WHITE PAPER

Constitutional AI Governance:

Why Policy Frameworks Will Always Fail Without Structural Architecture

Kavanagh Industries LLC
Clinton Township, Michigan • kavanaghind.com
Published: April 2026 • USPTO Provisional Patent #63/991,057

EXECUTIVE SUMMARY

The American founders didn't just pass laws — they built structural architecture with constraints no simple majority could quietly dissolve. We need the same shift in AI governance. Today's frameworks treat data sovereignty as a policy preference, subject to terms-of-service updates, corporate lobbying, and the next court ruling. Kavanagh Industries proposes something fundamentally different: sovereign infrastructure where the protections are not contractual promises but engineered constraints — the way Asimov conceived his Three Laws — not as guidelines, but as hardcoded, immutable architecture. When sovereignty is built into the foundation, no platform update, no congressional session, and no judicial interpretation can move the walls.

That is not a product. That is a governing principle.

© 2026 Kavanagh Industries LLC — kavanaghind.com

FOREWORD

The Synthesis

> *Both of these ideas — the founders' and Asimov's — were written for moments exactly like this one.*

The founders were staring at a technology problem too — not AI, but governance itself. They'd watched every previous attempt at organized society collapse because the rules lived on paper, and paper burns. So they stopped writing rules and started building architecture. Separation of powers, checks and balances, supermajority requirements — those aren't policies. They're friction engineered into the system so that bad actors have to overcome structure, not just persuade people.

Asimov was doing the same thing from the other direction. He looked at the future of intelligent machines and said the same thing the founders said about power: you cannot trust a system to choose to behave. You have to make misbehavior structurally costly or structurally impossible.

This framework is the synthesis. It takes the founders' insight — build the walls, don't just write the laws — and Asimov's insight — hardcode the constraints, don't just publish the guidelines — and applies both to the exact moment where they're needed most.

> *Nobody else in this conversation is an engineer. The lawyers are writing briefs. The academics are writing papers. The politicians are writing legislation. The engineers are the only ones in the room who actually build load-bearing things for a living and understand what "structural" means in the physical world.*

That is not a coincidence. That is a credential.

Shaun Kavanagh
Founder & CEO, Kavanagh Industries LLC
Clinton Township, Michigan • April 2026

© 2026 Kavanagh Industries LLC — kavanaghind.com

I. The Problem: Policy Cannot Keep Pace

Every legislative and regulatory body currently engaged in AI governance is operating with the same flawed assumption: that rules written today can constrain a technology that rewrites itself tomorrow.

The pace of AI development has consistently outrun every attempt to codify it. The European Union's AI Act, enacted after years of debate, was already being stress-tested by capabilities that didn't exist when drafting began. The United States has relied on executive orders and agency guidance — instruments that change with every administration. Court rulings, as seen in the landmark 2026 cases of Morgan v. V2X, Warner v. Gilbarco, and United States v. Heppner, are filling the vacuum one fact pattern at a time.

This is not a criticism of lawmakers. It is a structural problem. Policy is reactive by nature. It must observe a harm before it can address one. Architecture is proactive. It prevents the harm from being possible in the first place.

The question facing every government, court, and institution engaging with AI is no longer "what rules should govern AI?" The correct question is: "How do we build systems where the protections are structural, not statutory?"

II. The Founding Principle: Architecture Over Legislation

The United States Constitution did not simply list rights. It engineered a system of competing authorities and structural constraints so that no single actor — no president, no congress, no court — could unilaterally dissolve what had been established. The Bill of Rights is not a policy document. It is load-bearing architecture.

The founders understood something that modern AI governance has forgotten: that the value of a protection is not in its words but in the cost of circumventing it. When circumvention requires overcoming structural resistance — supermajorities, independent branches, time — the protection has real weight. When circumvention requires only updating a terms-of-service document at 2:00 AM on a Tuesday, it has none.

Applied to artificial intelligence and data sovereignty, this principle demands a new category of thinking. Not "what policy governs this platform?" but "What is built into this platform such that no policy change can undo it?"

> *You cannot write policies fast enough to contain a technology that evolves faster than legislation. The answer isn't better rules — it's constitutional architecture, where the protections aren't written on paper, they're load-bearing walls engineered into the foundation.*

This is not a philosophical position. It is an engineering mandate. And it is the animating principle behind every system Kavanagh Industries builds.

© 2026 Kavanagh Industries LLC — kavanaghind.com

III. The Three Laws Precedent: Asimov's Structural Vision

In 1942, science fiction author Isaac Asimov introduced what he called the Three Laws of Robotics. Most people treat them as a literary device. That reading misses the point entirely.

Asimov's Three Laws were never intended as guidelines to be considered when convenient. They were conceived as hardcoded, immutable constraints — architecture that a robot could not override, rationalize around, or petition to have modified. The entire dramatic tension in Asimov's fiction arises not from robots choosing to violate the laws, but from the impossibility of doing so and the unintended consequences that emerge from perfect compliance.

Asimov saw, eighty years before the current AI governance debate, that the safety of intelligent systems is not a governance question. It is a design question. A system that "follows the rules" because it chooses to is fundamentally different from a system that cannot violate them because they are part of its foundation.

Modern AI governance has almost universally chosen the former: systems that promise compliance, that contractually agree to protect data, that publish policies about what they will and won't do. Kavanagh Industries is building the latter.

The RigidTrust framework — a constitutional Nine Bills of Rights governing every KI product and platform — operationalizes Asimov's insight at the infrastructure level. The protections are not promises. They are the architecture.

IV. The 2026 Legal Landscape and What It Reveals

Three federal court rulings in early 2026 have brought the AI governance question into sharp relief for practitioners, courts, and policymakers:

United States v. Heppner (S.D.N.Y., Feb. 2026)

A criminal defendant's communications with a public AI platform were ruled not protected by attorney-client privilege or the work product doctrine. The court found that materials created using consumer AI tools, without attorney direction, did not meet the elements required for protection.

Warner v. Gilbarco (E.D. Mich., Feb. 2026)

The Eastern District of Michigan reached the opposite conclusion in a civil case involving a pro se litigant, ruling that AI tools are instruments, not persons — and that disclosing information to an AI tool does not constitute disclosure to an adversary capable of breaking privilege.

Morgan v. V2X, Inc. (D. Colo., Mar. 2026)

A federal court in Colorado ruled that a pro se litigant's AI-assisted legal materials are protected under the work product doctrine under FRCP 26(b)(3), provided the litigant maintains a reasonable expectation of privacy. The court explicitly noted that electronic interaction passing through third-party systems does not automatically forfeit privacy protection.

© 2026 Kavanagh Industries LLC — kavanaghind.com

What these three cases reveal, taken together, is not a settled framework. They reveal a vacuum. Courts are reasoning from first principles on a case-by-case basis because no structural framework exists to guide them.

The Morgan court went further, explicitly acknowledging that its ruling creates a “technological gap” between litigants who can afford enterprise-grade secure AI infrastructure and those who cannot. That gap is not a court problem. It is an infrastructure problem. And infrastructure problems require infrastructure solutions.

V. The Policy Vacuum in Real Time

While the three court cases above were unfolding, the legislative and executive branches were simultaneously demonstrating — in real time — exactly why policy cannot solve a structural problem.

On March 20, 2026, the White House released its National Policy Framework for Artificial Intelligence. The document outlines the Trump Administration’s legislative recommendations to guide Congress toward a unified federal approach to AI regulation. It is four pages long. By the administration’s own description, it is non-binding and does not create new legal obligations.

On the same day, Representative Beyer introduced the GUARDRAILS Act, which would repeal the administration’s underlying executive order and block efforts to impose a federal preemption moratorium on state AI regulation. Two days earlier, Senator Blackburn had released a 291-page updated draft of the TRUMP AMERICA AI Act, taking a directly opposing prescriptive approach.

Two opposing legislative visions. One day. Neither addresses architectural sovereignty. Both are policy documents that will be contested, amended, lobbied, and potentially reversed by the next administration.

> *On a single day in March 2026, two directly opposing AI governance frameworks were introduced in Congress. Both are written on paper. Both can be dissolved by a change in administration. Neither asks the architectural question. This is not a failure of effort. It is structural proof that legislation cannot do what architecture can.*

The executive branch is not faring better. Executive Order 14365, signed December 11, 2025, directed the Commerce Department to evaluate “onerous” state AI laws within 90 days. That evaluation was due March 11, 2026. As of this writing, it has not been publicly released. The AI Litigation Task Force established to challenge state laws has been announced but has not yet acted.

Meanwhile, the Harvard Law Review published a case note on Heppner in March 2026 acknowledging that Judge Rakoff’s ruling “veers toward categorically excluding” AI from privilege protection, and calling for a more fact-dependent analysis that would protect AI use when confidentiality is properly maintained. Even elite legal scholarship is circling the same answer that infrastructure has already built: the key variable is not whether AI was used. It is whether the platform provides genuine, structural confidentiality — not contractual promises.

This is the vacuum. Courts deciding case by case. Congress divided. Executive orders non-binding. The White House’s own framework admits Congress needs to act — while Congress

© 2026 Kavanagh Industries LLC — kavanaghind.com

introduces legislation to undo the framework on the day it is released. The entire apparatus of policy governance is spinning in place while the technology accelerates.

Constitutional architecture is not waiting for Congress. It is being built now.

VI. RigidTrust: Constitutional AI in Practice

Kavanagh Industries' RigidTrust framework is the only current implementation we are aware of that applies constitutional architectural principles — rather than policy promises — to AI data sovereignty and system governance.

RigidTrust is not a product. It is the substrate — the connective architecture — through which every KI platform operates. It functions as a Nine Bills of Rights, with each bill encoding a structural constraint rather than a behavioral guideline:

- Data resides where the owner designates. The architecture enforces this; no policy update can override it.
- Processing occurs under owner-controlled parameters. The system cannot route data through unauthorized pathways.
- Deletion is verifiable and permanent. The architecture provides cryptographic proof of deletion, not a contractual promise.
- No training occurs on owner data. This is not a terms-of-service clause. It is a structural impossibility within the platform.
- Access is tiered and auditable. Every interaction is logged with immutable provenance through RigidVault.
- Sovereignty is portable. Owners can migrate from cloud-hosted to on-premises sovereign deployment without losing protection continuity.

The philosophical foundation of RigidTrust draws explicitly from Asimov's Three Laws. Where Asimov described laws that a robot "must not" violate, RigidTrust encodes constraints that the system architecturally cannot violate. The distinction is the entire point.

VII. The Tiered Sovereignty Model

Kavanagh Industries operationalizes constitutional AI governance through a three-tier deployment architecture designed to meet organizations and individuals wherever their sovereignty requirements place them:

Tier 1 — Cloud Node

Data is processed and stored within KI's sovereign infrastructure. All RigidTrust protections apply. The customer controls parameters; KI controls the architecture that enforces them.

Tier 2 — Sovereign Node

The customer's NAS hardware operates on their premises. AI processing occurs at KI via encrypted tunnel. The customer holds physical custody of their data; KI provides the intelligence layer under constitutional constraints.

© 2026 Kavanagh Industries LLC — kavanaghind.com

Tier 3 — Full Sovereign Node

Both NAS storage and Jetson-based AI inference operate on the customer's premises. KI monitors remotely. The customer owns the entire stack. Data never leaves their physical environment.

This tiered model is the infrastructure answer to the "technological gap" identified in Morgan v. V2X. It makes constitutional-grade AI sovereignty available at price points accessible to individual litigants, small businesses, and municipal governments — not just enterprise legal teams.

VIII. A Call to Governing Bodies

We are not asking courts, legislators, or regulatory agencies to recognize Kavanagh Industries. We are asking them to recognize the principle.

The current case-by-case adjudication of AI privilege, work product, and data sovereignty will not produce a coherent framework. It will produce an archipelago of conflicting decisions that serves neither individuals nor institutions. The legal community has identified this problem. Morgan v. V2X named it explicitly.

What is required is not more policy. What is required is a governing standard that distinguishes between:

- AI platforms that promise compliance through contractual language
- AI platforms that enforce compliance through structural architecture

These are not the same thing. Treating them as equivalent — as current policy frameworks generally do — fails both the courts trying to adjudicate privilege claims and the individuals trying to protect their data.

We propose that any meaningful AI governance framework — whether judicial, legislative, or regulatory — adopt structural sovereignty as the baseline standard. Not: "Does this platform have a privacy policy?" But: "Does this platform make privacy physically impossible to violate?"

> *The measure of a protection is not the strength of the promise. It is the cost of breaking it.*

When the cost of breaking a data protection is updating a terms-of-service document, the protection is not real. When the cost is overcoming cryptographic architecture, immutable audit trails, and owner-controlled physical hardware, the protection is constitutional in the only sense that matters: it is structural.

IX. Conclusion

The American founders built a system that has survived for nearly 250 years not because they wrote better laws than their predecessors, but because they built better architecture. They made the cost of tyranny structural — not merely illegal.

© 2026 Kavanagh Industries LLC — kavanaghind.com

Kavanagh Industries is applying that same principle to the most consequential technology challenge of this century. Not AI that follows rules. AI that is constitutionally incapable of breaking them.

The courts are building a framework one case at a time. The legislatures are writing rules that will be obsolete before they are signed. The corporations are publishing privacy policies that change with their business interests.

We are building walls.

That is not a product. That is a governing principle.

ABOUT KAVANAGH INDUSTRIES

Kavanagh Industries LLC is a Michigan-based sovereign technology and manufacturing company headquartered in Clinton Township, 5 miles from the Detroit Arsenal/TACOM. Founded by Shaun Kavanagh, a mechanical design engineer with 30+ years of experience spanning automotive, aerospace, defense, and biomedical industries, KI is a vertically integrated ecosystem built around 14 registered brand pillars united by a single governing principle: enabling others to own their destiny.

KI holds USPTO Provisional Patent #63/991,057 and 15 registered trade names with the Michigan LARA. Its RigidTrust framework has been in active development since the company's founding and is implemented across all KI products and platforms.

shaun@kavanaghind.com • kavanaghind.com/constitutional-ai • Clinton Township, Michigan

© 2026 Kavanagh Industries LLC — kavanaghind.com