

Updating Use Guidance for Generative Artificial Intelligence

By Hande Güven and Jessica Wang

November 2025

This policy project was developed by Aspen Policy Academy fellows while participating in the Science and Technology Policy Fellowship. They were asked to recommend updates to the New York City Office of Technology and Innovation's [Artificial Intelligence Action Plan](#). This policy brief provides an overview of the fellows' proposed solution. [The full project is available here](#). Please note that all authors' opinions published here are their own. This publication does not reflect the views of the Aspen Policy Academy or the Aspen Institute.

Overview

In connection with New York City's Office of Technology and Innovation (OTI), **this project recommends that OTI update its [Preliminary Use Guidance for Generative Artificial Intelligence](#) by introducing an action-oriented “guiding questions” framework grounded in the City's Artificial Intelligence (AI) Principles.** Updating the Preliminary Use Guidance, which was released as part of the AI Action Plan to support New York City agencies in their exploration of generative AI (GenAI) tools, would support broader adoption of the guidance by city agencies. Released in 2023, New York City's AI Action Plan was the nation's first comprehensive plan for the responsible use of AI in municipal government.

This update would enhance the utility of the GenAI Use Guidance by more concretely addressing emerging risks, incorporating lessons learned from early adoption of GenAI tools, and improving usability for agency staff across a range of technical backgrounds. It would provide a clear and concise framework built on guiding questions for each of the [5 AI Principles the City enumerated by the City in 2023](#), thereby creating a foundation for New York City's AI governance efforts. These questions would help readers by first inviting them to ask key questions about their GenAI use, then explaining why each issue matters, and, finally, providing practical mitigation recommendations alongside real-world use cases to illustrate safe and effective use of GenAI.

We also recommend improving the usability and readability of the guidance by making easy-to-implement changes, such as structuring the document to move from high-level, executive-oriented content to practical and actionable advice illustrated with end user-focused use cases; adding a concise Quickstart Guide explaining how best to consume the document; and using visual cues to guide the user through the content.

Background

As GenAI tools become more widespread and accessible, cities must establish clear, practical guidance to ensure their safe, reliable, and responsible use. As part of this effort, New York City has already taken an important first step with the release of its [Preliminary Use Guidance for Generative AI](#) in March 2024. With an update to this guidance scheduled to be completed in 2025, OTI has a valuable opportunity to expand the document's scope to address emerging issues and lessons learned since the release of the preliminary version while providing practical guidance and improving the usability of the guidance.

Other US cities have also been working to provide guidance for city agencies' GenAI use. For example, San Francisco released [Generative AI Guidelines](#) in July 2025, and [San Jose's Generative AI Guidelines](#) were last updated in March 2024. These guidelines share a number of features that make the content easier to consume and more user-friendly, such as the use of risk categories to classify different use cases and provide corresponding mitigation measures and guidance. Similarly, [California created an online GenAI hub](#) that offers guidance and resources presented in an accessible, interactive format. The platform includes a dynamic "choose your GenAI journey" experience that helps agencies assess their experience with AI and determine where they are in the GenAI adoption process.

In addition to reviewing publicly available guidance from other cities, our conversations with policymakers, researchers, and city staff have provided valuable insights into common pitfalls when developing and rolling out GenAI use guidance for a broad audience. One of the most frequent challenges identified by these users of GenAI tools is ensuring that guidance regarding their use is accessible to individuals with different levels of technical expertise while keeping it as clear, actionable, and practical as possible. Another key theme that emerged from these conversations is the need for regular, ongoing updates to the guidance, given the rapid pace of technological change. Equally important is establishing clear feedback channels to foster trust and collaboration between city agencies and the authoring department, such as OTI.

Recommendations

As OTI updates its GenAI Use Guidance, it has a valuable opportunity to clarify its strategic approach and leverage its expertise to help city agencies navigate the rapidly evolving and often poorly understood landscape of GenAI. To that end, the guidance should reflect the agency's approach to GenAI use and technological innovation by offering guiding questions about issues such as the user's attention to cybersecurity, trust and transparency, and reliability issues, thereby engaging users in critical conversations rather than giving unnecessarily constraining guidance. This can be accomplished by focusing on four key objectives:

1. First, the updates should clarify risks and offer actionable mitigation recommendations while building on the foundation of New York City's AI Principles. The guidance should clearly map the risks that GenAI poses for [New York City's AI Principles](#), providing a concise framework that city agencies can use to assess new or proposed use cases. Rather than being overly prescriptive, the guidance should provide a holistic set of questions for readers to ask themselves, then offer examples of risks (e.g., hallucination, bias, privacy risks) and corresponding mitigation strategies (e.g., human review, transparency mechanisms, reviewing data use policies). This approach would encourage consistent application of AI principles across city agencies while allowing flexibility to adapt to specific contexts. A strong link to existing values and frameworks would reinforce trust in the guidance and align it with the City's broader digital policy goals.

2. Second, the updates should offer guidance with varying levels of detail and technical depth for different audiences. The audience for the GenAI Use Guidance document has varying levels of technical fluency and responsibilities, so the guidance should be layered accordingly. For example, the guidance could begin with high-level framing and decisionmaking considerations for executives and policy leads and then transition to specific technical constraints and example use cases for practitioners. This layered structure would ensure relevance across roles without overwhelming users who need only the big-picture view. Tailoring the depth of guidance in this way would support more effective and widespread adoption.

3. Third, the updates should improve usability and adoption by making low-lift changes to the structure and format of the guidance document. Structuring the document to progress from high-level strategic guidance to detailed examples – and placing use cases in an appendix – will help users quickly find relevant content. A concise Quickstart Guide,

placed at the beginning of the guidance or linked from OTI's website, would support time-constrained users by summarizing key takeaways and decisionmaking steps. In either the upcoming iteration or in future versions, adding visual aids like flowcharts or checklists would also improve navigation and comprehension. These low-effort, high-value changes would make the guidance more actionable, approachable, and likely to be adopted by a wider audience across city agencies.

4. Finally, the updates should establish clear feedback channels to support continuous improvement and include a plan for regular updates to keep pace with the evolving technology. To remain relevant and effective, the guidance should be treated as a living document, supported by regular updates and clear mechanisms for user feedback. Agencies should be encouraged to share challenges, successes, and edge case scenarios through designated channels, such as an online submission form or quarterly working sessions. OTI could reinforce this approach by committing to revisit and revise the guidance at consistent intervals, incorporating lessons learned and responding to the evolving GenAI landscape. This participatory model would foster a sense of shared ownership among agencies and help ensure that the guidance remains practical and timely.

If you'd like to learn more, see the full project, including Best Practices for GenAI Use Guidance Rollout and Draft Use Guidance for GenAI, at aspenpolicyacademy.org/project/nyc-genai-guidance-2025.



About the Aspen Policy Academy

The Aspen Institute's Policy Academy helps community leaders and experts across the political spectrum elevate their voices, influence key decisions, and strengthen democracy from the ground up. Our innovative training programs and resources equip people across sectors – from tech to the environment, science to civic engagement – with the skills to shape critical policy efforts. Learn more at aspenpolicyacademy.org.