

NASCIO State IT Recognition Awards 2026

Entry Submission



State / Jurisdiction: Washington, DC

Agency: Office of the Chief Technology Officer (OCTO)

Category: State CIO Office Special Recognition

Project Title: Governing AI at City Scale: DC's End-to-End AI Governance Framework

Project Dates: February 2024 – Present

Contact Name: Laura-Celine Patrick

Contact Title: Public Information Officer, OCTO

Contact Email: laura-celine.patrick@dc.gov

Contact Phone: 202-913-3525

Executive Summary

Under the leadership of CTO Stephen N. Miller, Washington, DC has built one of the most structured and accountable AI governance frameworks of any major American city — and the Office of the Chief Technology Officer initiated and led every element. The framework begins with Mayor's Order 2024-028, signed February 2024, which established OCTO as the District's authoritative body for AI governance and created six core AI Values: Clear Benefit to Residents, Safety and Equity, Accountability, Transparency, Sustainability, and Privacy & Cybersecurity. From that mandate, CTO Miller enforced access restrictions on unauthorized AI tools in January 2025, deployed the AI Values Alignment (AIVA) review process across 80+ agencies, launched enterprise-wide mandatory responsible AI training — making DC the first major U.S. city to require it — and rolled out Microsoft Copilot Chat enterprise-wide following a rigorous 314-person, 37-agency pilot. The result is a governance-first AI adoption model that treats rigor and innovation not as competing forces, but as inseparable pillars of responsible public service.

Project Description

IDEA

What problem does this project address?

As AI tools proliferated rapidly across government, most jurisdictions faced a common dilemma: the pressure to adopt AI quickly, with few mechanisms to ensure that adoption was safe, equitable, or accountable. Without a structured governance framework, agencies risked deploying tools that had not been evaluated for bias, privacy vulnerabilities, or alignment with resident interests. At the same time, a workforce adopting AI without grounding in responsible use principles risked compounding those risks at scale. Washington, DC recognized that governance policies alone were insufficient — employees needed to understand why the District's approach was designed the way it was, not just what the rules required.

Why does it matter?

AI in government is not an abstract technology question — it is a question of public trust. When AI systems touch child welfare caseworkers, 311 routing, housing deed review, and resident engagement, the stakes of getting it wrong are real. DC's framework directly addresses three dimensions of that risk: it ensures every AI tool is evaluated before deployment; it ensures every government employee understands their accountability in using AI; and it ensures residents have a voice in how AI is used on their behalf. No other major U.S. city had required all three simultaneously.

What makes it different?

Most AI governance efforts in government are either purely policy (frameworks without enforcement) or purely technical (tool procurement without workforce accountability). DC's framework closes both gaps. CTO Miller's January 2025 decision to restrict all employees from unauthorized AI tools at the system level — not just on paper — was a deliberate signal that governance was real. The mandatory training, completed at zero cost to the District through the InnovateUS partnership, paired that enforcement with education. And the AIVA public advisory body institutionalized resident voice in a way that remains rare in municipal AI programs nationally.

What makes it universal?

Every state and jurisdiction is grappling with AI governance. The core questions DC has answered — how do you evaluate AI tools before deployment, how do you build workforce readiness at scale, and how do you maintain public accountability — are universal. DC's model is explicitly designed to be replicable: the AIVA framework, the training partnership with InnovateUS, and the six AI Values are all transferable to any jurisdiction. This directly aligns with the State CIO Top Ten Priority of digital transformation leadership and responsible technology adoption.

IMPLEMENTATION

What was the roadmap?

The framework was implemented in deliberate, sequenced phases — each building on the last:

- February 2024: Mayor's Order 2024-028 signed. OCTO designated as authoritative body for AI governance. Six AI Values established. AI Taskforce and AIVA Advisory Group created.
- 2024: AI Values Alignment (AIVA) review process developed and deployed. Agencies required to complete structured evaluation before deploying any AI tool.
- Q3–Q4 2024: Three-phase Microsoft Copilot pilot conducted across 314 participants from 37 agencies. Pulse surveys and outcome reviews at each phase informed the enterprise rollout strategy.

- January 2025: CTO Miller restricts all DC Government employees from unauthorized AI tools at the system level — enforcement, not just policy.
- November 2025: Microsoft Copilot Chat deployed enterprise-wide following pilot completion.
- February 12, 2026: Mandatory responsible AI training launched for all DC Government employees and contractors. DC becomes the first major U.S. city to require this.
- May 2026: 15,300+ employees have begun training within 60 days; 48 AIVA requests under active review with goal to publish by June 2026.

Who was involved?

CTO Stephen N. Miller initiated and led the framework from the outset, with AI governance positioned at the cabinet level. Within OCTO, the AI Taskforce — spanning AI, data, cybersecurity, law, and software expertise — guides agencies through the AIVA process and ensures compliance. The AI Values Alignment Advisory Group (AIVA) is a public body that conducts community engagement and advises the Mayor on ethical alignment. OCTO's dedicated AI staff member coordinates implementation across 80+ agencies. The InnovateUS partnership — based at the Governance Lab at Northeastern University — provides the training curriculum, maintained by a nationwide coalition of public sector learning leaders with full editorial independence.

How did you do it?

The framework's implementation relied on three reinforcing mechanisms: policy (the Mayor's Order), enforcement (unauthorized tool restrictions, training deadlines with access consequences), and enablement (free training, enterprise-licensed tools, AIVA support for agencies). The mandatory training was delivered at zero direct cost to the District — InnovateUS's existing public sector curriculum eliminated curriculum development costs, while OCTO's investment focused on customization: a 'Module Zero' tailored to DC's own governance framework and a leadership video from CTO Miller. Employees complete two self-paced modules in under two hours via Single Sign-On credentials.

IMPACT

What did the project make better?

Washington, DC now has a city government in which every AI tool has been evaluated before deployment, every employee has been trained in responsible AI use, and residents have a formal mechanism to shape AI governance through the AIVA Advisory Group. The gap between AI policy and AI practice — which plagues most jurisdictions — has been systematically closed. Agencies that previously had no framework for AI adoption now have a structured review process, published AIVA reports, and assigned responsible officials for every deployed AI system.

How do you know?

The outcomes are documented and measurable:

- Mandatory training: DC became the first major U.S. city to require responsible AI training for its entire government workforce — all employees and contractors — on February 12, 2026.
- Reach: 15,300+ employees and contractors (44% of DC Government's workforce) began training within approximately 60 days of launch.
- Cost: Implemented at zero direct cost to the District.
- Enterprise deployment: Copilot Chat deployed following a three-phase, 314-participant, 37-agency pilot — one of the most rigorous pre-deployment evaluations of any enterprise AI tool in U.S. municipal government.
- AIVA pipeline: 3 AIVA Reports published; 2 additional ready to publish; 48 requests under active review with a goal to publish by June 2026.

- Productivity potential: Pilot data found government employees could potentially gain 2.5 hours of productivity per week using generative AI tools responsibly.
- External recognition: OCTO's AI Taskforce received the inaugural AI 50 Award from the Center for Public Sector AI.

"DC's approach demonstrates that strong governance and bold innovation are not in tension. They are, in fact, inseparable." — CTO Stephen N. Miller

What now?

OCTO is evaluating making AI training a recurring annual requirement — establishing responsible AI literacy as an ongoing pillar of DC's workforce development rather than a one-time compliance exercise. Curriculum expansion is under evaluation to cover emerging topics including agentic AI and agency-specific use cases. OCTO will continue publishing AIVA Reports and supporting all DC Government agencies in authoring their own AI Strategic Plans. Most significantly, this model is designed to be shared: the combination of a no-cost training platform, a governance-first module, a mandatory completion framework, and a public advisory body is a blueprint any city, county, or state could adopt.

Supplemental Information

Additional references and documentation:

- Mayor's Order 2024-028 and DC AI Values: <https://techplan.dc.gov/page/mayors-order-articulating-dc%E2%80%99s-artificial-intelligence-values-and-establishing-artificial>
- DC AI Values and Strategic Plan: <https://techplan.dc.gov/page/dcs-ai-values-and-strategic-plan>
- Mandatory AI Training announcement: <https://octo.dc.gov/release/dc-becomes-first-major-us-city-require-responsible-ai-training-government-workforce>
- InnovateUS Partnership: <https://innovate-us.org/partner/dc>
- AI 50 Award: <https://octo.dc.gov/release/artificial-intelligence-task-force-wins-inaugural-ai-50-award>
- Media coverage — Government Technology: <https://www.govtech.com/artificial-intelligence/residents-want-ai-to-simplify-government-d-c-cto-says> public
- Media coverage — Government Technology: <https://www.govtech.com/artificial-intelligence/washington-d-c-mandates-ai-training-for-staff-contractors>
- Media coverage — Cities Today: <https://cities-today.com/dc-becomes-first-major-us-city-to-require-ai-training/>
- Media coverage — Technical.ly: <https://technical.ly/civics/dc-mandates-generative-ai/>