G1VE TECHNOLOGIES

Why Washington Is Planning Agentic AI Faster Than It Can Assign Authority

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

May 10, 2026

Original source

https://www.linkedin.com/pulse/why-washington-planning-agentic-ai-faster-than-can-w-b--fzyge/

I. The Federal AI Debate Has Entered A More Dangerous Phase

The federal government is no longer just experimenting with artificial intelligence in narrow, low-stakes settings. It is now moving toward agentic AI — systems designed not merely to assist, summarize, or recommend, but to pursue multi-step goals and take bounded actions across software environments with increasingly limited human touch. In May 2026, a survey of more than 200 technology executives across civilian and defense agencies found that 53% of agencies were exploring agentic AI or actively planning pilots, while another 15% said they were already implementing these systems or had effectively done so already. That is not marginal momentum. That is a clear signal that the federal enterprise is preparing to operationalize machine action, not just machine analysis.

That distinction matters because once a system can act — initiate workflows, move information, trigger downstream processes, or shape operational tempo without waiting for repeated human prompts — the governance problem changes completely. The issue is no longer whether agencies understand AI in theory. The issue becomes whether they have built a command structure for machine autonomy inside government. And right now, the evidence suggests Washington is moving faster on capability than it is on control.

II. Appetite Is Racing Ahead Of Authority

The most troubling part of the May survey was not the speed of adoption. It was the gap between what agencies say they need and what they have actually implemented. Seventy-seven percent of respondents said oversight frameworks for agentic AI are essential, yet fewer than a third have put such frameworks in place. Only 20% said their agencies have defined policies for pre-deployment testing or generic agentic AI use. Only 8% reported having a defined framework for incident response. Only 6% had a formal framework for third-party or vendor governance. And fewer than one-third — 29% — had documented kill-switch procedures.

That last figure should disturb anyone who takes mission assurance seriously. Because if a government system can act at speed, then the core governance question is not whether leaders can describe responsible AI in a memo. The question is whether they can interrupt, constrain, suspend, or terminate machine action quickly when a system drifts, misfires, or exceeds its intended role. If there is no clearly documented technical and organizational path to do that, then the government does not truly control the tool. It is simply relying on optimism and vendor promises.

III. “Responsible AI” Means Very Little If Nobody Owns "The Stop Button"

A recent Nextgov/FCW commentary cut to the heart of the matter with unusual clarity: the federal government is spending more time talking about responsible AI than deciding who gets to stop it. That is exactly the right critique. Washington’s AI language still leans heavily on

familiar abstractions — fairness, explainability, accountability, risk management, transparency. Those principles matter, but they are not operational answers. They do not tell a deputy secretary, CIO, mission owner, or risk executive who has the authority to halt a live system when new conditions make continued reliance dangerous.

The same commentary argued that government does not encounter risk at the level of principle. It encounters risk at the point of use. A model may pass legal review, procurement review, internal policy review, and still become unsafe when assumptions change. Maybe the data quality degrades. Maybe the environment becomes more adversarial. Maybe a workflow originally considered low-risk expands into a mission-critical lane. Maybe a vendor changes the system in ways not fully visible to the government. At that point, "responsible AI" as a slogan is worthless unless someone has clear override authority, and unless the institution can reconstruct what was approved, what limits applied, and who chose to continue relying on the tool.

IV. This Is A Command-And-Control Issue, Not Just A Technology Issue

For national security professionals, the right frame is not "digital transformation." It is command and control.

No military organization worthy of the name would deploy a fast-moving operational system without clarity about authorities, intervention rights, escalation thresholds, and fallback procedures. Yet that is disturbingly close to what parts of the federal government are now doing with agentic AI. Tools are being discussed as productivity enablers and modernization accelerators, but their most important characteristic is often underplayed: they can alter decision velocity inside institutions that were not built for machine tempo.

That means federal leaders need to stop treating agentic AI as if it were merely a better software suite. It is closer to an internal autonomy layer. And once you recognize that, the required governance model becomes clearer. Every meaningful agentic deployment needs:

- Clearly assigned operational ownership;
- Human override rights with real authority;
- Auditable action logs;
- Decision thresholds tied to mission risk;
- Tested intervention pathways; and,
- Contractual rights that preserve government control over logs, changes, and recovery.

Those are not "nice-to-haves." Those are the minimum conditions for credible use.

V. The Federal System Actually Knows What High-Risk Oversight Should Look Like

To be fair, the federal enterprise appears to understand some of this instinctively. The May survey showed that oversight expectations rise dramatically when the data or mission is high consequence. For national security, critical infrastructure, and emergency response data, 79% of respondents said agencies require human-in-the-loop approval for every action performed by AI. For high-risk data such as benefits claims or financial data, 78% said agencies require

formal human approval before high-risk actions are taken, though not for every action. By contrast, more than 90% favored lower direct involvement for low- and moderate-risk data.

That is a promising foundation because it shows federal leaders do not actually believe all AI autonomy should be treated equally. They already understand that mission context matters. But the fact that agencies recognize risk gradations does not mean they have built the architecture to enforce them. Intent without mechanism is not governance. It is policy aspiration.

And this is where the kill-switch gap becomes truly important. Human approval requirements are only meaningful if the institution can technically and procedurally intervene when reality diverges from planning. If a system is authorized to act in a bounded way, then the government needs a repeatable way to narrow those bounds when indicators change. Without that, "human oversight" becomes ceremonial.

VI. Standards Are Becoming A Battlefield For Control

This is one reason the GSA – National Institute of Standards and Technology (NIST) partnership announced in March matters so much. The two agencies launched a joint effort through the Center for AI Standards and Innovation to create common benchmarks, testing methods, guidelines, and checklists for AI models and services used in federal operations. Their aim is to standardize how agencies evaluate AI before moving it into real mission environments and to reduce redundant effort as pilots move toward deployment.

At first glance, this may sound bureaucratic. It is not. In an agentic environment, evaluation standards become part of mission assurance. If the government lacks common ways to test how a system behaves under stress, what its failure modes look like, how its outputs degrade when data changes, or how it should be monitored after deployment, then speed becomes dangerous. An agentic system that moves fast without a matching evaluation discipline is not modernization. It is exposure at machine speed.

The strategic point is simple: standards are not paperwork. They are one of the only ways large institutions preserve control when systems become too dynamic to manage informally

VII. The Workforce Problem Is Deeper Than Most Leaders Admit

Even if the policy and standards picture improves, there is another structural weakness beneath the surface: human capital. In April, Nextgov/FCW reported that 88% of federal decision-makers see AI as a critical tool for modernization, but 44% identified workforce skill gaps as a major barrier. 95% said they were focused on upskilling current employees, and 96% said agencies should focus on streamlining and refining internal processes before implementing new technology broadly.

The hidden issue here is that agentic AI requires more than technical users. It requires institutional supervisors of autonomy. It requires people who know how to interpret logs, define intervention triggers, investigate drift, challenge vendor claims, distinguish minor anomalies from systemic failure, and understand when a workflow is being pushed beyond the boundaries originally approved.

That is not just a software skill set. It is a governance skill set. And it is one the federal government does not yet have at sufficient depth.

If Washington wants to move fast on agentic systems while also reducing headcount pressure, it risks creating a false economy: fewer humans overseeing more autonomous processes in more consequential settings. That may produce nominal efficiency while degrading actual control.

VIII. Procurement Is Now A Governance Weapon

Another truth too many agencies still ignore: procurement is no longer just a purchasing process. It is one of the government's strongest governance tools.

The authority-gap commentary made this point directly: contracts determine whether the government retains access to meaningful logs, receives change notices, can demand intervention rights, sees usable testing information, and preserves visibility into system behavior. Weak terms mean the government may buy functionality while quietly surrendering control.

That is especially dangerous in the agentic era because the system does not sit still after award. Vendors update models. Interfaces evolve. orchestration layers change. Dependencies multiply. If contracts do not preserve the government's right to inspect, intervene, pause, test, and demand post-incident transparency, then the institution will be governing a moving target with stale tools.

In a defense context, that would be recognized immediately as unacceptable. The same standard should apply here.

IX. What A Warfighter-Grade Agentic Governance Model Should Include

If Washington wants to move from aspiration to control, it needs a harder-edged governance model built around operational realities rather than ethics branding.

That model should include at least six things.

First, pre-deployment testing rules tied to mission and data criticality. Second, explicit override authorities assigned before fielding. Third, documented kill-switch procedures that are tested, not merely drafted. Fourth, risk-tiered autonomy, where national security and critical-infrastructure actions encounter tighter human approval than administrative workflows. Fifth, contractual rights that preserve government access to logs, system-change visibility, and intervention pathways. Sixth, post-incident review architecture so failures become contained, repeatable workflows rather than public crises.

This is not overkill. It is the baseline for using agents in any institution that takes operational credibility seriously.

X. Bottom Line

The kill-switch gap is not a niche policy issue. It is the clearest available indicator that the federal government is moving into agentic AI faster than it is building the institutional power to govern it. Agencies are piloting and planning agentic systems aggressively, but oversight

maturity, incident-response design, vendor governance, and true intervention capability remain weak.

Washington increasingly understands the need for human approval in high-consequence domains, and standards work is moving in the right direction, but until the government can answer one blunt question before deployment — who gets to stop the system, when, and how — it is not governing agentic AI. It is fielding autonomy ahead of authority.

✨ 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