G1VE TECHNOLOGIES

The NSA Wants Every Model: Why AI Governance Just Became An Intelligence Mission

W.B. Carrington, MBA, LSSMBB, CTPRA, GRCA, GRCP

Aug 31, 2026

Original source

https://www.linkedin.com/pulse/nsa-wants-every-model-why-ai-governance-just-became-w-b--dzz3e/

I. The Pivot: Last Week AI Moved Into Operations. This Week The Intelligence Community Claimed The Keys.

Last week this newsletter argued that the AI war had moved out of policy papers and into live operations — FedRAMP authorized frontier models in federal workflows, AI agents unifying missile sustainment data, and machine speed cyber pressure reshaping defense. That argument aged well in exactly seven days.

Because on August 27, National Security Agency Deputy Director Tim Kosiba said the quiet part out loud at an Intelligence and National Security event in Bethesda:

- "We want access to all the models, and we're going to take advantage of that."

Read that sentence again. Not some models. Not approved models. All the models. The nation's premier signals intelligence and cyber organization — the same agency that has used AI in various forms for decades — is now formally positioning itself at the center of how the United States government evaluates, tests, and deploys frontier AI.

This is no longer a story about adoption. It is a story about who governs the governors. And for every leader in National Security, National Intelligence, and defense technology, the implications land this week, not next year.

II. What "Access To All The Models" Actually Means

Tim Kosiba was deliberately vague about mechanics, and that vagueness is instructive. "Access" can mean three very different things:

1. Using a model through a company controlled service,
2. Deploying a version inside a government environment, or,
3. Receiving pre-release access for security testing against classified networks.

Each arrangement carries radically different risk profiles, legal exposure, and operational value. The National Security Agency is pursuing all three simultaneously, and it is doing so under two of The White House 's instruments (June 2026) that most executives have not read closely enough:

1. The sweeping executive order directing National Security agencies to build classified tests that determine when an AI system has developed hacking capabilities advanced enough to warrant additional scrutiny.
2. The National Security Presidential Memorandum instructing intelligence and defense agencies to form "deep, proactive partnerships" with industry and make the most advanced models available to National Security personnel without delay — while explicitly warning against dependence on any single supplier.

The practical mechanism is a voluntary channel:

- Developers hand qualifying models to the government for up to 30 days before wider release, and the National Security Agency's director decides which systems cross the

"covered frontier model" threshold, in consultation with the Office of the National Cyber Director, The White House , and the Cybersecurity and Infrastructure Security Agency .

III. The Anthropic Saga Is The Governance Stress Test Nobody Wanted But Everybody Needed

If you want to see what AI governance looks like when it collides with National Defense reality, study the Anthropic dispute. The United States Department of War designated Anthropic a supply chain risk after the company declined to relax safeguards restricting certain military uses of its models. Anthropic sued. A federal judge temporarily blocked parts of the government's action.

...Meanwhile — and this is the part that should focus every boardroom — the National Security Agency carved out exemptions and is already running Anthropic 's advanced Mythos model experimentally, using it to test defenses protecting U.S. military networks and to evaluate weaknesses in networks belonging to China, Russia, North Korea, and Iran.

Then it got messier. Parts of the National Security Agency temporarily lost access to Mythos 5 during an export control dispute between the company and the administration — a live demonstration that commercial policy decisions outside any agency's control can interrupt intelligence operations mid-stream.

The restrictions were later lifted, but the lesson is permanent:

- When your mission capability rides on a vendor's terms of service, your readiness posture has a single point of failure wearing a corporate logo.

This is precisely why the June memorandum's anti-dependence clause — diversify suppliers, avoid single-company reliance — is the most strategically important sentence in the entire framework.

IV. Why The Intelligence Community Is Moving This Fast: The Capability Threshold Already Crossed

The National Security Agency is not doing this out of curiosity. Anthropic withheld the first version of Mythos from public release specifically because its ability to find and exploit previously unknown software vulnerabilities was judged too dangerous to ship. Since then, Anthropic , OpenAI , and other developers have fielded a growing family of models with similar capabilities — forcing policymakers to figure out how to test these systems, control access, and harness them for cyber defense without handing adversaries the same advantage.

OpenAI 's GPT-5.6 series went live for federal customers this month through its FedRAMP authorized enterprise service, with the most advanced variant aimed at software development, biological, and cybersecurity workloads.

Notice the asymmetry. Open weight models — the kind anyone can download and modify — were excluded from the voluntary pre-release testing program, even as Office of the National Cyber Director, The White House , publicly pushes for American open source AI to become the world's preferred standard.

...So the governance picture as of this week is:

1. Closed frontier models get classified pre-release scrutiny through the National Security Agency,
2. Open weight models get promoted globally but tested less formally, and,
3. Adversary states are shipping increasingly capable open models of their own.

If that sounds less like a coherent doctrine and more like three doctrines colliding in real time, that is because it is. The organizations that thrive will be the ones that build their own internal AI governance discipline instead of waiting for Washington's version to stabilize.

V. The Human Check Is Now The Most Important Control In The Stack

Tim Kosiba made one commitment that deserves to be quoted in every AI governance charter in the country:

- "What the models are putting out, how we're leveraging technology today, will always require that human to double-, triple-check and make sure we are in complete compliance with the law."

Coming from the deputy director of an agency whose business is intercepting the world's communications, that is not boilerplate. It is a doctrinal statement:

- Mission command over AI outputs remains human, accountable, and legally bounded — even when the machine is faster than the analyst.

Translate that to your organization. If the National Security Agency — with the deepest technical bench in the federal government — still insists on layered human verification before acting on model output, then any company or agency claiming its AI "runs itself" is either naive or selling something.

The operating model emerging from this week's disclosures looks like this: frontier models for speed and scale, classified testing for capability thresholds, diversified suppliers for resilience, and human adjudication for consequence.

That is not a technology stack. That is a governance doctrine, and JP 3.0 style mission command logic — decentralized execution, centralized intent, disciplined trust — maps onto it almost perfectly.

VI. The Strategic Criticality of the Next 90 Days

What serious leaders do in the next 90 days —

1. Map your organization's exposure to single-vendor AI dependence the same way you would map a single-source munitions component: The Anthropic access interruption proved this risk is operational, not theoretical.
2. Build an internal frontier model evaluation discipline before the government's voluntary program becomes a de facto procurement expectation: Because vendors that participate

in the 30 day pre-release channel will quietly accumulate a trust advantage in federal markets.
3. Assign a named executive owner for AI governance who answers questions in the language of mission assurance, not innovation theater: What happens when the model is wrong, unavailable, or legally constrained at the worst possible moment?
4. Watch United States Congress , the United States Senate Committee on Armed Services , U.S. House of Representatives Armed Services Committee, and appropriators.

...The intelligence community's aggressive posture — access to all models, classified testing regimes, supplier diversity mandates — will translate into contract language, reporting requirements, and supply chain certifications that land on industry within budget cycles, not decades.
The firms reading this week's signals as procurement pre-positioning will be the ones holding favorable positions when the language hardens.

VII. Thought-Provoking Close: The Bottom Line

Last week:

- AI moved into operations.

...This week:

- The Intelligence Community moved into AI.

The National Security Agency 's demand for access to every model — backed by executive orders, a 30 day pre-release testing channel, and live experimental deployments against adversary networks — marks the moment AI governance stopped being a compliance conversation and became an intelligence mission.
The winners will be the organizations that treat model access, vendor independence, and human verification as readiness functions. Everyone else will discover, the way the National Security Agency briefly did with Anthropic 's Mythos 5, that capability you do not control is capability you do not have.

✨ Call to Action: Step Into the Arena ✨

Most professionals will read about the future. A select few will shape it.
If you are serious about staying ahead of the curve in Artificial Intelligence, Regulatory Compliance, and Enterprise Risk, this is where you need to be.
📩 Subscribe to Future Trends & Insights: AI to gain direct access to:

1. Strategic foresight on AI Governance & Ethical Deployment
2. Real-world applications across National Security & Mission Assurance
3. Advanced GRC, TPRM, and Regulatory Compliance frameworks
4. Emerging Fintech and Risk Intelligence innovations

⚔️ The landscape is shifting—fast. The gap between those who understand and those who execute is widening.
This is built for professionals responsible for decisions that carry consequence—not commentary.
📅 Ready to operationalize these insights and move beyond theory?
👉 https://calendly.com/g1ve-technologies
Let us translate insight into measurable, defensible outcomes—and build capability, not just conversation.

🔍 About the Author

W.B. Carrington, MBA, LSSMBB, CTPRA, GRCA, GRCP is a National Security-aligned Artificial Intelligence Strategist/Tactician and Governance, Risk, and Compliance (GRC) Authority/Subject Matter Expert with over a decade of experience operating at the intersection of National Security, National Defense, Critical Infrastructure, Cybersecurity, AI Governance, Mission Assurance, Mission-Critical Systems, and Predictive Funding/Grant Intelligence.
As Managing Partner at G1VE Technologies, LLC and a Strategic Board Advisor, W.B. Carrington, MBA, LSSMBB, CTPRA, GRCA, GRCP architects and operationalizes enterprise-scale risk, compliance, and AI governance ecosystems across both Public and Private Sector environments—including Fortune 500/1000 organizations and Federal-aligned missions, where failure is not an option.
His work is centered on delivering Decision Advantage, strengthening Operational Resilience, and enabling organizations to navigate high-stakes, high-uncertainty environments with precision and confidence.
📡 Connect with W.B. Carrington, MBA, LSSMBB, CTPRA, GRCA, GRCP here to stay at the forefront of:

- Artificial Intelligence
- National Security
- GRC
- Cyber Risk
- Regulatory Strategy
- Mission Assurance