G1VE TECHNOLOGIES

Who Has the Power to Stop It?: Why AI Authority, Override Rights, and Federated Governance Are the New Chain of Command

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

Jul 13, 2026

Original source

https://www.linkedin.com/pulse/who-has-power-stop-why-ai-authority-override-rights-w-b--pzkfe/

I. The Uncomfortable Question No One Wants To Hold

The federal AI conversation in 2026 has advanced far beyond “Should we use it?” It has jumped straight to procurement volume, model deployment inside classified environments, agentic workflows, and mission acceleration. That is progress. But it has skipped over a question that will decide whether the coming wave of federal AI deployments ends in operational advantage or institutional embarrassment: who, exactly, has the authority to stop an AI system when it starts to drift, misfire, exceed its role, or produce consequential harm?

That is not a philosophy question. That is a chain-of-command question. It is a mission-assurance question. It is a Joint Force question. It is a Senior Executive Service question. It is a question about JP 3.0-style operational authority applied to systems that increasingly touch procurement, cyber defense, intelligence workflows, logistics orchestration, benefits adjudication, health screening, ISR triage, and command decision support. If leaders cannot answer it in plain English before deployment, they will be forced to answer it in a hearing after failure.

The federal AI story in 2026 is not a story about model capability. It is a story about authority, override rights, non-human identity, delegation logic, and executive accountability. Everything else is decoration.

II. Speed Without Authority Is Not Modernization. It Is Exposure.

An April 2026 commentary in Nextgov/FCW framed the problem sharply: the federal government is spending a great deal of time talking about responsible AI, but not nearly enough time deciding who gets to stop it. Governance principles like fairness, transparency, and explainability matter, but they operate at the level of ideas. Real risk lives at the point of use. And at the point of use, agencies keep confusing “we approved this system” with “we know who has the authority to pause it.” Those are not the same thing.

The commentary lays out three anchors that federal agencies still need to install. First, explicit internal override rights so someone can slow, narrow, or halt the use of a system without navigating an internal maze of legal, procurement, and technical review after harm has already occurred. Second, auditable decision trails so accountability is real rather than performative — meaning someone can reconstruct who approved a system, what limits attached to that approval, what changed over time, and who kept relying on it. Third, procurement as governance, meaning contracts must preserve access to logs, change notices, testing evidence, intervention rights, and usable explanations rather than surrendering those levers to vendors.

The most quotable line — and the one that should be on the wall of every AI governance office in the executive branch — is this: “Speed without decision control is not modernization. It is exposure.” That is the entire thesis. Any agency, program office, combatant command support element, or civilian department that cannot answer who has the power to stop an AI system before it fields that system is not governing AI. It is buying it.

III. Meanwhile, The White House Told National Security Agencies To Accelerate

The tension is that the strategic environment is pushing the other way. President Trump's June 2026 national security memorandum directed the United States Department of War , Office of the Director of National Intelligence , the Federal Bureau of Investigation (FBI) , the Office of the National Cyber Director, The White House , and other national-security agencies to move faster on advanced AI adoption — deeper partnerships with AI companies, faster availability of frontier models to national-security personnel, and hard focus on protecting U.S. models from foreign theft, manipulation, and industrial-scale distillation.

That is a speed memo. It sits on top of an AI security executive order that encourages developers to give the U.S. government 30 days of pre-release access to frontier models, establishes classified benchmarking, and gives National Security Agency a formal role in evaluating which systems cross into a higher-risk category. The two documents together tell agencies: adopt faster, protect harder, and evaluate more seriously.

The problem is not the direction. The problem is what happens if adoption outruns internal authority structures. If frontier models enter classified networks, agentic tools begin acting inside acquisition, cyber defense, ISR triage, and intelligence workflows, and agencies still cannot cleanly answer who has the power to stop them, then acceleration becomes a governance liability rather than a mission advantage.

IV. Zero Trust Was Not Designed For Entities That Never Sleep

The next layer of the problem sits deeper than governance memos. It sits inside the identity architecture itself. A July 2026 commentary made a sober observation: after four years rebuilding federal cybersecurity posture around zero trust under EO 14028 and OMB M-22-09, the current architecture was built on a human assumption. Zero trust in its current form assumes the entity behind a request is a human — logging in at human speed, holding permissions that change on human timescales, and clicking through interfaces the way humans do. Agentic AI invalidates every one of those assumptions at once.

The numbers are jarring. An ATARC (Advanced Technology Academic Research Center) analysis cited in that piece calculated that 1,000 agents operating at machine speed can generate roughly 7.4 million authentication events per day — about 148 times the volume an equivalent human population would produce. No federal identity stack is sized for that load. Even worse, most stacks do not carry the audit metadata needed to forensically investigate what an autonomous agent did, on whose behalf, under what authority. That is a governance vacuum operating at machine speed.

The default temptation for many agencies will be to onboard agents as service accounts. That would be a strategic mistake. Service accounts are already among the weakest links in federal zero trust — static credentials in vaults, broad standing privileges, infrequent rotation, thin attribution. Scaling that pattern into the largest non-human identity population in federal history would take the weakest link in the chain and make it the biggest. That is not a technical inconvenience. That is a mission-assurance problem waiting to happen.

V. What Agent Identity Should Actually Look Like

The same commentary offers a genuinely useful blueprint that federal leaders across the United States Department of War , IC, and civilian components should be reading, not skimming. Agents need the opposite of what service accounts get: cryptographically verifiable identity that travels with them, narrow permissions scoped to specific tasks, and credentials that expire on the order of minutes. That is not a nice-to-have. That is baseline hygiene for a mission environment where machine-speed action can create machine-speed damage.

Three additional shifts should happen in the next two budget cycles. First, cryptographic identity should be a baseline requirement for any production agent, using decentralized identifiers and verifiable credentials, so an agent can prove what it is, who deployed it, and what it is authorized to do — without pinging a central registry on every request. Second, delegation chains between agents and sub-agents should be auditable end-to-end, with permissions narrowing at each hop and revocation cascading in seconds. Without that, incident response after a compromised agent has no way to bound the blast radius. Third, the federal zero trust mandate should be formally extended to non-human autonomous identities, backed by tailored Office of Management and Budget guidance and National Institute of Standards and Technology (NIST) profiles — not shoehorned into a service-account category never designed to absorb entities that set their own sub-goals.

These are not exotic ideas. They are the natural extension of mission-assurance logic into an environment where AI agents will increasingly act inside logistics, sustainment, ISR triage, benefits processing, IC workflows, cyber defense, procurement, and JADC2-style information flows. If agencies fail to build this now, the postmortem will write itself.

VI. What This Actually Means For The Joint Force And The IC

For the The Joint Staff , combatant commands, the United States Department of War / DoD, service headquarters, National Security Agency , Defense Intelligence Agency , U.S. Cyber Command , Joint Special Operations Command , and the intelligence community writ large, the operational takeaway is direct. AI is entering mission systems faster than the identity, authority, and audit architecture around it is maturing. That gap is now the primary strategic risk in federal AI, not model quality.

The mission-command consequence is that Federated AI cannot be treated as an IT program. It has to be treated as a command-and-control architecture — one that defines authority, escalation, override, revocation, and accountability at every layer, from the model provider down to the individual agent action. Without that framing, agencies risk building a shadow chain of command where the machines act, humans watch, and no one is entirely sure who owns the halt authority when things break. That is unacceptable for mission-assurance and unacceptable for national defense.

For SES leaders, agency directors, undersecretaries, secretaries, and program executives, the personal exposure question is even sharper. If your agency's AI stack is compromised, degraded, or misused, and no one can produce a clean paper trail showing who approved it, what limits attached, and who had the power to stop it, then the accountability defaults to leadership. That is not a hypothetical. That is how oversight actually works.

VII. What Premier Leaders Should Do Now

There are five moves that separate premier organizations from the ones that will be reactive after their first incident.

First, designate an executive sponsor for AI authority, not just AI adoption. Adoption belongs to program owners. Authority to stop belongs to executive leadership. Those must be different chairs.

Second, build an override doctrine before deployment, not after. Every AI system supporting mission-relevant work should ship with a documented halt authority, a defined escalation ladder, and a clearly identified revocation point of contact.

Third, treat agent identity as a first-class citizen. Cryptographic identity, minute-scale credentials, narrow task scopes, and cascading revocation should be procurement requirements, not aspirations.

Fourth, use procurement as a governance instrument. Contracts should lock in log access, change notices, testing evidence, override rights, and interpretability provisions. If a vendor cannot supply them, they should not be inside mission systems.

Fifth, protect the model layer as a national-security asset. Foreign distillation, insider misuse, prompt exfiltration, weight leakage, and data-center exposure are counterintelligence risks now — not just IT problems. The White House already said so out loud.

VIII. Thought-Provoking Close: The Final Assessment

The premier AI-enabled organizations of the next five years will not be defined by which model they chose. They will be defined by whether they built a legitimate chain of command around it. Model access is now a commodity conversation. Authority is not.

Who has the power to stop it? That is the question 2026 is asking every federal leader — quietly, but with rising volume. The ones who answer it before an incident forces the question will hold the operational advantage. The ones who wait will be answering it under lights, in front of a committee, with a very different tone of voice.

✨ 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