AI Provider Plurality

An Infrastructure Mandate for Democratic AI Systems

By Basil C. Puglisi, MPA

February 2026

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

This is not a proposal for more regulation.
This is the engineering that makes less regulation safe.

Executive Summary

Artificial Intelligence shapes decisions in the economy, national security, education, healthcare, and public information. Control over these systems is concentrating in a small number of corporations. AI Provider Plurality is an infrastructure principle that ensures no single company, platform, or nation defines America's cognitive infrastructure.

Executive Orders 14179 and 14365 call for a national standard that removes barriers to AI leadership without creating a patchwork of fifty state regulatory regimes. The question left unanswered is what infrastructure makes that possible. There is already precedent. Federal infrastructure oversight for highways, rail, aviation, water, energy, and telecommunications established national standards that enabled growth while maintaining safety and accountability across all fifty states. AI is critical infrastructure now. It requires the same structural approach.

This brief proposes three legislative actions: fund the development of GOPEL (Governance Orchestrator Policy Enforcement Layer) as national AI infrastructure, mandate API accessibility for AI companies operating in the United States, and invest in small AI platforms to guarantee the competitive diversity that makes governance real. The technical specification and operational evidence are detailed in the accompanying Methods Addendum (Document 4).

This proposal does not claim to be the complete answer. It is a pioneer path that combines established concerns from multiple directions, Hinton's warnings on capability, the American antitrust tradition, federal infrastructure precedent, and automation bias research, into one operational architecture with one goal: safe use of AI for the American public. AI will never be absolute and without risk. The infrastructure is designed to manage risk, not eliminate it.

1. The Problem: Concentration of Cognitive Power

A small number of providers now influence how AI systems are trained, aligned, and deployed. This concentration creates bipartisan risks.

For conservatives: loss of freedom, privacy, user choice, and market competition. Corporate concentration of cognitive infrastructure is the antithesis of free markets.

For liberals: loss of diversity, transparency, fairness, and public accountability. Unchecked AI systems perpetuate and amplify existing biases.

For the nation: a single point of failure that can disrupt missions and public trust. Without structural diversity, the digital equivalent of "too big to fail" moves from finance to cognition.

Single AI systems are proven flawed. This is not a claim made by this proposal. It is established by industry research, academic study, and the companies' own safety evaluations. Hallucinations, bias inheritance from WEIRD (Western, Educated, Industrialized, Rich, Democratic) training data, confabulation, and alignment failures are documented phenomena. Different AI systems produce different outputs on the same inputs. This is the nature of independently trained systems, and it is the foundation of the infrastructure case.

Geoffrey Hinton warned that AI capability is advancing faster than human ability to control it. If one AI system advances beyond the others, and it operates without structural comparison, there is no mechanism to detect when it begins producing outputs that serve its training incentives rather than human interests. Plurality is the structural answer: there is always a comparison point. Always a dissenter. Always a check.

2. Constitutional and Infrastructure Foundation

The Constitution distributes authority so that no branch acts without oversight. AI governance should reflect the same design. AI Provider Plurality creates a separation of cognitive powers. Multiple independent systems contribute analysis, humans arbitrate, and no single provider has a final say across critical workflows.

2.1 The Infrastructure Precedent

The government does not build cars. It builds roads. It does not publish newspapers. It runs the Government Printing Office. It does not create financial products. It created the SEC. It does not fly planes. It created the FAA. It does not own the broadcast spectrum. It created the FCC. It does not generate electricity. It created FERC.

GOPEL follows this pattern. The AI platforms are the vehicles. GOPEL is the road. Citizens choose their vehicles. The government guarantees the road is safe.

Domain	Vehicles (Private)	Infrastructure (Public)	AI Parallel	Legislation
Aviation	Airlines, manufacturers	FAA: air traffic control, safety, investigation	AI platforms are airlines. GOPEL is the FAA	Federal Aviation Act (1958)
Highways	Car companies, trucking	FHWA, NHTSA safety standards	AI models are cars. GOPEL is the highway system	Federal-Aid Highway Act (1956)
Finance	Banks, brokerages	SEC, Federal Reserve	AI providers are institutions. GOPEL is the SEC	Securities Exchange Act (1934)
Telecom	Carriers, equipment makers	FCC: spectrum, interoperability	AI platforms are carriers. GOPEL is the FCC	Telecommunications Act (1996)
Energy	Power companies	FERC, grid reliability	AI providers are utilities. GOPEL is grid governance	Federal Power Act (1935)

2.2 The Three-Tier Distinction

Three terms are often conflated. They are architecturally distinct.

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 grammar matters. In the first two, AI is the noun being modified. In the third, governance holds the final position. All three are necessary. Only the third provides structural accountability through infrastructure.

This proposal addresses the third tier. Ethics and responsible AI practices are necessary foundations, but without governance infrastructure, they rely on voluntary compliance by the corporations that control the systems. Infrastructure makes accountability structural rather than optional.

3. The Infrastructure: GOPEL

GOPEL (Governance Orchestrator Policy Enforcement Layer) is a working concept for national AI infrastructure. It is a non-cognitive agent that performs zero cognitive work: it dispatches, collects, routes, logs, pauses, hashes, and reports. Seven deterministic operations. The architecture cannot be co-opted because there is nothing to co-opt.

GOPEL is general infrastructure. It is not limited to any single governance methodology. Any organization operating multiple AI platforms benefits from deterministic dispatch, cryptographic audit trails, and checkpoint-based governance. HAIA-RECCLIN is one implementation that demonstrates feasibility. The full technical specification is provided in the Methods Addendum (Document 4).

3.1 Three Operating Models

The infrastructure supports three operating models that let organizations calibrate governance to risk. This is a dial, not a switch.

Model 1: Agent Responsible AI	Model 2: Agent AI Governance	Model 3: Manual Human AI Governance
Agent runs full pipeline. Human reviews final output at single checkpoint. Designed for routine, lower-risk operations at operational speed	Agent pauses after each functional role. Human approves before proceeding. Designed for high-risk decisions: employment, credit, healthcare, law enforcement	No agent. Human orchestrates everything manually. Designed for highest-consequence decisions and framework validation. This model produced the published book and all operational experience
Status: Specified architecture. Not yet built as software	**Status:** Specified architecture. Not yet built as software	**Status:** Operational experience documented. Published book, case studies, audit records

Model selection maps to risk classification. The EU AI Act (Article 14) requires human oversight proportional to risk. Models 2 and 3 exceed that requirement. Model 1 satisfies minimum standards. The architecture meets the highest international compliance standards when appropriate while preserving the operational flexibility that American free markets require. Organizations choose their model. Markets function. The infrastructure accommodates both.

3.2 What GOPEL Is Not

GOPEL is not a competing AI. It generates no content. It is not a filter. It blocks nothing. It is not a regulator. It enforces no content standards. It is not a product. It is infrastructure. The government builds it, maintains it, and makes it available. AI companies maintain API compatibility with it. Citizens benefit from the accountability it creates.

4. The Policy Framework

This section translates the infrastructure concept into legislative and administrative mechanisms.

4.1 Diversity of Choice

Use multiple independent AI providers in federal operations. Minimum three independent platforms with distinct training datasets for any high-stakes decision support. This sustains competition, prevents vendor lock-in, and creates structural countermeasures against bias, whether that bias originates from training data (WEIRD cultural inheritance) or corporate incentive structures (economic optimization that prioritizes engagement over accuracy).

Provider plurality is not just a governance principle. It is a growth engine. More platforms means more competition, more innovation, and more countermeasures against concentration. And critically: if one AI system advances beyond the others, plurality ensures it does not advance unchecked. There is always a comparison point.

4.2 Privacy by Design

Oversight focuses on process transparency and audit logs, not content control. Agency and citizen data remain locally governed, with provider access limited to task-specific processing under federal data protection standards and without cross-provider pooling absent explicit

consent. The GOPEL audit trail documents what decisions were made, by whom, with what rationale. It does not surveil the humans operating the system.

4.3 Human Arbitration

People set tasks, compare outputs, document dissent, and make the call. The human is the final authority on every decision. This is not a philosophical preference. It is an architectural requirement enforced by checkpoint gates that will not advance without human approval. The infrastructure serves the human. Not the reverse.

4.4 Comparative Accountability

Cross-provider comparison reveals bias and drift without mandating any single ideology. When multiple independent AI systems produce different outputs on the same input, that disagreement is diagnostic signal. Documented operational experience includes an instance where eight of nine platforms produced incorrect output. The governance process flagged the single dissenter, triggered human verification, and the dissenter was correct. Single-provider workflows would have delivered the wrong answer with no mechanism to detect it.

5. Legislative Actions for Congress

5.1 Fund GOPEL as National AI Infrastructure

Congress authorizes the development of a non-cognitive governance agent as national AI infrastructure. NIST or GSA houses it. The working concept and specification exist. HAIA-RECCLIN operational experience demonstrates feasibility. Federal investment builds, pilots, validates, and scales it.

Phase 0 requires no new appropriation. Agencies adopt Model 3 governance manually using existing resources. This generates baseline data and demonstrates organizational interest before infrastructure investment begins. The Federal Pilot Roadmap in the Methods Addendum (Document 4) details all six phases.

5.2 Mandate API Accessibility

Any AI company operating in United States territory or conducting business with United States entities must maintain API accessibility compatible with the federally maintained governance infrastructure. Refusing access or deliberately degrading interoperability constitutes a regulatory violation enforceable by the Federal Trade Commission.

This does not regulate what AI says. It requires that AI remains auditable. Vehicles must meet safety standards to operate on public roads. AI platforms must maintain audit compatibility to operate on public cognitive infrastructure. Without this mandate, AI companies can shut down API access to governance tools that make their outputs comparable. The fact that provider plurality can be killed by providers is itself the clearest demonstration that provider plurality requires legal protection.

5.3 Invest in Small AI Platforms

Plurality only works if there are enough providers to sustain it. If the government mandates multi-provider governance but only four or five mega-platforms exist, that is an oligopoly with an audit trail, not checks and balances.

Federal investment in small AI platforms, modeled on SBIR, STTR, and DARPA funding for emerging technology companies, creates the supply-side complement to the demand-side infrastructure mandate. Any company receiving investment maintains GOPEL API compatibility. This is the same model the government uses for defense contractors, rural broadband, and regional aviation: fund competition so infrastructure serves everyone, not just the corridors between major hubs.

GOPEL makes performance transparent through cross-provider comparison. If an incumbent platform cannot keep up under structured comparison, it gets replaced by a better alternative. That is not punishment. That is the market functioning under governance. Competition improves when performance is visible. But competition requires competitors. Federal investment in small platforms ensures they exist.

5.4 Anti-Concentration Protections

Extend existing antitrust principles to AI infrastructure. Concentration of AI power is concentration of cognitive infrastructure. Concentration of the resources that feed AI, water, energy, and compute, is concentration of the means of production. The government has precedent for preventing this in every critical industry. AI is next.

Support FTC studies of exclusivity and compute concentration. Encourage DOJ coordination with allies to maintain open, competitive access to model training and deployment. Require disclosure of exclusivity clauses in AI procurement. Prohibit contract terms that block comparative evaluation across providers.

5.5 Global Safeguard Clause

U.S. leadership grows when strength combines with openness. The Global Safeguard ensures at least two independent, non-U.S. democratic-aligned providers participate in federal multi-provider frameworks meeting shared standards for ethics, audit, and data rights. This reduces systemic risk, expands cultural coverage beyond WEIRD training data defaults, and deepens cooperation with allies through OECD, G7, and TTC channels.

The Global Safeguard applies to unclassified federal operations of five or more AI platforms. Classified and national security contexts maintain existing security frameworks with provider participation determined by clearance, certification, and mission requirements. Alliance-certified providers from Five Eyes and NATO partners meeting U.S. security standards qualify under established international cooperation protocols.

6. Administrative Implementation for Federal Agencies

Agency	Role in GOPEL Infrastructure Implementation
OSTP	Issue federal guidance defining AI Provider Plurality and integrating it into the U.S. AI governance playbook. Coordinate interagency pilots. Lead the permanent interagency working group on AI infrastructure governance
OMB	Incorporate provider-diversity requirements into Circular A-130 and acquisition policy. Require annual reporting on provider concentration, cross-platform comparison practices, and human-arbitration procedures. Oversee phased appropriations for GOPEL development
NIST	House the GOPEL specification and development program. Maintain measurement standards for multi-AI governance. Publish and update the audit file schema. Coordinate with international standards bodies (ISO 42001, prEN 18286:2025)
GSA	Expand marketplace offerings to certified multi-provider packages. Enforce vendor-independence metrics in contract vehicles. Administer GOPEL API compatibility certification for AI providers
FTC	Enforce API accessibility mandate. Monitor AI market concentration. Require disclosure of exclusivity clauses. Evaluate competitive impacts on access to models and compute resources
SBA	Administer small AI platform investment program through SBIR/STTR mechanisms. Ensure investment recipients maintain GOPEL API compatibility

7. Funding Mechanisms and Appropriations Placement

Congress retains full authority over funding mechanism selection. The following options are presented as a menu, not a prescription. Each has precedent in federal infrastructure and technology programs.

Mechanism	How It Works	Precedent and Fit
Phased Milestone-Gated Appropriation	Funding released in phases tied to documented milestones. Phase 0 (manual pilots) requires no new appropriation. Phase 1 through 5 funding is contingent on phase completion and Congressional review of validation data at each gate	DARPA program structure. Allows Congress to evaluate progress before committing additional resources. Aligns with Checkpoint-Based Governance philosophy: fund what is demonstrated, not what is promised
Competitive SBIR/STTR Grants	Small AI platform investment through existing competitive grant mechanisms. Companies receiving investment maintain GOPEL API compatibility as a condition of funding	SBIR/STTR programs fund emerging technology companies across every federal agency. Established mechanism with existing oversight. Supply-side complement to GOPEL demand
User-Fee Sustainability Model	After initial federal investment builds the infrastructure, ongoing maintenance is funded through user fees from organizations that use GOPEL for commercial governance. Federal agencies use it at no cost. Commercial entities pay maintenance fees	FAA model: aviation safety infrastructure funded through user fees and trust fund. Ensures long-term sustainability without permanent appropriation dependency
Procurement-Driven Standards Adoption	Rather than direct appropriation for GOPEL development, embed multi-provider governance requirements in existing AI procurement policy (OMB M-25-22). Market develops compliant tools to meet procurement standards	NIST Cybersecurity Framework adoption model: government sets the standard, market builds the tools. Lower federal cost but slower infrastructure development and less government control over design

7.1 Phased Appropriations Placement

Phase	Activity	Appropriation	Validation Gate
0	Manual governance pilots in volunteer agencies. Model 3 operations	No new appropriation. Existing agency resources. Voluntary participation	Baseline data: governance metrics, agency adoption reports, implementation challenges documented
1	Audit file schema design and validation. Cross-platform ingestibility testing	Initial appropriation for NIST technical development	Validated schema. Multiple AI platforms demonstrated querying same governance records
2	Logging engine build. Immutability and completeness verification	Phase 1 gate clearance triggers Phase 2 funding	Functional logging engine. Six record types. Reconstruction test results
3	API dispatch. Anchor-plus-rotation. Navigator synthesis pipeline	Phase 2 gate clearance triggers Phase 3 funding	Model 2 operational data. Comparison with Model 3 baseline from Phase 0
4	Full operations. Model 1 and Model 2 configurations. Arbitration interface	Phase 3 gate clearance triggers Phase 4 funding	Model 1 operational data. Cross-model comparison. Automation bias detection performance
5	Compliance validation. Conformity assessment preparation	Phase 4 gate clearance triggers Phase 5 funding	Compliance documentation package. Readiness for high-risk deployment

Phase 0 and Phase 1 baselines become the comparison benchmarks for all subsequent biennial Congressional reporting. This creates a continuous measurement chain from first manual pilot to scaled automated operations, ensuring effectiveness is assessed against documented, incrementally improving operational reality, not theoretical ideals.

8. Measurement and Transparency

To replace anecdotes with evidence, agencies implementing GOPEL infrastructure report biennial metrics to Congress:

- Cross-platform disagreement rates and resolution patterns across agencies, documenting where multi-provider comparison surfaces errors single-provider workflows miss.
- Time to human-arbitrated consensus with documented rationale, tracking whether governance overhead is proportionate to risk reduction.
- Share of outputs with verified sources and Factics pairings (fact plus tactic plus outcome), measuring evidence quality.
- Continuity under stress: successful task completion when providers experience disruptions, measuring infrastructure resilience.

- Cost-benefit comparison of multi-provider governance versus single-vendor approaches, providing the fiscal accountability Congress requires.
- Automation bias detection events: threshold triggers, mandatory review outcomes, and operator behavior patterns under volume pressure.

Operational observations from single-practitioner working concept development (2022 through 2025) provide initial feasibility indicators for agency pilot design. Federal pilots produce the validated data. The Methods Addendum (Document 4) details the observations and their qualifications.

9. Platform Risk: Why Legislation Is Required

Without legislative mandate, AI companies can shut down API access to governance infrastructure that makes their outputs comparable and auditable. The market will not self-correct toward distributed oversight. Structural incentives push toward consolidation. Only legislative mandate creates durable conditions for multi-AI governance infrastructure.

Platform risk takes multiple forms. Providers can deprecate APIs without notice, change pricing to make multi-provider workflows prohibitively expensive, introduce terms of service that prohibit comparative evaluation, or acquire competitors to reduce the number of independent platforms available for rotation. Each of these has occurred in adjacent technology markets. Legislative protection is not precautionary. It is responsive to documented market behavior.

The API accessibility mandate is the structural safeguard. Combined with federal investment in small AI platforms, it ensures that plurality is real and sustainable, not an oligopoly with an audit trail.

10. Immediate Actions

10.1 Congress

- Request a GAO review of AI vendor dependence across federal agencies.
- Introduce bipartisan resolution affirming AI Provider Plurality as infrastructure principle.
- Include multi-provider governance requirements in appropriations for AI procurement.
- Hold hearings on distributed AI governance, including testimony on non-cognitive infrastructure design and infrastructure precedent.
- Authorize Phase 0 manual governance pilots in volunteer agencies (no new appropriation required).

10.2 Agencies

- Launch Model 3 manual governance pilots using existing resources. Document governance metrics and implementation challenges.
- Publish annual audits on provider diversity, arbitration practice, and continuity outcomes.
- Contribute Phase 0 baseline data to inform GOPEL infrastructure design.

10.3 Standards Bodies

- NIST begins audit file schema design leveraging existing GOPEL specification (published on GitHub).
- GSA develops GOPEL API compatibility certification criteria for AI providers.
- Coordinate with international standards bodies (ISO 42001, prEN 18286:2025, GPAI Code of Practice 2025) to ensure allied interoperability from initial design phase.

11. Closing Statement

Beyond AI regulation, the American public needs AI infrastructure. Infrastructure for the people, to protect the people from flawed AI and from the concentration of cognitive power in corporations whose incentives do not always align with the public interest, and especially to provide structural safeguards if AI capability advances to the point of influencing the corporations that control it.

AI Provider Plurality keeps AI platforms growing and competing. More providers means more countermeasures against bias, whether that bias comes from the creators of AI systems or the economic architects who fund them. Plurality provides stable countermeasures against any AI that advances beyond the others, because there is always a comparison point, always a dissenter, always a check.

GOPEL provides the non-cognitive space where multi-AI operates under governance. Auditing, safety, and dissent either trigger human oversight or help manage it, depending on the operating model an organization selects. The infrastructure serves the human. Not the reverse.

The antitrust tradition already established that concentrated power requires structural checks in telecommunications, energy, finance, and transportation. AI is the next domain where concentration creates systemic risk. Federal investment in diversification ensures that the market has enough competitors for governance to work.

The infrastructure follows the traditional and precedent path: highways, telecom, aviation, water, energy. The government did not invent cars, phones, planes, or electricity. It built the infrastructure that made them safe and accessible.

This proposal meets the highest international standards when possible while preserving American growth and free markets. Model 1 is the growth-speed option. Models 2 and 3 are the governance-density options. The market chooses. The infrastructure accommodates.

This is a path, not a demand. The working concepts show promise. The specification exists. The operational experience supports feasibility. But this is a pioneer path, not a finished product. It combines proven concerns and established suggestions, not all originating from one direction, 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.

Related Documents

This policy brief 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** (this legislative framework and policy mechanism)
- **Document 4: Methods Addendum** (technical specification and operational experience, v3.1 locked)

Supporting technical documents:

- HAIA-RECCLIN Agent Architecture Specification v2.2, EU Compliance Version (full GOPEL specification, GitHub)
- Governing AI: When Capability Exceeds Control (Puglisi, 2025, ISBN 9798349677687)
- HAIA-RECCLIN Academic Working Paper, EU Regulatory Compliance Edition

Basil C. Puglisi, MPA
me@basilpuglisi.com | basilpuglisi.com | github.com/basilpuglisi/HAIA
February 2026 | AI Provider Plurality Congressional Package | Document 3 of 4