

The Assurance Economy

Everyone is watching whether Congress passes a sweeping AI law. The real signal is the assurance market Washington is building underneath the argument, piece by piece.

SHANE SCHRECK FOUNDER, ALEETH JUNE 2026

[PUBLICATIONS.ALEETH.COM/THE-ASSURANCE-ECONOMY](https://publications.aleeth.com/the-assurance-economy) PDF

Everyone is watching the wrong fight. The question consuming AI policy is whether Congress will pass one sweeping law, and whether Washington will overrule the states. Neither answer is the signal. The signal is in the plumbing the federal government is laying underneath the argument.

Across the White House, Congress, NIST, and the defense committees, a quieter project is underway. Not a regulatory regime. An assurance ecosystem. Standards, evaluations, evidence requirements, accountability rules, and pre-deployment review are being assembled one piece at a time, and together they point at a market that does not have a settled name yet.

— WHAT WASHINGTON IS ACTUALLY BUILDING

Look at what moved this cycle, not at what made the headlines.

- › The White House issued an executive order on advanced AI innovation and security.
- › NIST's Center for AI Standards and Innovation, CAISI, signed agreements to run frontier-model evaluations and is becoming the technical center of gravity for federal AI standards.
- › Major frontier developers are now submitting models to voluntary federal evaluation before broad deployment.
- › A bipartisan House draft, the Great American AI Act, sketches a national governance framework.

- › The Responsible AI in Defense Act would require the Pentagon to test, document, and keep humans accountable for the AI it fields.
- › NIST and CAISI launched an AI Agent Standards Initiative for systems that act on a user's behalf.

No single one of these is a law that reshapes the market. Together they are something more durable. The federal government is deciding, in practice, that advanced AI has to be evaluated, documented, and proven governable before it is trusted in any setting that matters.

— ASSURANCE IS REPLACING COMPLIANCE

That is a different demand than the one most companies prepared for. Compliance asks what your policy says. Assurance asks what you can prove.

The federal conversation has already made the turn. It is no longer "show me your responsible-AI principles." It is "show me the evaluation, the logs, the human-accountability chain, the deployment review, the evidence that this system is under control."

A policy is a claim. Assurance is a record. Washington has started asking for the record.

— THE STACK, AND THE EMPTY LAYER

The architecture taking shape is legible once you stack it. Standards at the base, set by NIST and CAISI. Controls in the middle, where a standard becomes something a system actually does. Evidence above that, generated as the system runs. Assurance at the top, where a third party can verify the whole chain without taking anyone's word.

Standards are becoming infrastructure. But a standard does not enforce itself, and an evaluation does not translate itself into a control a company can run every day. Someone has to convert standard into control, control into evidence, and evidence into something a board, a regulator, or a procurement officer can trust. That execution layer is where the market is forming, and it is largely empty.

— DEFENSE MOVES FIRST. AGENTIC AI FORCES THE ISSUE.

Two forces will pull this forward faster than any comprehensive statute could. The first is defense. The Pentagon is moving to acquire, test, and deploy AI under explicit safety, accountability, and human-oversight requirements. Military AI is about to become one of the first large-scale operational assurance markets, because the cost of an ungoverned autonomous system there is not a redo.

The second is agentic AI. The moment AI stops answering questions and starts taking actions on behalf of people and institutions, governance can no longer be a document reviewed once a year. It becomes a control plane that runs continuously. NIST and CAISI naming agent standards a priority is the tell. The category is shifting from periodic and documentary to architectural and continuous.

— THE EXECUTION LAYER HAS A NAME

This is the layer ALEETH built. Institutional Control Architecture converts standards into controls, controls into evidence, and evidence into trusted deployment. A certification standard for autonomous systems. Seven control layers a system must establish to be certifiable. Seven failure modes that map how agents actually break. Cryptographically signed records any third party can verify without trusting the issuer. Continuous governance, not a one-time attestation.

It is not AI compliance software. It is institutional AI assurance infrastructure. The exact thing the federal architecture is about to require, and the layer that architecture does not supply on its own.

The comprehensive AI law everyone is waiting for may take years, and it may never arrive in one piece. The assurance market is being built right now, through standards, evaluations, procurement language, agency guidance, and defense requirements. By the time the category has a settled name, the institutions that can already prove governability will be the ones cleared to deploy. The rest will be retrofitting it under deadline.

Washington is no longer asking institutions to promise their AI is safe. It is starting to ask them to prove it.

That is the market ALEETH was built for.

ALEETH>

The Institutional Control Architecture. Court-grade certification for autonomous systems. Mathematically signed. Independently verifiable. Forever.

NOT PITCHED. NOT PROMISED. PROVEN.

Source: an open-source scan of U.S. federal AI-governance activity as of June 2026, including the White House executive order on advanced AI innovation and security (June 2, 2026); NIST CAISI frontier-evaluation agreements and program materials; the Responsible AI in Defense Act (Coons, Reed); the Great American AI Act House discussion draft (Oberholte, Trahan); the Future of AI Innovation Act (Young, Cantwell); and contemporaneous legal and policy analysis. This is strategic interpretation, not legal advice; primary sources should be verified before external citation.

SHANE SCHRECK · FOUNDER · ALEETH

INTELLIGENCE · INSTITUTIONAL · INEVITABLE