
The AI Adoption Strategy Gap in Policing: An Unmet Need with Real-World Consequences

WHAT LAW ENFORCEMENT LEADERS SAY AMERICAN POLICING
NEEDS TO GET AI RIGHT

This brief draws on practitioner input gathered at NPI's April 2026 Executive and Leaders Roundtable in San Francisco, co-hosted by Microsoft, combined with research and examples from other sectors and jurisdictions. It makes the case for a national AI adoption strategy for law enforcement and describes what agencies can begin doing now within their current constraints. The order in which the content is organized reflects the order of the Roundtable agenda and is not reflective of the priority or importance of the topics.

Chatham House rules applied throughout the convening. No statements are attributed to named individuals or agencies. The views expressed are those of NPI.

Suggested Citation: National Policing Institute. (2026). The AI Adoption Strategy Gap in Policing: An Unmet Need with Real-World Consequences. Arlington, Virginia.

ABOUT NPI

NPI partners with law enforcement agencies, leaders, and stakeholders to elevate policing so officers and staff are better supported, organizations are higher-performing, and communities are safer. Learn more at www.policinginstitute.org.

Correspondence about this report should be addressed to the National Policing Institute at info@policinginstitute.org.

© 2026 National Policing Institute.

KEY TAKEAWAYS

- **AI adoption in law enforcement is already underway.** Eighty-three percent of agencies represented at NPI's April 2026 roundtable had formally deployed at least one AI tool. Not a single participant reported zero AI presence.
- **Acquisition has outpaced preparation.** Forty-four percent of those same agencies had done no AI-specific training for any personnel. Governance infrastructure, data readiness, and community engagement lag far behind.
- **Decision frameworks for evaluating specific tools are emerging.** What the field lacks is an adoption strategy: a practical plan for building the organizational capacity, workforce literacy, and governance structures (policies and practices) that make responsible use possible.
- **Other sectors and other nations have built these strategies.** American law enforcement has not. Every agency is navigating this alone.
- **This brief makes the case for a national AI adoption strategy for law enforcement** and describes what agencies can begin doing now, within their current constraints, while that strategy is developed.

ABOUT THIS BRIEF AND THE CONVENING

In April 2026, the National Policing Institute convened an invitation-only Executive and Leaders Roundtable in San Francisco, hosted by Microsoft. Participants included chiefs, sheriffs, senior command staff, crime analysts, technology strategy officers, and command-level leaders from 13 agencies ranging from large urban departments to regional sheriffs' offices and small-city agencies, representing a geographic cross-section of primarily West Coast law enforcement leadership. NPI is grateful to Microsoft for their financial support and hosting, and for their sustained commitment to responsible AI adoption in public safety.

The session was a structured listening exercise. NPI did not bring a predetermined policy agenda. The goal was to hear directly from practitioners about where their agencies stand



on AI, what is working, what is not, and what they believe the field needs most. Chatham House rules applied throughout; what follows reflects the themes and substance of the conversation without attributing statements to named individuals or agencies.

This brief draws on that conversation, combined with research and examples from other sectors and jurisdictions that have moved further along the preparation curve. It is not a technology review or a vendor guide. For purposes of this brief, 'AI' refers to generative AI, agentic AI, and the machine learning tools that underpin them -- a scope that reflects both the current wave of accessible AI entering agencies and the near-term horizon where the most consequential adoption decisions will be made. This brief and the convening effort is intended to make an argument, grounded in practitioner input, for why law enforcement needs a national AI adoption strategy and what that strategy must address.

ADOPTION IS HAPPENING. PREPARATION HAS NOT.

AI adoption in law enforcement is not a future question. It is a present reality. Across the public sector broadly, the conversation has already shifted from cautious piloting to active implementation: 55% of public sector leaders now report using agentic AI in their organizations (Google Cloud & National Research Group, 2026, as cited in Teale, 2026). Law enforcement is not behind that curve. Every participant at NPI's roundtable had some form of AI presence, and 83% had formally deployed at least one tool.

Preparation has not kept pace. As NPI's president framed it at the opening of the roundtable: "What's missing is not prohibition, not regulation against the use of AI, but a responsible framework for how to prepare your agencies to better leverage AI for more efficient and effective public safety. We're not talking about where to start necessarily. We're talking about how do you prepare your governance processes, your accountability processes, your staff, and how do you engage the community early and consistently in this process."

"Leadership needs to get on board now because our officers and deputies are using it. They may not be using it per policy. They may not be using it with anybody's approval. But they are using it, and it is going to be sooner or later before it bites us."



The risk is real. Without governance, training, and accountability structures in place, agencies face real exposure: Freedom of Information Act (FOIA) requests for AI-generated content, tools that fail publicly, labor grievances, and political restrictions imposed by oversight bodies that have grown skeptical of technology they do not understand. The agencies most likely to lose access to AI are not the ones that adopted it. They are the ones that adopted it without preparation.

The field has no shortage of guidance on what AI can do. What it lacks is a clear adoption strategy. Many existing strategies begin by identifying problems AI might solve or capabilities agencies may want to acquire. **NPI's view is different: policing first needs a strategy for building its capacity to adopt AI responsibly. That means strengthening organizational readiness—governance, workforce capability, data infrastructure, and community trust—before focusing on specific tools and solutions.** Decision frameworks for evaluating those tools and other resources are emerging. For example, the Council on Criminal Justice Task Force on Artificial Intelligence published a user decision framework in 2026 covering evaluation, procurement, implementation, and ongoing monitoring (CCJ, 2026). The Future Policing Institute's Center for AI has published AI adoption principles and a variety of other resources for agencies with the input of dozens of law enforcement leaders across the nation (Future Policing Institute, 2026). These are valuable contributions for agencies that are ready, but an adoption strategy addresses how agencies become ready in the first place.

THE BENEFITS ARE REAL AND WITHIN REACH

Participants at the San Francisco roundtable who have deployed AI tools are more confident in the technology, not less. The use cases generating the most interest were not speculative: report writing efficiency, crime analysis and data mining, staffing and scheduling support, early warning indicators for officer wellness, automated compliance checks in reporting workflows. These are areas where the technology is mature enough to deliver real value now.

The most striking example from the room came from an agency that has built substantial in-house AI and machine learning capacity over several years. A participant described the shift this way: "No longer are we starting to rely on products. We're building these ourselves for a fraction of the cost and a fraction of the time." That same agency had a deputy build a transcription application for investigative interviews over two weekends that



is now in active use. Annual cost projections for the system: under \$100,000. The application was built on existing infrastructure, using AI-assisted development.

“One of the biggest seismic shifts I’m seeing is that no longer are we starting to rely on products. We’re building these ourselves for a fraction of the cost and a fraction of the time.”

Agencies that invest in internal technical capacity are increasingly able to build tools faster, cheaper, and better-fitted to their operations than what the vendor market delivers. But this capacity does not appear by accident. It is built through sustained investment in people, data infrastructure, and governance. That is the preparation work most agencies have not yet done.

Realizing the benefits of AI is contingent on that preparation. The agencies capturing value built the organizational foundation first or built it alongside deployment. The field needs a roadmap for doing that work systematically, not agency by agency, from scratch.

YOU CAN’T JUST TURN IT ON: WHAT PREPARATION REQUIRES

The five dimensions of preparation described here are not a deployment checklist. They are the organizational foundation that agencies need before they can responsibly make decisions about which tools to acquire, how to govern their use, and how to sustain accountability over time. Decision frameworks address what to do with a specific tool. Preparation addresses what an agency needs to be ready to use any tool responsibly.

DATA: THE FOUNDATION EVERYTHING ELSE DEPENDS ON

“AI is only as good as the data behind it. If your data foundation isn’t ready, AI won’t just underperform — it might introduce real operational and reputational risk” (Marcolis, 2026). Participants at the San Francisco roundtable encountered this directly. One described a National Incident-Based Reporting System (NIBRS) implementation where reporting requirements conflicted with academy training, producing data inconsistencies that



compounded over time: “What they’re taught in the academy doesn’t match what NIBRS is. So your data is not going to be there, not going to be accurate, to be able to pull all that.”

A second example carried the same lesson. When one agency examined the data behind a vendor’s success claims, the analysis revealed a fundamental flaw in the vendor’s methodology. As a participant described it: “When you don’t know what you’re looking for, or how they’re determining success, you’re gonna be fooled. You can’t defer to somebody who doesn’t have a background in law enforcement.” The data was being used to demonstrate effectiveness. The conclusions were wrong.

What data readiness requires:

- **Ownership:** Know who controls your data, what your vendor contracts say about data rights, and what happens to your data if you change providers. One participant from a well-resourced agency offered clear guidance: “Don’t give your data to anyone. Have the ownership of the data but deliver access to that data using appropriate infrastructure and security measures.”
- **Governance:** Documented policies and standards for how data is collected, maintained, and used, with accountability when those standards are not met.
- **Quality checks:** Automated validation that surfaces problems before they become AI output problems.
- **Assessment:** An honest agency-level review of data maturity before any tool is acquired or deployed. Assessment before automation is not caution. It is competence.

GOVERNANCE: STRUCTURE BEFORE DEPLOYMENT

Sixty-one percent of agencies at the San Francisco roundtable had a designated person or committee for AI governance. But governance is more than an org chart entry. It is a set of structures that determine who approves which use cases, what the audit trail requires, and what happens when an AI output is wrong.

Risk tiering is a practical starting point: classifying AI use cases by potential for harm before deploying them. Georgia requires state agencies to categorize every use case as low, medium, or high risk before moving forward (Kinkade & Knell, 2026). Law enforcement has no equivalent standard. The Council on Criminal Justice (CCJ) user decision framework establishes a five-phase process, from readiness assessment through ongoing monitoring, that provides a replicable approach to tool-level governance (CCJ, 2026). But that process requires an agency to have governance infrastructure in place before it begins. That infrastructure is what an adoption strategy builds.



Legal exposure arrived in the room as a concrete example, not a hypothetical. As recounted in the roundtable discussion, a sheriff's office communications officer in a southeastern state had AI content released under a records act request because no policy governed its use. The content was taken out of context. The staff member had done nothing wrong. But the agency had no framework to defend against the scrutiny, because it had no policy to point to.

"NIST has a framework for risk, and I think police agencies need to get to the same place. We need a use-of-force continuum for AI risk, and that needs to be in our policies. Translation services aren't the risk that facial recognition is. Putting that on a scale matters."

To NPI's knowledge, no American law enforcement agency has developed a proven, replicable accountability framework for AI that can serve as a field-wide standard. The UK's National Police Chiefs' Council has published a governing covenant with six principles every force must follow: lawful, transparent, explainable, responsible, accountable, and robust. American law enforcement has no equivalent.

WORKFORCE: THE MOST IMMEDIATE GAP

Forty-four percent of agencies at the San Francisco roundtable had done no AI-specific training for any personnel. At the same time, officers and deputies are using AI on personal devices or accounts, writing with it, and in some cases acting on AI-generated information, with no policy framework and no training to guide responsible use.

Leaders are systematically overestimating workforce readiness. A Boston Consulting Group (BCG) survey found that 76% of executives believed their employees were enthusiastic about AI adoption; the actual figure among employees was 31% (Zaki, 2026). That gap is not only about morale. When AI is introduced without genuine communication and involvement, adoption breaks down. Research found that nearly a third of employees, and 44% of Gen Z workers, admitted to actively opposing or sabotaging their organization's AI strategies (Zaki, 2026). In a policing context, that means tools deployed at significant expense, but not used or used incorrectly, diminish return on investment.

Supervisors face a specific challenge that goes beyond basic training. AI accelerates the pace at which work is produced in ways that traditional management structures are not designed to handle. Research on workplace AI adoption describes a growing bottleneck:



managers overwhelmed by the volume and pace of AI-assisted output, without corresponding changes in how review and oversight are structured (Fosslien & West Duffy, 2026). Police supervisors reviewing AI-assisted reports, analyses, and communications will face exactly this pressure, and most are not yet equipped for it. In other words, AI may accelerate the work in one task or domain, but other parts of the system are not ready for the increased capacity.

Workforce preparation must be sequenced deliberately:

- **Procurement staff first.** The Department of Defense, when it developed its AI adoption strategy, made procurement training the first priority. As NPI's president noted at the roundtable: "That's the front door, and the front door has to be educated and brought on board." Every vendor claim, every data rights clause, every performance standard gets negotiated at the contract stage. Agencies that do not have AI-literate procurement staff are exposed from the first signature.
- **Leadership and command staff second.** Command staff who cannot evaluate vendor claims, assess AI outputs, or answer basic questions from oversight bodies are not positioned to lead responsible adoption or to defend it when challenged.
- **Supervisors third.** They are the accountability checkpoint for AI-assisted work product, and they need the skills to know what to look for and when to push back.
- **Frontline personnel throughout,** with emphasis on responsible use, verification habits, and knowing when not to trust the output.

One agency at the table had navigated this well by starting with the workforce population most burdened by administrative load. As a participant described: "We looked at AI not just as a product. We looked at AI as a strategy in transforming the workforce. We targeted the individuals that were really hungry to relieve the burden, and then moved up from there." The labor union, which had scheduled meetings to push back on the initiative, stopped showing up because members saw the tools relieving burden rather than threatening jobs.

"Our frontline employees are going to know where the gaps are and where the efficiencies are. We're creating a framework with a clear policy and turning the working cops loose: what are the efficiencies you're looking for in your workspace? If we trust the old guard to predict what the up-and-coming team is looking for, we're always going to be behind the curve."



ACCOUNTABILITY: DESIGNING WHAT DOES NOT YET EXIST

“Human in the loop” is a design question, not a policy declaration. Declaring that human oversight exists is not the same as designing where in the workflow that oversight sits, who is responsible when an AI output is wrong, and how that review is documented. Agencies that do not design this deliberately will find that oversight becomes nominal as familiarity with AI outputs grows and the impulse to review carefully fades.

One participant described the pattern directly: tools that moved through the system generating outputs that were accepted without scrutiny, not because anyone chose to abdicate review, but because the culture had not been built to demand it. The failure was not technological. It was organizational.

Globally, 57% of public sector agencies have multiple AI proof-of-concept projects running. Only 12% have moved those projects to production systems (Microsoft, shared by a company participant at the San Francisco roundtable). The gap reflects what happens when agencies adopt tools without the governance infrastructure to move from experimentation to accountable deployment. A good adoption strategy addresses that transition explicitly, before the tool is turned on.

COMMUNITY AND STAKEHOLDERS: LEGITIMACY IS NOT A CHECKBOX

Policing authority rests on community trust. AI adoption that bypasses community engagement puts that foundation at risk. Several participants at the San Francisco roundtable were already managing the consequences of decisions that moved ahead of community awareness, dealing with skeptical oversight bodies, alarmed communities, and elected officials who responded by imposing restrictions.

The agencies that have navigated this most successfully did it through early engagement and concrete explanations of what the technology does not do. As one participant from a city with significant community AI skepticism described: “We’ve done a lot of policy work, a lot of work with our Privacy Commission and City Council to get everybody kind of understanding what the technology actually does, so it’s not as scary when we bring this stuff forward.”



“We’re in the middle of Silicon Valley. We see all this technology. We see how it can enhance what we do. But yet our elected officials just don’t see it the same way. How do we get to a point where we can educate not only the elected officials but the community at large?”

Research on organizational AI adoption identifies a consistent pattern: people are more likely to accept new policies and technologies when decision-makers communicate clearly and listen to those affected, what researchers describe as “procedural justice” around new technologies (Zaki, 2026). This is not a communications tactic. It is the same procedural justice principle that shapes how communities evaluate policing more broadly. It applies here.

Community co-design of AI tools, proactive public communication strategies, and structured engagement with the populations most directly affected by AI-assisted policing decisions represent gaps in preparedness. A national adoption strategy needs to address them directly.

WHAT OTHERS ARE DOING

Other sectors and other nations are not waiting for a national standard to emerge. They are building the organizational preparation infrastructure now, and the lessons apply directly to law enforcement in the U.S.

GEORGIA: TEST BEFORE YOU DEPLOY

Georgia Technology Authority opened a dedicated Innovation Lab as a controlled environment for state agencies to test AI use cases before deploying them to production. The lab requires ethics training before any solution reaches that production stage. The operating philosophy, as described by the state’s CIO, is explicit: like baking a cake, you test, tweak, and perfect in a controlled setting before serving anything to the public (Silver, 2026). The lab is expanding because demand from cities, universities, transit authorities, and state agencies has exceeded its original capacity.

Georgia’s CIO identified what makes the difference: “AI literacy is the most critical characteristic of AI governance. When they’re able to use the tools responsibly, then we don’t have to worry about issues coming out” (Kinkade & Knell, 2026). The state classifies



every AI use case as low, medium, or high risk before moving forward, requires staff to use only the data necessary for each specific application, and makes AI governance everyone's responsibility, not just the technology team's responsibility.

THE UNITED KINGDOM: NATIONAL INFRASTRUCTURE FOR LOCAL ADOPTION

The United Kingdom launched PoliceAI on June 10, 2026, a national center hosted by the College of Policing and backed by 75 million pounds (approximately \$95 million U.S.) over three years, dedicated to the responsible development, testing, and scaling of AI across all 43 forces in England and Wales. In its first year, PoliceAI is running large-scale pilots in up to ten forces focused on digital evidence triage, disclosure, and summarization, one of the most time-consuming parts of any investigation, with plans to scale to all forces by 2027 and a projection that the technology will free up the equivalent of 3,000 officers annually. The center will publish a public registry of AI tools in use across policing, and all AI models are independently tested for accuracy and bias before deployment, a standard the Policing Minister described as essential to earning public consent at every step (UK Home Office, 2026; Office of the Police Chief Scientific Adviser (OPCSA), 2026). American law enforcement has no equivalent at any level.

SAN DIEGO COUNTY SHERIFF'S OFFICE (SDSO): IN-HOUSE SOLUTION BUILDING

The San Diego County Sheriff's Office did not acquire its way to AI capability -- it built that capability over more than a decade through deliberate organizational investment. Starting around 2019, the agency hired a dedicated AI and machine learning engineer, partnered with the UC San Diego Data Science Institute for technical talent, and continued to grow their internal technology team that now includes approximately 80 individuals. That investment made possible results that the vendor market could not match at comparable cost: a license plate reader built by an intern and deployed on existing patrol vehicle hardware; a jail interview transcription system built by a deputy over two weekends, with projected annual costs under \$100,000. The agency also approaches data governance and vendor contracting with unusual discipline -- retaining ownership of all agency data while enabling regional sharing through secure API access and full audit trails, and building data rights, API access, and performance milestones into every vendor contract from the start. The SDSO story is not a blueprint most agencies can replicate immediately. It is evidence that law enforcement agencies willing to invest in organizational capacity and internal technical infrastructure can realize many benefits -- and that responsible adoption starts with building the foundation, not buying the tool.



WHAT LAW ENFORCEMENT LEADERS CAN DO NOW

Many agencies face real constraints on AI adoption: city or county restrictions, state legislation limiting specific technologies, oversight body moratoriums, union agreements governing technology introduction, and political environments where any visible AI use carries risk. A credible adoption strategy acknowledges these realities. None of them, however, prevent the preparation work. Agencies that can document their governance processes, training programs, and community engagement are in a stronger position with skeptical oversight bodies than those that cannot.

Here is what agencies can do now, within most existing constraints:

- **Conduct an honest internal assessment.** Before acquiring anything new, understand what is already in use, authorized or not. Several participants at the roundtable described discovering AI already embedded in workflows their agencies had not formally sanctioned. Understand what AI systems have access to agency data and what is being done with it. Know what policies currently govern those uses, and where the gaps are.
- **Address data before deployment.** Identify where data quality problems exist. Review current vendor contracts for data ownership provisions. Understand what data your AI tools are using and how it is governed. The cost of not doing this shows up later, publicly.
- **Build leadership and procurement literacy now.** Ensure basic knowledge and understanding of AI frameworks and how the underlying technology works. Command staff who cannot evaluate a vendor's claims, assess an AI output for reliability, or answer basic questions from an oversight body are not equipped to lead responsible adoption. Procurement officials who represent the "front door" must be equally capable. That literacy can be built without deploying a single new tool.
- **Know your legal and labor landscape.** Understand what your jurisdiction currently permits. Engage legal counsel before contracts are signed, not after. Know what your labor agreements require before introducing new technology and develop a workforce communication strategy that addresses displacement concerns directly and honestly.



- **Start with low-risk, internal use cases.** As the roundtable discussion made clear, agencies that have navigated AI adoption most successfully started with administrative and internal applications such as documentation support, compliance checks, and scheduling, where risk is manageable and benefits are visible to the workforce. The safer end of the pool is the right entry point: demonstrate proficiency and accountability before moving to higher-stakes applications.
- **Engage the community before it becomes required.** The agencies that have navigated oversight bodies successfully built relationships before controversies arose. Proactive transparency is a much stronger position than reactive explanation.

THE CASE FOR A NATIONAL ADOPTION STRATEGY

Every agency in the room in San Francisco is solving the same set of problems independently: how to assess data quality, how to categorize AI risk, how to train procurement staff, how to structure human oversight, how to engage community stakeholders. There are no shared standards, no common frameworks, and no replicable models. The solutions developed at one agency cannot be easily transferred to another. The failures at one agency do not prevent the same failure at the next.

That is what a national AI adoption strategy for law enforcement would change. Not a federal mandate. Not a vendor playbook. A field-developed, practitioner-grounded strategy developed independent of preferences for agency type, size, location, or function, that gives agency leaders a practical framework for assessing where they stand, building what they need, and governing responsibly, developed by people who understand operational policing and have heard directly from those navigating it in real time.

Other sectors have built these strategies. Healthcare has frameworks for AI implementation that address clinical safety and organizational readiness together. The Department of Defense has an AI adoption strategy that begins with workforce training before deployment. State government in Georgia has built a test-first approach that is expanding because it works. Policing has not built a comparable strategy, not because the problem is less serious, but because no organization has taken ownership of it.



NPI is making the case that this work needs to happen, and it needs to happen now, before more agencies deploy tools without the organizational preparation to govern them responsibly, and before the failures that result from that gap produce political and public backlash that sets back responsible adoption for years.

For law enforcement leaders: your participation in convenings like this one, your willingness to share what is working and what is not, and your demand for field-tested practical guidance are what make a credible strategy possible. NPI will continue to convene and to listen.

FOR PHILANTHROPY AND CIVIC SECTOR LEADERS

Tools without strategy produce risk. The investment that makes every other investment in public safety technology more likely to succeed is the one that builds the field's capacity to adopt responsibly. That is the work NPI is organized to support.

REFERENCES

Council on Criminal Justice (CCJ). (2026). Assessing AI for criminal justice: A user decision framework. Council on Criminal Justice. <https://counciloncj.org/assessing-ai-for-criminal-justice-a-user-decision-framework/>

Fosslien, L., & West Duffy, M. (2026, May 25). Managers are struggling to keep up with the AI productivity boom. Harvard Business Review. <https://hbr.org/2026/05/managers-are-struggling-to-keep-up-with-the-ai-productivity-boom>

Future Policing Institute. (2026). Center on AI. <https://www.futurepolicing.org/center-on-ai>

Kinkade, L., & Knell, N. (2026, May 22). 'AI is a team sport': Georgia upskills for AI, tests agents. Government Technology. <https://www.govtech.com/artificial-intelligence/ai-is-a-team-sport-georgia-upskills-for-ai-tests-agents>

Marcolis, C. (2026, February 16). How state and local agencies can build AI-ready data foundations in 2026. StateTech Magazine. <https://statetechmagazine.com/article/2026/02/how-state-and-local-agencies-can-build-ai-ready-data-foundations-2026>



National Police Chiefs' Council (NPCC). (2023). Covenant for using artificial intelligence (AI) in policing (Version 1.1). National Police Chiefs' Council.

https://science.police.uk/site/assets/files/4682/ai_principles_1_1_1.pdf

Silver, A. (2025, July 18). Georgia's AI lab cooks up tech ideas in new test kitchen.

Government Technology. <https://www.govtech.com/artificial-intelligence/georgias-ai-lab-cooks-up-tech-ideas-in-new-test-kitchen>

Teale, C. (2026, April 30). Government leaders see 'momentum' in agencies' AI adoption.

Route Fifty. <https://www.route-fifty.com/tech-data/2026/04/government-leaders-see-momentum-agencies-ai-adoption>

Jones, S. (2026, June 10). PoliceAI to speed up investigations and fight crime - GOV.UK.

<https://www.gov.uk/government/news/policeai-to-speed-up-investigations-and-fight-crime>

Zaki, J. (2026, April 30). Empathetic leadership can make or break AI adoption. Harvard

Business Review. <https://hbr.org/2026/04/empathetic-leadership-can-make-or-break-ai-adoption>

This brief reflects proceedings from the NPI Executive and Leaders Roundtable convened in San Francisco on April 16, 2026, and is intended as a practitioner-facing policy document. Chatham House rules apply; no statements are attributed to named individuals or agencies. The views expressed are those of NPI.