From Adoption to Accountability:

Building the Trusted Foundation for Agentic AI in Government

Market Connections
Research you can act on.

SPONSORED BY servicenow

As federal agencies pivot from generative AI pilots to agentic systems—AI capable of taking independent action to achieve goals—the mission is shifting from "what is possible" to "what is governed." While the promise of autonomous workflows offers a solution to the public sector's most pressing resource constraints, it also introduces a new set of risks. True digital transformation in the federal space requires more than just high-performance models; it demands an integrated control tower that bridges the gap between ambitious AI pilots and secure, enterprise-scale execution.

THE MOMENTUM PARADOX

The AI landscape is fundamentally different than it was even a year ago, moving beyond simple chatbots toward systems that interact with core mission data. Yet while the appetite for agentic AI systems is high – only 6% of agencies are not considering the technology at all — structural execution remains in its infancy. The gap between those piloting and those in complete workflow integration indicates a "pilot purgatory" where agencies struggle to move AI from isolated sandboxes into the production environment. The window to establish foundational governance is now, before these disparate efforts crystallize into unmanageable silos.

of agencies are in early stages, exploring or planning pilots.

of agencies are currently running a pilot.

of agencies have fully or partially implemented agentic systems.

of agencies have completely integrated agentic AI into their workflows.

AGENCY PERSPECTIVE

"Our eyes are bigger than our stomach. I don't necessarily know that we've gotten our arms around the problem, but we have ambitious plans for AI." – Senior Enterprise Architect, FedCiv

BUILDING GOVERNANCE THAT WORKS

The primary obstacle to scaling AI is a disconnect between the perceived necessity of oversight and the actual existence of institutional frameworks. While the majority (77%) of federal leaders view defined oversight frameworks as "essential," only a small fraction (20-30%) have actually implemented them.

This oversight gap is most visible in the specialized areas of risk management. While nearly a third of agencies have frameworks for continuous monitoring of AI systems—the top current use case—adoption drops significantly for policies related to pre-deployment testing and general AI use. Even fewer agencies have defined frameworks for incident response, and less than one in 10 have established governance for third-party or vendor oversight.

The deeper, more foundational barrier is not just policy, but data readiness. Agencies are increasingly realizing they are not prepared to feed their data to AI agents safely or effectively. Without a unified platform to orchestrate both data and policy, AI remains an uncoordinated set of risks rather than a mission-enabling asset.

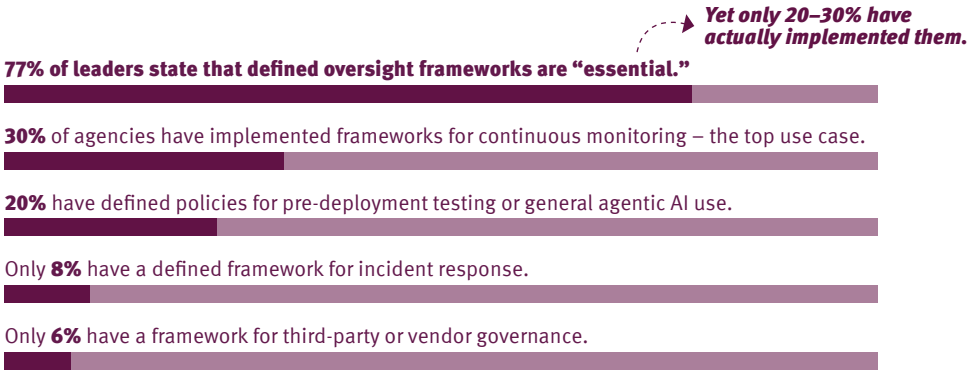

AGENCY PERSPECTIVE

"One of the very big questions that you have to ask is 'how am I getting my data ready for AI consumption?' That governance piece becomes critical [to] making sure that your data within your organization and your AI are working together." – Director, IT Modernization, FedCiv

AUTOMATING HUMAN-IN-THE-LOOP

Trust in agentic AI is built through granular control. Federal leaders already tier their risk to determine exactly where and how, throughout an agentic AI process, a human must intervene. The challenge today is not defining these risk tiers, but operationalizing an agency's response when intervention is needed. Humans need to be in the loop for mission-critical decisions, but those processes can be unreliable. As agentic AI takes a larger role in automating decisions, agencies need equally automated guardrails and decision trees to reliably trigger these interventions when needed most.

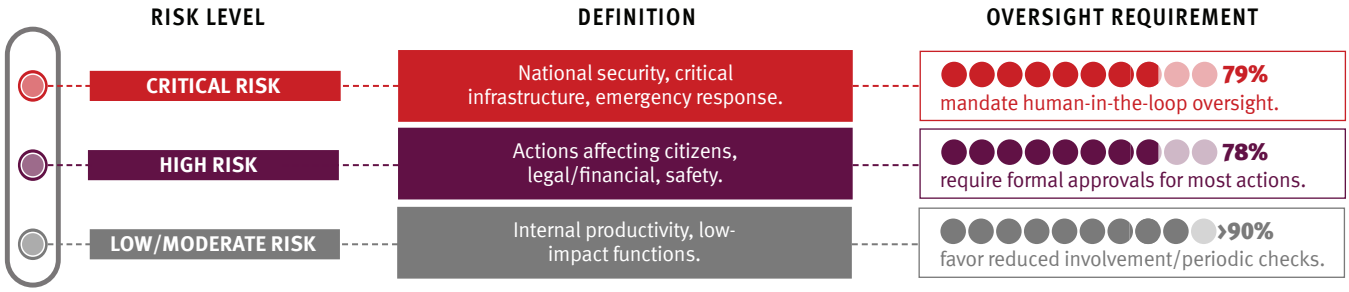

AGENCY PERSPECTIVE

"There have to be some parameters. Who are you going to hold accountable if AI makes a wrong decision? There's some things that cannot be easily fixed. If we allow [AI] to make autonomous decisions, then we need traceability to check back without a labor-intensive investigation."

– Program Manager, Defense

BUILT-IN, NOT BOLTED ON

Current accountability practices often look back, focusing on what went wrong rather than what needs to go right to prevent the error. To scale, governance must be built into the workflow platform from the very beginning, rather than "bolted on" via administrative checklists. True accountability requires a "control tower" approach: a single pane of glass that provides real-time visibility into every AI-driven action. By unifying legacy systems and AI agents into a single orchestrated flow, agencies can ensure that every autonomous action is recoverable, transparent, and aligned with federal mandates.

89% of agencies require logging and audit trails for all AI actions.

84% utilize documented escalation policies to manage accountability.

79% rely on structured post-incident reviews.

29% have a documented kill switch procedure — even as 79% say they want human-in-the-loop control for the "highest-risk functions."

AGENCY PERSPECTIVE

"We keep logs of things so we can track what [our AI] is doing. There's accountability there. Is it foolproof? No. Is it still going to make a couple mistakes along the way? Nothing is infallible. But as long as it's recoverable, and we can learn from it, those are the key aspects."

– Executive Director, FedCiv

Conclusion

Federal leaders say they want human control over high-risk AI—but less than a third have a kill switch to enforce it, leaving a dangerous gap between intent and capability where trust is won or lost. Closing that gap doesn't mean starting over; it means treating governance as core infrastructure, as essential as networks and data. Agencies must act now to define intervention triggers, ensure data readiness, and unify oversight into a single platform—or risk losing control as systems scale. When failures happen, they cannot be crises; they must be contained, repeatable workflows. Built-in accountability is no longer optional—it is the prerequisite for any agency serious about deploying agentic AI at mission speed.

To learn more visit www.servicenow.com and www.marketconnectionsinc.com