AI Provider Plurality: Infrastructure for Safe AI in America

Congressional Summary | February 2026, revised November 2026 | Basil C. Puglisi, MPA

THE ASK

1. Fund GOPEL as national AI infrastructure (the government builds the road, AI platforms are the vehicles)

2. Mandate API accessibility for AI companies serving federal agencies, federally regulated sectors, and high-consequence decision pipelines (vehicles must meet safety standards to use public roads)

3. Invest in small AI platforms to guarantee the competitive diversity that makes governance real (SBIR/STTR model)

4. Mandate verifiable inference attestation for AI platforms processing regulated data (VAISA, Document 5)

Phase 0 requires no new appropriation. Agencies adopt manual governance using existing resources to generate baseline data.

The Problem

AI shapes decisions in finance, healthcare, education, and national security. Control over these systems is concentrating in a small number of corporations. Executive Orders 14179 and 14365 call for a national standard that removes barriers to AI leadership without creating fifty state regulatory regimes. Single AI systems are proven flawed: hallucinations, bias, confabulation, alignment failures. Different AI systems produce different outputs on the same inputs. Geoffrey Hinton warns capability is advancing faster than control. Without structural safeguards, the digital equivalent of "too big to fail" moves from finance to cognition.

The Infrastructure Answer

The American public does not need more AI regulation. The American public needs AI infrastructure.

GOPEL (Governance Orchestrator Policy Enforcement Layer) is a non-cognitive governance agent: it dispatches, collects, routes, logs, pauses, hashes, and reports. Seven deterministic operations. Zero cognitive work. It cannot be co-opted because there is nothing to co-opt. If AI capability advances to the point of influencing its operators, a non-cognitive layer has no cognition to manipulate. This is not a proposal for more regulation. This is the engineering that makes less regulation safe.

Aviation	Highways	Finance	Telecom	AI
Airlines fly. FAA governs.	Cars drive. FHWA sets safety.	Banks trade. SEC oversees.	Carriers transmit. FCC regulates.	**AI platforms produce. GOPEL governs.**

Why This Is Different

Ethical AI establishes values (what AI should do). **Responsible AI** shapes machine behavior (how AI should operate). **AI Governance** exercises human authority (who decides). The first two rely on voluntary corporate compliance. Only the third builds structural accountability through infrastructure. That is what we are proposing.

Three agent models calibrate governance to risk: **Agent Model 1** (growth-speed, single checkpoint) for routine operations. **Agent Model 2** (structured gates, human approval at each role) for high-risk decisions. **Agent Model 3** (full manual orchestration) for highest-consequence work. Organizations choose. Markets function. The infrastructure accommodates both speed and accountability. This meets EU compliance standards while preserving American free markets.

Legislative Actions

Action	Mechanism
Fund GOPEL Infrastructure	Authorize NIST/GSA development of non-cognitive governance agent. Phased milestone-gated appropriation: funding released only after each phase clears validation gate. Phase 0 costs nothing
Mandate API Accessibility	AI companies serving federal agencies, operating in federally regulated sectors, or participating in high-consequence decision pipelines maintain audit compatibility with federal governance infrastructure. FTC enforces. Voluntary adoption elsewhere. Vehicles must meet safety standards to use public roads
Invest in Small AI	SBIR/STTR-style competitive grants for emerging AI platforms. Recipients maintain GOPEL API compatibility. Creates supply-side competition so governance has enough providers to work

Mandate Inference Attestation	AI platforms processing regulated data return verifiable evidence of processing conditions. NIST and HHS publish the standard. Phased from federal procurement to contingent mandate (Document 5)
Anti-Concentration	FTC/DOJ studies of exclusivity, compute concentration, resource consolidation (water, energy, compute). Extend antitrust tradition to cognitive infrastructure
Global Safeguard	At least two non-U.S. democratic-aligned providers in federal multi-provider systems (five or more platforms, unclassified). Expands cultural coverage beyond WEIRD training data. Five Eyes and NATO partners qualify

The Bipartisan Case

For Conservatives	For Liberals
Protects freedom, privacy, user choice. Prevents corporate cognitive cartels. Sustains market competition. Reduces government regulation through infrastructure. No content control	Protects fairness, transparency, accountability. Counters WEIRD bias through provider diversity. Prevents AI from amplifying existing inequities. Structural safeguards, not voluntary compliance

Evidence Base

Proven by others: Single AI systems produce flawed outputs (industry research, academic study, corporate safety evaluations). Automation bias causes humans to defer to AI under volume pressure (EDPS, Goddard et al., Banovic et al.). Concentration threatens every critical industry without structural checks (American antitrust tradition).

Built and operated as working concepts: HAIA-RECCLIN multi-AI governance under human arbitration, using Factics methodology (facts paired with tactics and measurable outcomes), across several hundred articles and a published book (2022 through 2025). Cross-platform disagreement observed in approximately 15% to 25% of tasks. Arbitration typically resolved under 30 minutes. One documented instance where eight of nine platforms were wrong and the dissenting platform was correct.

Proposed for federal development: GOPEL infrastructure at national scale. Specification and reference code published. Federal pilots produce the validated data.

These are working concept observations from single-practitioner development, not validated benchmarks. They provide feasibility indicators for agency pilot design.

The Package

Doc	Title	Purpose
1	Summary Flyer (this document)	Elevator pitch for infrastructure proposal
2	Ethics for Oversight	Constitutional and philosophical case
3	AI Provider Plurality	Legislative framework, policy mechanism, funding, appropriations
4	Methods Addendum	Technical specification and operational evidence
5	Verified AI Inference Standards Act (VAISA)	Attestation API requirements and legislative framework for AI inference data protection
6	Overwatch: The Cognitive Monitoring Shield for GOPEL	Working paper to support AI Provider Plurality

This is a pioneer path, not a finished product. It combines established concerns (Hinton on capability, antitrust on concentration, infrastructure precedent, automation bias research) into one architecture with one goal: safe use of AI for the American public. AI will never be absolute and without risk. The country needs to start.

Basil C. Puglisi, MPA
me@basilpuglisi.com | basilpuglisi.com | github.com/basilpuglisi/HAIA
Contact to discuss implementation, draft legislative language, or schedule a briefing.

#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). AI Provider Plurality: Infrastructure for Safe AI in America. basilpuglisi.com.