors for compliant entities, and sets a federal floor with state authority preserved above it. The law is technology-neutral, phased over 30 months, and built on infrastructure that already exists. It converts the AI governance promise from trust us to prove it.

Methodology and Disclosure

This document was developed under the HAIA-RECCLIN (Human AI Assistant) framework, which assigns structured roles to AI platforms operating under human authority. The author served as Tier 0 arbiter with final editorial, analytical, and publication authority over all content. AI platforms operating in assigned RECCLIN roles (Researcher, Editor, Coder, Calculator, Liaison, Ideator, Navigator) contributed to drafting, technical verification, legal research, and structural review. The paper was adversarially reviewed under the HAIA-CAIPR (Cross AI Platform Review) protocol, where each platform reviews independently without access to the others' findings. Platforms participating across drafting and review rounds included Claude, ChatGPT, Gemini, Grok, DeepSeek, Kimi, Perplexity, Meta AI, Mistral, and CoPilot. Convergence and dissent were documented, and the Dissent Register at the end of this paper records every substantive disagreement raised during review with the rationale for how each was resolved. Under Checkpoint-Based Governance (CBG v4.7), every editorial decision, structural change, legal correction, and publication approval required explicit human authorization at defined checkpoints. No AI platform operated autonomously. No content was published without Tier 0 arbiter sign-off. The methodology, governance architecture, and review protocols are published at basilpuglisi.com and github.com/basilpuglisi/HAIA.

I. The Invisible Moment: A Gap Current Law Does Not Cover

Data protection law in the United States has always followed the data. HIPAA follows protected health information wherever it travels. FERPA follows student education records. Gramm-Leach-Bliley and SEC cybersecurity rules follow financial data. The Defense Production Act follows government procurement data. Each framework requires that whoever holds the data protect it.

None of them adequately governs what happens during active AI computation.

The model for data protection assumes two states: data at rest and data in transit. Encryption standards, access controls, audit logging, and retention limits all address data in one of those two states. When a healthcare organization sends a clinical note to an AI platform for summarization, the note is encrypted in transit under TLS. When it arrives at the platform's inference endpoint, it is decrypted into the platform's processing environment. That environment is opaque. The healthcare organization cannot inspect it. No technical standard requires the platform to prove what happened inside it. A business associate agreement allocates liability after the fact. It does not verify the processing conditions before or during inference.

This is the Invisible Moment. It is not hypothetical. It occurs billions of times daily across every regulated sector that uses an AI API. The same structural gap appears in financial services when banks and investment firms use AI platforms to process account data, transaction records, and credit information. It appears in education when student learning data, disciplinary records, and assessment results are processed by AI tutoring and analytics platforms. It appears in every federal agency using AI platforms to process data that procurement law already defines as government property. The Invisible Moment is a cross-sector infrastructure failure, not a healthcare-specific edge case.

The HIPAA Security Rule at 45 CFR 164.312(b) requires covered entities and their business associates to implement hardware and software activity controls that record and examine activity in systems that contain electronic protected health information. That requirement applies to AI inference systems that handle ePHI. No major AI platform currently provides a standard, customer-callable attestation interface that satisfies that requirement at the inference boundary. The audit control the law requires does not technically exist at the layer where the law demands it.

FERPA requires institutional protection of student records processed by third-party systems. The Protecting AI and Cloud Competition in Defense Act of 2025 (H.R.3434/S.1775), if enacted, would require that the government maintain exclusive rights to access and use of all government data in covered DoD contracts with foundation model providers. Gramm-Leach-Bliley requires financial institutions to ensure the security and confidentiality of customer records processed by service providers. Each of these statutes, whether enacted or pending, implies the same technical capability at the AI API layer. None of them can be satisfied by a platform that processes regulated data in an opaque environment and offers only contractual assurances in return.

VAISA does not create a new security obligation. It specifies the technical evidence standard that makes existing security obligations enforceable at the AI inference boundary. The underlying duty to protect regulated data during processing already exists in law. What does not exist is the technical proof that the duty was met during the most vulnerable moment of processing. VAISA creates that evidence requirement.

II. Why Contracts Are Not Enough

This argument requires direct engagement because it is the objection compliance officers, lawyers, and platform vendors raise first and most confidently. Business associate agreements, enterprise privacy commitments, and zero-retention promises are meaningful instruments. They are not the same instrument as verifiable confidential inference, and the distinction is not semantic.

A contract allocates responsibility and liability. It governs what a party is obligated to do and what consequences follow if they fail. A BAA between a healthcare organization and an AI platform establishes that the platform must protect ePHI, must not use it for model training without authorization, must notify of breaches, and must return or destroy data at contract termination. Those are enforceable obligations.

What a BAA cannot do is prove that a specific inference request was processed only inside an intended trusted execution environment. It cannot prove that the decrypted clinical note was visible only to approved code in an approved operating state. It cannot prove that no unauthorized process observed the plaintext data during computation. Those are not matters of contractual obligation. They are matters of technical fact, and the only way to establish technical fact is through technical evidence.

RFC 9334, the IETF's Remote ATtestation procedureS architecture, provides the precise framework for this distinction. Under that architecture, an attester provides evidence about its operating state, a verifier evaluates that evidence against known good values, and a relying party, in this context the healthcare organization or other regulated entity, makes a decision about whether to release data based on verified evidence rather than contractual promise. The attester cannot fabricate a valid attestation quote without access to the hardware private key held inside the TEE. The evidence is cryptographically bound to the physical hardware state at the moment of attestation. No contract produces that quality of assurance.

The practical stakes are concrete. OpenAI, Anthropic, and Google all publish meaningful privacy and security commitments. OpenAI offers healthcare BAAs for eligible customers. Those are genuine market signals and they matter. They are not the same as a universal, customer-verifiable attestation mechanism on mainstream managed LLM APIs. A platform could honor every term of its BAA and still experience an insider threat, a misconfiguration, or a compromised dependency that exposes ePHI during inference. The BAA governs the aftermath. Attestation governs the conditions. VAISA requires the latter.

III. The Standards Ecosystem Is Converging Without a Federal Mandate to Follow

The direction of the standards landscape is unambiguous. Every major governance body working on AI data protection has identified the same technical gap and each has moved toward the same class of solution. What is missing is the federal requirement that turns convergence into adoption.

NIST's February 2026 concept paper, "Accelerating the Adoption of Software and Artificial Intelligence Agent Identity and Authorization," identifies identity, credentialing, and authorization as core requirements for AI agents operating across APIs and data systems. The initiative explicitly anticipates mechanisms to authenticate AI agents, define and limit their permissions, and ensure they can be properly scoped, monitored, and governed within systems. Attestation is the technical instrument those mechanisms require.

The Cloud Security Alliance AI Controls Matrix, released July 2025, covers 243 controls across 18 security domains including data security, privacy, model robustness, and AI lifecycle management. It is positioned as the foundation for an AI STAR certification and attestation pathway program. That pathway is described as coming in the near future. Federal legislation accelerates near future to now.

The EU AI Act's Annex VII requires that notified bodies evaluating high-risk AI systems receive full access to training, validation, and testing data through API or other relevant technical means enabling remote access. That provision establishes a regulatory precedent for mandatory technical access to AI system evidence, though it addresses conformity assessment of training data rather than customer-callable runtime inference attestation. The principle of technically accessible evidence for AI governance extends naturally to the inference boundary. U.S. enterprises operating globally will face these requirements regardless of domestic law.

The UK House of Lords Communications and Digital Committee, reporting in March 2026, called for a statutory obligation requiring AI developers to adopt open technical standards for data provenance and labeling, stating that without those foundations there is no reliable way to establish whether protected works or data have been used. The technical standards that committee described extend directly to inference-time data handling.

The FTC Health Breach Notification Rule, which extends coverage beyond HIPAA to non-HIPAA businesses with unsecured individually identifiable health information, signals that Congress and regulators have consistently chosen to close gaps between regulated and unregulated handlers of the same class of data. VAISA applies that same logic to the inference boundary.

The pattern is consistent. The direction is established. The market is not moving fast enough on its own because no commercial incentive drives individual platforms to invest in customer-verifiable attestation when contracts and privacy policies already satisfy current regulatory requirements. Federal legislation exists precisely to close that gap.

IV. Multi-AI Governance, Bias, Control, and the Case for Open Infrastructure

The two-layer argument above is sufficient to justify VAISA on compliance grounds alone. The third layer explains why the standard must be open infrastructure rather than a proprietary feature, and why acting now rather than waiting for another mandate is the correct strategic choice.

Enterprise AI has moved from single-model assistance to multi-platform orchestration. A governance decision in healthcare, finance, law, or government may now pass through several AI platforms in sequence, with each platform producing outputs that feed the next. Bias can enter at any junction. A corrupted or compromised system can influence downstream outputs without detection. Human arbiters reviewing the final output cannot see what happened at each intermediate inference step unless attestation exists at every boundary.

Provider Plurality, the governance principle that no single AI actor should control the decisions that matter most, has a technical corollary at the inference layer. If multiple AI systems process sensitive data across platforms and no attestation standard exists, a single compromised or biased platform can influence an entire governance chain invisibly. The risk that market-level plurality addresses, namely concentration of AI capability and the corruption or capture of that capability, manifests at the inference boundary. Attestation is what makes plurality operationally meaningful rather than merely contractual.

Governance frameworks that require human checkpoint decisions to be based on verifiable evidence of what AI systems did with the data they processed cannot function without attestation. Without it, every checkpoint rests on the AI provider's self-report. That is not governance. It is trust with a governance label. The same structural requirement applies to any future framework that takes human oversight of AI seriously, regardless of which framework architecture an organization adopts.

The superintelligence risk argument extends this further. The scenarios that most concern AI safety researchers, systems that pursue goals misaligned with human values, systems that manipulate their environment invisibly, systems that deceive operators about their internal states, all depend on a common enabling condition: the system can act in ways that are invisible to oversight at the moment of action. Attestation at the inference boundary does not eliminate that risk. It is a necessary layer of the infrastructure that makes oversight technically possible rather than merely aspirational.

This is the reason VAISA must require an open standard rather than permitting proprietary attestation implementations. A proprietary attestation system allows a single platform to offer compliance verification as a competitive moat. Enterprises dependent on that platform for both AI capability and compliance verification cannot switch providers without losing their audit trail. That reintroduces the concentration risk that Provider Plurality exists to prevent, now at the compliance infrastructure layer rather than the capability layer. An open standard, callable through any compliant provider's API and verified against a public specification, serves every governance framework, every regulated sector, and every size of enterprise equally.

The statutory model is HL7 FHIR under the 21st Century Cures Act. Before that mandate, health data interoperability was technically possible. The market was not producing it voluntarily because proprietary data formats created lock-in that benefited large vendors. The mandate created a clear technical requirement, a standards body responsible for the specification, a phased compliance timeline, and an enforcement mechanism. The market followed. AI inference attestation is the same problem one layer deeper in the same technology stack, and the Cures Act precedent is understood by the same HHS and congressional staff who would receive VAISA.

V. The Technical Solution: Four-Profile Classification

VAISA's operational core is a four-profile classification system that governs every AI inference dispatch involving regulated data. The profiles are operational states with specific technical requirements and specific permissions, not aspirational categories.

Profile 1: Attested Confidential Inference

The AI platform operates within a hardware-enforced Trusted Execution Environment. Before any sensitive data leave the sending organization's boundary, the orchestrating system performs a deterministic check of a hardware-signed attestation quote from the platform. Full regulated data may be processed under this profile. The signed attestation quote serves as the audit evidence that 45 CFR 164.312(b) and equivalent standards require.

A compliant attestation quote must include five mandatory technical elements. First, a hardware vendor signature from a recognized TEE architecture, currently AMD SEV-SNP, Intel TDX, ARM CCA, or NVIDIA H100 Confidential Computing, with the list updatable by NIST as hardware evolves. Second, an enclave measurement consisting of a cryptographic hash of the approved code image running inside the enclave, binding the attestation to specific approved software. Third, a freshness nonce generated by the relying party and included in the quote, preventing replay of prior valid attestations. Fourth, a TCB version establishing the security

patch level of the trusted computing base, allowing the relying party to verify the enclave is not running on downgraded firmware with known vulnerabilities. Fifth, a signed processing receipt that binds the specific inference transaction to the attested environment through the nonce and an input hash, creating a per-transaction audit artifact rather than a generic platform certification.

Profile 2: Minimized Processing

Attestation is unavailable or has not been verified for the specific data class at issue. The orchestrating system applies deterministic tokenization to structured regulated identifiers before dispatch. The platform processes a redacted representation. The tokenization map remains in the sending organization's custody. The platform never holds the key to reverse the tokenization. Approved minimization must cover all structured PHI elements under HIPAA Safe Harbor, all FERPA-defined personally identifiable information fields, and equivalent identifier classes for financial and government data. The tokenization ruleset must be versioned, hashed, and subject to annual third-party audit. This profile permits processing of minimized representations, not raw regulated data.

Profile 3: Human-Gated Emergency Processing

Neither attestation nor sufficient minimization is available for the specific data class at issue, and a documented operational necessity requires processing. A mandatory human pause gate activates. A named human arbiter must review and explicitly approve the inference dispatch. The approval is logged with the arbiter's identity, the timestamp, the data class involved, the volume of records, and the documented justification. This profile is available only under defined emergency conditions, is time-limited to 72 hours per authorization, and requires post-use review within 14 days. It is not a routine operational profile and repeated use without migration toward Profile 1 or 2 constitutes non-compliance.

Profile 0: Prohibited

No attestation, no approved minimization, no human gate approval. Raw regulated data do not leave the organization's boundary for inference at an external platform. This is the default state. Every inference dispatch involving regulated data begins at Profile 0 and must affirmatively qualify for a higher profile before data are released.

Scope limitation: Attestation under VAISA verifies that approved code ran in an approved hardware environment. It does not verify that the approved code behaves as the relying party expects. A platform operator who controls the code running inside the TEE could deploy software that passes measurement checks but acts against the data subject's interests, including unauthorized logging, retention beyond declared periods, or side-channel extraction within the enclave. The defense against this scenario is independent code audit combined with attestation, not attestation alone. VAISA addresses the environment verification layer. Code behavior assurance requires complementary audit mechanisms that VAISA does not prescribe but that the NIST standard should reference as a recommended practice for the highest-assurance deployments.

VI. Post-Quantum Integrity

The audit trail that attestation creates must remain trustworthy across the retention periods that regulated data require. Health records carry multi-decade retention obligations. Government records are longer. Financial records frequently exceed ten years. Current digital signature standards including Ed25519 and RSA are vulnerable to Shor's algorithm running on a cryptographically relevant quantum computer. An adversary who harvests signed audit logs

today and forges or decrypts them after quantum capability arrives can retroactively alter the governance record at the exact moment those records are most likely to be needed for enforcement, litigation, or historical audit.

VAISA requires post-quantum cryptographic readiness for all attestation signatures, signed inference receipts, and audit log integrity chains. NIST has finalized FIPS 203, FIPS 204, and FIPS 205 as the initial post-quantum cryptography standards. The transition profile for VAISA compliance requires hybrid signatures on all new deployments after Phase One, combining a classical signature with an ML-DSA signature in a non-separable composite. Existing deployments complete migration by month 36. External anchoring of audit log hash chains to a quantum-resistant public timestamping authority prevents retroactive history rewriting even if signing keys are eventually compromised.

VII. Legislative Framework

VAISA requires six statutory provisions to function as a coherent operating system.

Provision One: Scope and Definitions

VAISA applies to any Covered AI Inference Service defined as any external system that receives protected health information, electronic protected health information, student education records, government procurement data covered by applicable federal law, customer financial records covered by Gramm-Leach-Bliley, individually identifiable consumer health data, or other statutorily defined sensitive data classes, for the purpose of generating, classifying, summarizing, transcribing, extracting, or otherwise transforming the data using machine learning or foundation model inference.

The law applies to covered entities and business associates under HIPAA, to educational institutions under FERPA, to financial institutions under Gramm-Leach-Bliley, to federal contractors under applicable procurement law, and to non-HIPAA consumer health businesses through parallel FTC authority. It applies to platforms wishing to serve regulated U.S. data regardless of the platform's country of domicile.

Provision Two: Technical Standards

Congress directs NIST and HHS to publish a Verified Confidential Inference Standard within 12 months of enactment. That standard defines the five mandatory technical elements for Profile 1 attestation as described in Section V, the approved tokenization methods and audit requirements for Profile 2 minimization, the logging and review requirements for Profile 3 human-gated processing, the post-quantum migration profile for all signature and key management functions, the evidence fields that orchestrating organizations must retain to demonstrate compliance, and the update mechanism by which NIST adds approved TEE architectures as hardware evolves.

The standard must also define three additional technical deliverables. First, the signed inference receipt specification, including transaction binding formats, per-request evidence schemas, and the cryptographic relationship between the attestation quote and the receipt, so that platforms build to a published standard rather than an undefined legislative requirement. Second, cross-vendor attestation format harmonization across Intel TDX, AMD SEV-SNP, ARM CCA, and NVIDIA GPU attestation, including the abstraction layers required for a unified API that governance systems can consume without vendor-specific integration. Third, coordinated CPU-GPU attestation chain requirements for distributed inference pipelines, addressing the binding between CPU TEE attestation (where TLS termination and prompt preprocessing occur) and

GPU TEE attestation (where transformer computation runs), including the data path integrity of bounce buffers and interconnect transfers that may traverse non-TEE memory.

The standard must include proportional compliance pathways for smaller AI platforms, open-source model providers, and research institutions. These pathways should include a reduced-burden attestation profile for platforms processing lower-risk regulated data, graduated requirements tied to revenue or processing volume thresholds, and reference architectures that allow smaller providers to achieve Profile 2 compliance through tokenization without requiring the hardware investment that Profile 1 demands. Without proportional pathways, VAISA risks concentrating the regulated market among large cloud providers, contradicting the AI Provider Plurality principle that the Congressional package is built on.

Provision Three: Multi-Agency Enforcement and Penalty Structure

HHS OCR enforces against covered entities and business associates under HIPAA. After the Phase Three deadline, processing raw ePHI through an unverified external AI inference endpoint without an approved Profile exception constitutes a Security Rule violation subject to civil monetary penalties.

Penalty tiers mirror the HIPAA civil monetary penalty structure, with per-violation minimums subject to annual inflation adjustment under the Federal Civil Penalties Inflation Adjustment Act. Tier One applies where the covered entity did not know and with reasonable diligence could not have known of the violation, with an annual cap of $25,000 per identical violation category. Tier Two applies where the violation is due to reasonable cause and not willful neglect, with an annual cap of $100,000. Tier Three applies where the violation is due to willful neglect that is corrected within 30 days, with an annual cap of $250,000. Tier Four applies where the violation is due to willful neglect that is not corrected, with an annual cap of $1,500,000. Per-violation daily minimums follow the current inflation-adjusted schedule published by HHS OCR.

A separate penalty category applies to AI platforms that market services as HIPAA-compliant, FERPA-compliant, or equivalent without providing a compliant attestation interface: $50,000 per violation per day that the misrepresentation continues, treated as a Tier Four violation and enforced jointly by OCR and the FTC.

The FTC enforces parallel duties for non-HIPAA consumer health businesses. False claims about confidential processing constitute unfair or deceptive conduct under Section 5 of the FTC Act. The FTC Health Breach Notification Rule is updated to require notification when a breach occurs at an AI inference endpoint processing covered health data, regardless of HIPAA applicability, aligning non-HIPAA enforcement with the same technical standard.

ONC integrates attestation disclosure and evidence requirements into certified health IT where AI functions appear in regulated clinical workflows, building on the algorithm transparency framework in HTI-1. CISA and HHS 405(d) publish operational adoption guidance. OMB, GSA, and FedRAMP mirror the requirement in federal procurement.

Criminal liability for knowing falsification of attestation evidence relating to healthcare data attaches under 18 U.S.C. 1035 (false statements relating to health care matters), with penalties up to five years imprisonment. For non-healthcare regulated data classes, VAISA creates a separate criminal provision for knowing falsification of attestation records, with penalties calibrated to the sensitivity of the data class and the scale of the falsification. This bill-created criminal language is necessary because no existing generally applicable federal statute covers knowing fabrication of technical compliance evidence across all regulated sectors VAISA addresses.

Provision Four: Private Right of Action

Any individual whose regulated sensitive data are processed by a Covered AI Inference Service in violation of VAISA may bring a civil action in federal district court. This provision fills the enforcement gap that HIPAA has never closed, namely the absence of a private cause of action for individuals harmed by violations that regulators may not prioritize or pursue.

Available remedies include actual damages suffered as a result of the violation, statutory damages of not less than $1,000 and not more than $10,000 per violation where actual damages are difficult to establish, injunctive and declaratory relief, and reasonable attorney's fees and costs.

Class actions are permitted where the same violation affected a defined class of individuals, subject to standard Rule 23 requirements. A two-year statute of limitations runs from the date the individual discovered or reasonably should have discovered the violation.

The private right of action does not displace regulatory enforcement. It supplements it, giving patients, students, account holders, and citizens a direct stake in the enforcement of the standard that protects their data.

Provision Five: Safe Harbor for Compliant Entities

No covered entity, business associate, educational institution, financial institution, or federal contractor faces civil monetary penalties or private right of action liability under VAISA where that entity demonstrates all of the following. First, it implemented a compliant profile classification system as defined in Section V. Second, it retained signed attestation evidence, tokenization audit records, or human gate approvals as applicable to each inference dispatch involving regulated data. Third, it promptly disclosed any attestation failure, TEE downgrade, or profile misclassification to affected individuals and to the relevant regulatory body within 30 days of discovery of the failure. The 30-day window runs from the date the entity first identifies the failure event, not from confirmation of its full scope; entities may supplement initial disclosure with updated scope assessments without losing safe harbor protection. Fourth, it maintained a current migration plan toward Profile 1 compliance for all regulated data workflows.

The safe harbor is not available where the entity knowingly falsified attestation records, knowingly processed regulated data at Profile 0 without a documented emergency exception, or delayed disclosure of a known violation beyond the 30-day window.

The safe harbor creates the market incentive that enforcement alone does not provide. Entities that invest in compliance gain concrete liability protection. The asymmetry between compliant and non-compliant entities drives adoption without requiring regulators to pursue every violation.

Provision Six: Federal Floor with State Authority Preserved

VAISA establishes a federal floor. It does not preempt stronger state protections. States that have enacted AI transparency, data provenance, confidential computing, or health privacy requirements at or above the VAISA standard are not affected. States may strengthen requirements for their own residents. The federal standard removes the compliance chaos of 50 different minimum requirements while preserving state authority to lead where Congress has not yet acted. This is the HIPAA model: federal minimum, state ceiling above it, and no preemption of state laws that offer greater protections.

VIII. Implementation Roadmap

Phase One: Standards Development (Months 1 through 12)

Congress directs NIST and HHS to publish the Verified Confidential Inference Standard, define sensitive data classes, specify evidence fields and approved attestation formats, identify compliant TEE architectures, and publish the post-quantum migration profile. During this phase, covered organizations must inventory which AI workflows touch regulated data, what contracts govern those flows, what audit logs currently exist, and whether any current endpoint offers verifiable confidential processing. That mapping exercise produces immediate governance value independent of the technical standards completion.

Phase Two: Infrastructure and New Deployment Requirements (Months 12 through 30)

New deployments involving raw regulated data must use a compliant profile before launch. Existing systems continue under transition conditions provided they document risk, limit data scope, publish a migration plan, and do not expand regulated data processing scope without profile compliance. Cloud providers expose TEE attestation APIs for AI workloads. AI platforms implement attestation endpoints and signed inference receipts. Covered organizations update procurement requirements and BAAs to include attestation verification as a compliance condition. Post-quantum hybrid signatures are required for all new attestation signing operations beginning at month 12.

Phase Three: Full Enforcement (Month 30 and Beyond)

Raw regulated data under VAISA's defined sensitive classes do not flow to unverified external AI inference endpoints in ordinary operations. Profile 3 emergency exceptions remain available but are rare, time-limited, logged, and subject to post-use review. Existing deployments without post-quantum migration complete that migration by month 36. OCR, FTC, and other enforcement bodies audit attestation evidence as a standard compliance verification element. Private right of action claims may be filed beginning at Phase Two (after the NIST and HHS standard is published and covered entities have had a defined implementation period), with the safe harbor available from that same date. Tying private suit eligibility to standards publication rather than enactment ensures that entities have a published technical standard to comply with before facing individual liability.

IX. Who Does What

Congress sets the rule, scope, deadlines, liability structure, and funding.

HHS OCR writes the healthcare regulatory implementation, updates BAA expectations to include attestation verification as a standard term, and enforces the penalty structure.

NIST issues the technical profiles, approved TEE architecture list, post-quantum migration profile, and update mechanism.

ONC integrates the requirement into certified health IT, surfacing profile status as a visible indicator in clinical workflows so clinicians and administrators see a binary status for each AI function.

FTC enforces for consumer health and non-HIPAA businesses, and updates the Health Breach Notification Rule to cover AI inference endpoint breaches.

CISA and HHS 405(d) publish operational adoption guidance for smaller and resource-constrained covered entities.

OMB, GSA, and FedRAMP align federal procurement, making the government an anchor customer for the compliant market rather than a purchaser below the statutory floor.

Large cloud and AI providers expose attestation APIs, attested key release mechanisms, and signed inference receipts. Hospitals, payers, digital health companies, educational institutions, financial institutions, and federal contractors update procurement, architecture, and compliance documentation. EHR and health IT developers surface processing status metadata inside their tools. Auditors and regulators verify the evidence trail without needing to inspect the protected content itself. That last point is critical. The enforcement model is evidence-based, not content-based.

X. Funding and Technical Assistance

Rural hospitals, community health centers, public hospitals, smaller digital health firms, smaller educational institutions, and community financial institutions require grants, technical assistance, and reference architectures to implement attestation verification in their orchestration layers. Without this provision, VAISA becomes a large-institution advantage rather than a national data protection standard.

Congress should appropriate funds through three channels. HHS 405(d) provides the existing operational support channel for healthcare entities. The Department of Education's Office of Educational Technology provides the parallel channel for educational institutions. A new NIST Small Entity AI Attestation Assistance Program provides reference architectures and implementation guidance for financial and other regulated entities outside the healthcare and education channels.

The FHIR interoperability mandate reached broad adoption partly because ONC funded implementation assistance for smaller providers who could not absorb the cost of custom development. The same investment logic applies here. Compliance adoption across the full regulated economy, not only among large institutions with dedicated compliance engineering teams, is both the policy goal and the enforcement precondition.

XI. Responding to the Standard Objections

VAISA will stifle AI innovation in regulated sectors.

VAISA creates two markets. Attested confidential AI for regulated data, and standard AI for non-regulated use cases. Innovation continues in both. The regulation accelerates confidential computing adoption that hardware already supports but that the market has not produced voluntarily. NVIDIA H100 Confidential Computing carries modest but non-zero performance overhead that varies by workload and deployment configuration. The safe harbor provision makes early compliance financially advantageous. The private right of action makes non-compliance financially risky. That combination drives adoption more effectively than enforcement alone.

Business associate agreements already cover this ground.

As Section II establishes in detail, a BAA allocates contractual liability. It does not prove inference conditions. Remote attestation under RFC 9334 produces cryptographically bound evidence of operating state. No contract produces that quality of assurance. The $12.5M HIPAA settlement recorded in 2025 against a major health system demonstrates that BAAs do not prevent breaches. VAISA governs the conditions under which data are processed. BAAs govern the consequences when those conditions fail.

This is too complex for small providers.

VAISA compliance for small organizations is orchestrator-mediated. Small practices use EHR systems. Small educational institutions use learning management platforms. Small financial institutions use core banking systems. Those platform vendors implement attestation verification transparently. The provider or administrator sees a verified or requires human approval status for each AI function. No technical expertise is required at the point of service. The funding provision ensures that even the infrastructure-layer complexity does not fall solely on under-resourced entities.

International platforms will not comply.

VAISA applies to the processing of U.S. regulated data, not to platform domicile. Platforms wishing to serve U.S. healthcare, education, financial services, or government must comply or restrict to Profile 2 minimization with reduced functionality. This is the established extraterritorial model that GDPR uses for EU personal data and that U.S. export controls use for sensitive technology. It is legally established territory.

XI-A. What VAISA Changes for Each Stakeholder

For Congress

The AI Provider Plurality Congressional Package proposes three actions: fund GOPEL as national infrastructure, mandate API accessibility, and invest in small AI platforms. VAISA provides the specific technical content of the API accessibility mandate. Document 3 of the package proposes the principle. VAISA specifies what the API must expose, why, and under what enforcement structure. Phase 0 of the Provider Plurality roadmap requires no new appropriation. VAISA Phase One operates within existing NIST and HHS budget authority for standard development. Phase Two and Three build on two to four years of federal operational evidence that the infrastructure works before the legislative mandate takes effect.

For Enterprise CISOs and Compliance Officers

Every governed workflow that sends sensitive data to a commercial LLM API today produces an audit trail gap. Governance infrastructure like the GOPEL Confidential Processing Extension makes that gap visible by classifying every dispatch and reporting its privacy status. Without VAISA, 100% of dispatches to commercial LLM platforms fall into Profile 0 (unverified) regardless of the platform's actual security posture, because no platform returns attestation evidence through its standard API.

VAISA changes the enterprise procurement conversation from a qualitative question to a binary one. Instead of asking "is our data safe?" and accepting a policy document, CISOs can ask a single question: "does your API return attestation evidence that our governance system can verify?" The answer is binary and auditable. Platforms that return attestation evidence qualify for Profile 1 (verified). Platforms that don't qualify for Profile 0 (prohibited for raw regulated data under VAISA). The procurement lever is the enforcement mechanism. Enterprise customers spending millions on AI platform contracts can include attestation API requirements in their procurement specifications today, without waiting for legislation.

For AI Platform Providers

The engineering requirement is significant but bounded. Cloud infrastructure already supports TEE attestation through AWS Nitro Enclaves, Azure Confidential Computing, Google Confidential VMs, and NVIDIA H100 Confidential Computing. Exposing attestation evidence through the API requires integration with the cloud provider's attestation service, and building the signed inference receipt requires additional platform-level engineering for per-transaction binding. The complexity increases for distributed inference across multi-tenant GPU clusters with CPU-GPU mixed attestation chains. The NIST standard in Phase One defines these specifications so platforms build to a published target rather than an undefined requirement.

The competitive advantage is concrete. The first major LLM provider to offer attestation-verified inference gains a significant advantage for the next wave of high-value regulated enterprise contracts: financial services, healthcare, government, legal, and insurance. Attestation is a threshold requirement for competing in these markets, not a sole differentiator, but every competitor without attestation falls into Profile 0 in every governance audit trail produced by every customer operating under VAISA, and that threshold disadvantage compounds as adoption grows.

The alternative is worse. Without voluntary adoption, the regulatory mandate arrives on the timeline VAISA specifies. Platforms that build attestation capability now shape the standard. Platforms that wait have the standard imposed on them. The safe harbor provision makes early compliance financially advantageous. The private right of action makes non-compliance financially risky. That combination drives adoption more effectively than enforcement alone.

XI-B. The Infrastructure Precedent

Every critical infrastructure domain follows the same pattern. The governance layer requires visibility into participant operations. The participants resist the cost and the exposure. Legislation mandates it. The market adjusts and the infrastructure becomes standard.

Aviation transponders were voluntary until they weren't. The FAA mandated Mode C transponders in controlled airspace (14 CFR 91.215) because air traffic control cannot separate aircraft it cannot see. ADS-B Out became mandatory in 2020. Airlines resisted the cost. The mandate held. The airspace is safer because the governance layer has visibility.

Financial institutions resisted automated reporting requirements under the Bank Secrecy Act. Suspicious Activity Reports are now mandatory. FinCEN requires visibility into transactions the institutions would prefer to keep opaque. The governance layer prevailed because the alternative was an unauditable financial system.

CALEA (Communications Assistance for Law Enforcement Act) requires telecommunications carriers to build interception capability into their networks. The carriers resisted. The mandate held because the governance layer requires structural access, not voluntary cooperation.

NERC CIP standards require utilities to maintain auditable records of access to critical cyber assets. Utilities resisted the compliance burden. The mandate held because the electrical grid is critical infrastructure and critical infrastructure requires governance visibility.

AI platforms processing regulated sensitive data are critical infrastructure participants. The governance layer requires visibility into processing conditions. The platforms resist because competitive opacity and liability asymmetry reward the status quo. The precedent across every other critical infrastructure domain is clear: the mandate follows, the market adjusts, and the infrastructure becomes safer.

XII. Relationship to Existing and Proposed Legislation

VAISA is complementary to the Health Information Privacy Reform Act and related HIPAA modernization proposals. It provides the technical implementation layer that privacy reform proposes without specifying. It is distinct from the AI LEAD Act, which addresses product liability for AI outputs. VAISA addresses processing conditions at the inference layer, not downstream product performance.

The December 2025 Executive Order on AI requires the development of a uniform national AI framework. VAISA provides the data infrastructure layer that a uniform framework requires to be technically enforceable rather than merely aspirational. It does so without restricting AI capability or creating sector-specific barriers to adoption, which aligns with the administration's stated innovation priorities.

The Protecting AI and Cloud Competition in Defense Act of 2025 (H.R.3434/S.1775), if enacted, would establish data rights for government procurement. VAISA provides the technical mechanism to enforce those rights at the inference boundary, converting contractual data rights into verifiable technical protections.

The 21st Century Cures Act and its FHIR mandate is the correct statutory precedent for VAISA in both structure and political economy. Before that mandate, health data interoperability was technically possible and the market was not producing it voluntarily. The mandate created a clear technical requirement, a standards body responsible for specification, a phased compliance timeline, an enforcement mechanism, and implementation funding for smaller entities. The market followed. The Cures Act precedent is understood by the same HHS, ONC, and congressional staff who would receive and act on VAISA.

XIII. Conclusion: From Trust Us to Prove It

The era of governing AI data handling through contractual promises has reached its limit. The promises are real. They are not sufficient. What regulated sectors need, what existing law already requires in principle, and what the standards ecosystem is building toward in practice, is customer-verifiable technical evidence that AI platforms protect sensitive data at the moment it is most vulnerable, during active computation.

VAISA converts that requirement from aspiration to infrastructure. It does not ask platforms to do something technically impossible. It asks them to expose what the hardware already supports, what major cloud providers already offer as optional services, and what governance frameworks already depend on. It gives enterprises a common procurement language. It gives regulators an auditable evidence trail. It gives patients, students, account holders, and citizens both verifiable assurance and a direct legal remedy when that assurance fails.

The standards are ready. The hardware is ready. The regulatory imperative is present in existing law today. Congress provides the mandate that converts readiness into adoption and converts compliance promises into compliance proof.

Appendix A: Technical Reference Standards

RFC 9334, Remote ATtestation procedureS (RATS) Architecture (Informational; conceptual foundation for attestation workflows) NIST FIPS 203, Module-Lattice-Based Key-Encapsulation Mechanism Standard (Final, August 2024) NIST FIPS 204, Module-Lattice-Based Digital Signature Standard (Final, August 2024) NIST FIPS 205, Stateless Hash-Based Digital Signature Standard (Final, August 2024) NIST IR 8547, Transition to Post-Quantum Cryptography Standards (Initial Public Draft, November 2024; public comment period closed January 10, 2025; final publication pending) 45 CFR 164.312(b), HIPAA Security Rule Audit Controls HHS HIPAA Security Rule NPRM, December 27, 2024 HTI-1 Final Rule, ONC/ASTP Algorithm Transparency Requirements Protecting AI and Cloud Competition in Defense Act, H.R.3434/S.1775, 119th Congress (introduced May 15, 2025; pending) FTC Health Breach Notification Rule, 16 CFR Part 318 Gramm-Leach-Bliley Act, 15 U.S.C. 6801 et seq. EU AI Act, Annex VII, Conformity Assessment Technical Documentation Access CSA AI Controls Matrix, July 2025 NIST Concept Paper, "Accelerating the Adoption of Software and Artificial Intelligence Agent Identity and Authorization," February 5, 2026 (Initial Public Draft) 21st Century Cures Act, Public Law 114-255, HL7 FHIR Interoperability Mandate

Appendix B: Governance Framework Context

VAISA is infrastructure. It does not prescribe any specific governance framework. Any multi-AI governance architecture that requires human checkpoint decisions to rest on verifiable evidence of what AI systems did with the data they processed gains that capability when VAISA's infrastructure standard exists. No governance framework is dependent on VAISA, and VAISA is not dependent on any framework. The relationship is technical enablement: governance frameworks gain verifiable evidence where they previously had contractual assurances. That shift from assurance to evidence is the difference between governance that is designed correctly and governance that functions correctly.

Supporting technical specifications for the four-profile classification system and the post-quantum cryptographic foundation referenced in this proposal are published at github.com/basilpuglisi/HAIA/tree/main/haia_agent. Federal validation through agency pilots would provide the independent operational evidence that congressional evaluation requires.

Appendix C: Bill Summary for Congressional Distribution

Short Title: Verified AI Inference Standards Act (VAISA)

Purpose: Close the data-in-use gap in existing federal data protection law by requiring that AI platforms processing regulated sensitive data provide customer-verifiable evidence of confidential processing conditions.

Key Provisions: Mandates a four-profile classification system for all AI inference involving regulated data, defaulting to prohibited without affirmative profile assignment. Directs NIST and HHS to publish a Verified Confidential Inference Standard within 12 months, defining five mandatory technical elements for compliant attestation. Establishes specific civil monetary penalty tiers mirroring HIPAA enforcement, with additional penalties for false compliance marketing. Creates a private right of action with statutory damages of $1,000 to $10,000 per violation plus attorney's fees. Provides a safe harbor for entities that implement compliant profiles, maintain signed evidence, and promptly disclose failures. Aligns enforcement across HHS OCR, FTC, ONC, CISA, and federal procurement. Establishes a federal floor with state

authority preserved above it. Requires post-quantum cryptographic readiness for all attestation and audit signing. Provides funding for implementation assistance for smaller covered entities through existing HHS 405(d) and new NIST channels. Phases compliance over 30 months with new deployment requirements beginning at month 12.

Existing Law Grounding: 45 CFR 164.312(b) (HIPAA audit controls), FERPA, Gramm-Leach-Bliley, H.R.3434 (Defense Competition Act, pending), 21st Century Cures Act (FHIR precedent).

Committee Jurisdictions: Senate HELP and Senate Commerce for healthcare, financial services, and standards dimensions. House Energy and Commerce and House Science, Space, and Technology for the technical standards and broader data protection dimensions.

Prepared for: 119th Congress **Contact:** BasilPuglisi.com **Date:** March 2026

Related Documents

This paper is part of the AI Provider Plurality Congressional Package:

Document 1: Summary Flyer (elevator pitch for infrastructure proposal)

Document 2: Ethics for Oversight (constitutional and philosophical case)

Document 3: AI Provider Plurality (legislative framework, policy mechanism, funding, appropriations)

Document 4: Methods Addendum (technical specification and operational evidence, v3.1 locked)

Document 5: Verified AI Inference Standards Act (this document; attestation API requirements and legislative framework)

Supporting technical documents:

- GOPEL Post-Quantum Cryptographic Agility Amendment v1.2 (March 2026, GitHub)
- GOPEL Confidential Processing Extension (CPE) v1.1 (March 2026, GitHub)
- HAIA-RECCLIN Agent Architecture Specification, EU Compliance Version (GitHub)
- Governing AI: When Capability Exceeds Control (Puglisi, 2025, ISBN 9798349677687)

Basil C. Puglisi, MPA

me@basilpuglisi.com | basilpuglisi.com | github.com/basilpuglisi/HAIA

March 2026 | AI Provider Plurality Congressional Package | Document 5 of 5