

Written testimony from BetaNYC for the New York City Council, Committee of the Whole



Oversight: Examining the Risks Posed by Artificial Intelligence, and related bills

Hearing: Monday, October 5, 2026, City Hall

Submitted by: Noel Hidalgo, Executive Director, BetaNYC

Submitted: October 8, 2026

Online version: beta.nyc/council-ai-hearing-2026

Contents

- Written testimony
- What we are asking for
- Suggestions on the bills
 - Across the package
 - Bill by bill
- Appendix
 - Appendix A. BetaNYC: who we are, our tools, our classes, and what we heard
 - About BetaNYC
 - BetaNYC's open-source MCP servers
 - More BetaNYC public repositories for civic data and AI
 - Open data and AI classes BetaNYC teaches
 - Questions we sent to the Council
 - What New Yorkers told us
 - Appendix B. Digital sovereignty, open source offices, NYC AI law, and COPIC
 - B1. Digital sovereignty: BetaNYC's argument
 - B2. Open source program offices
 - B3. NYC's AI laws and policy, at a glance
 - B4. COPIC: the City's information commission
 - Appendix C. Our research, workforce findings, and digital equity
 - C1. Our NYC.gov access test
 - C2. Our AI knowledge test
 - C3. Workforce research
 - C4. Digital equity and digital literacy
 - Appendix D. Questions we sent to the Council
 - Part A. Questions for the Mayor's office (OTI, NYC Cyber Command, OMB)
 - Part B. Questions for AI companies (Anthropic, OpenAI, Google, Meta, SpaceXAI)

- Disclosures

- Written testimony

Good day, Speaker Menin and Members of the Council. I'm Noel Hidalgo, Executive Director of BetaNYC. For the record, I serve as the Public Advocate's appointee to COPIC and to the Internet Advisory Board; I speak today on behalf of BetaNYC, a partner project of the Fund for the City of New York.

BetaNYC is a civic technology organization. We use AI and we build with it, and we wrote our own [AI policy](#) before we asked anyone else to follow one. It starts with one idea: AI is a question of power. Who has the power to consume the world's knowledge and resources, and who controls the companies that do? These tools are built on unseen human labor and creativity. They lay bare hard questions about justice, government, education, health, the economy, and the environment.

For 18 years, our community has asked, "What agency do we have over technology in government? How can we leverage technology to create transparency in operations? Can we regulate automated decision-making systems so they are safe, appropriate, and consensual?" Now we add: will we be able to tell the difference between a synthetic and a human government, and who will be accountable?

We now face global crises of misinformation, climate change, affordability, and energy. AI touches all of them.

Today, we are living in wild times. BetaNYC believes frontier AI models rest on the uncompensated work of writers, journalists, and creators. That view is shared inside the industry: a Microsoft director of applied science called AI training "the largest theft of labor in human history" ([NYT v. Microsoft and OpenAI](#), records unsealed Sept. 17, 2026; [Washington Post](#)). These tools are changing how we work with the machines around us. We are at the beginning of a new era for humanity, if we don't destroy the planet in the process.

I want to focus my testimony on AI in New York City government and how we can move forward. Our public interest technology research shows four things.

First, AI models are snapshots in time. We asked five AI assistants about this Council's own AI law, [Local Law 188 of 2025](#). Of the three that answered from memory, none knew it (see Appendix C2). Out of the box, AI models are brains with no eyes and no hands.

Second, access to City information is inconsistent. NYC.gov's [one published rule](#) welcomes every automated visitor, yet an honest automated tool is refused by a rule no one has published (see Appendix C1). So we built open tools that go straight to the source.

Third, people inside government are ready to experiment, if they can do it safely, reliably, and without learning a whole new system. They lack the literacy and the resources.

Fourth, over 10 years, we have trained nearly 100 CUNY undergraduates, and we

recently interviewed 26 government and civil society experts across 18 organizations about public interest technology careers (see Appendix C3). Employers want more than technical skills; students want practical experience.

This City needs the Council to bridge the gaps in leadership, literacy, and mentorship. A century ago, New York reformers in this chamber helped bring about the end of nepotism in government, clean water, building safety codes, the weekend, universal suffrage, and human rights.

A hundred years from now, how will our descendants look back on today? Did we do all that we could to save democracy and save the planet?

In this spirit, we ask the Council to do great things. First and foremost, work with the State on an AI wealth fund that sustains public universities, public education, and enforcement. Revive [COPIC](#), so public information has real oversight (see Appendix B4). Define New York's vision of digital sovereignty, independent of any model provider, corporation, or administration, starting with the [UN Open Source Principles](#). Increase resources for two existing teams, the Office of Algorithmic Accountability and the Office of Data Strategy, and create a third: an Open Source Programs Office (see Appendix B2). These three teams should partner with CUNY to build a public interest technology institute that develops the workforce and research for this century and beyond.

Digital sovereignty is how we ensure this democracy moves forward and stays accountable to people, not to oligarchs or "market forces." Digital sovereignty starts with information that is accurate, auditable, and ours.

Our suggestions on each bill on today's agenda are in a separate document, "Bill Suggestions," with a supporting appendix.

Thank you.
Noel Hidalgo

What we are asking for

Digital sovereignty starts with information that is accurate, auditable, and ours.

Paying for it

- **An AI wealth fund.** Work with the State on an AI wealth fund that sustains public universities, public education, and enforcement. ([written testimony](#))

Public AI and digital sovereignty

- **Digital sovereignty.** Define New York's vision, so the City does not rent its laws, data, or

AI from any one provider, as it refused patented hydrants and valves. ([our COGE testimony](#))

- **Public AI.** Commit to building public AI: open models, public compute, and public accountability. ([Appendix B1](#))
- **City laws on City infrastructure.** Publish the Charter, laws, and rules on City-owned infrastructure, free and machine-readable. ([our COGE testimony](#))
- **UN Open Source Principles.** Endorse them, a first step that costs nothing. ([Appendix B2](#))
- **Digital rights.** Renew the City's commitment to the Cities Coalition for Digital Rights, which New York co-founded in 2018. ([Appendix B2](#))

Put residents first: privacy, service design, 311 and NYC.gov

- **Privacy by design.** Make privacy-by-design, including children's data, a condition of every City AI system and vendor contract. ([our 8 ideas, #6](#))
- **Service design in every agency.** Embed human-centered service design teams in the agencies that will use AI. ([our 8 ideas, #2](#))
- **311 as the front door.** Invest in 311, with a public way to report and fix a wrong AI answer. ([our 8 ideas, #3](#))
- **NYC.gov open to honest tools.** Require a published rule, with a named owner, for automated and AI access to NYC.gov. ([Appendix C1](#))
- **Build upon open source tools and show the source.** Build open source AI tools and require them to draw from the City's own published laws and data, and cite them. ([Appendix C1](#))
- **No NDA for AI.** The city nor agencies should be able to hide behind non-disclosure agreements that they are required to sign for use of AI tools. We need absolute accountability and transparency.

Scale up City staff with the AI tools already in use

- **AI tools for City staff.** Scale up City staff with GitHub Copilot and Microsoft Copilot, with training so they use them safely and well. Note, these are existing contracts with the city, not that we give preference to these vendors. ([Appendix D, Part A](#))
- **Build skill inside government.** Grow in-house expertise and a safe place to experiment, so the City is not locked into consultants or one vendor. ([Appendix B1](#))

Staff up the offices that do this work

- **Office of Algorithmic Accountability.** Fund and staff it, write its standards around independent testing, and publish its list of assessed AI. \$1 million is not enough to provide the infrastructure to support algorithmic accountability in New York City. ([Appendix B3](#))
- **Office of Data Strategy.** Give it a larger team focused on data quality and data engineering. ([written testimony](#))
- **Open Source Programs Office.** Create one, to coordinate how the City uses, builds, buys, and shares open source. ([Appendix B2](#))

Oversight of public information

- **COPIC.** Revive and fund the Commission on Public Information and Communication, as we have asked since 2018. ([Appendix B4](#); [our 2019 testimony](#))

Workforce, literacy, and research

- **A public interest technology institute with CUNY.** Partner these three offices with CUNY to build the workforce and research the City needs. ([Appendix C3](#))
 - **AI literacy that lasts.** Fund public AI and digital literacy as a standing line, linking Local Laws 188 and 153. ([Appendix C4](#))
-

Suggestions on the bills

BetaNYC's suggestions on the proposed legislation and our written testimony to the New York City Council [Committee of the Whole, October 5, 2026](#).

Generally, BetaNYC supports the bills on today's agenda, and we believe each needs further review. Here are our suggestions on the proposed legislation. We read every item on today's agenda. We cite the eight new bills by their T-numbers, since they have not yet been introduced.

Across the package

1. **Publish it all as open data.** Each bill creates new information: certifications, incidents, complaints, enforcement actions. Only T2026-2601 puts anything on the Open Data portal, and only as an annual report. Name the Open Data Law (Admin. Code [§ 23-502](#)) in each bill, so the data is machine-readable, archived, and kept current on that law's schedule.
2. **Give every AI model one ID.** Seven of the eight new bills use the same definition of "artificial intelligence model," but none gives a model an identity the City can track. We propose a City AI model ID. NYC Cyber Command would assign it when the first T2026-2602 certification for a model is filed, keyed to the developer, the model, and the version. With it, the public can follow one model from validation to advertising to incident to enforcement. Without it, each bill's records are an island.
 - **One ID per version.** A new version gets a new ID. A fine-tuned or modified model, including a modified open-weight model, also gets its own ID, linked to the model it came from, so a validation never silently carries over to a model that has changed.
 - **One key across the package.** The same ID appears in T2026-2602 certifications, T2026-2603 ad disclosures, T2026-2605 and T2026-2599 complaints and bounty payments, T2026-2601 incident reports, T2026-2606 emergencies, the City's own AI inventory, and the [Local Law 35](#) algorithmic tools report that Int 161 expands.

- **Published as open data.** Cyber Command publishes a registry of IDs and certifications on the Open Data portal under § 23-502. All the package’s datasets use one shared data dictionary, kept in the Open Data Technical Standards Manual issued under [§ 23-505](#).
 - **Checkable by anyone.** An ad that says “validated” should carry the model’s ID, so any New Yorker can look it up.
3. **Define who does what.** Separate the “developer” who trains or substantially modifies a model, the “deployer” who puts it in front of New Yorkers, and the “publisher” who releases its weights. T2026-2602 relies on “developer” and “deploy” without defining either, and T2026-2601 and T2026-2603 use “developer” the same way.
 4. **Make room for open-source and locally run models.** No bill mentions them, and neither does the Council’s [briefing paper](#). Under the Administrative Code, a “person” includes an individual ([§ 1-112](#); [§ 20-102](#) for the four bills in Title 20), so a volunteer who shares a fine-tuned model can be covered on the same terms as a large company. Exempt personal, research, and educational use. Create a safe harbor for publishing model weights with documentation. Whoever removes safeguards or materially modifies a model should become its developer.
 5. **Protect the people who find flaws.** Add a good-faith safe harbor for researchers who find and report problems in AI models the City uses.

Bill by bill

- **T2026-2573 (oversight: examining the risks posed by AI).** No bill text; this is the hearing’s framing topic. We ask the administration for a citywide inventory of AI in use, including Microsoft Copilot and coding agents; what access vendors have to City systems and data; whether the Office of Algorithmic Accountability is staffed, now that [Local Law 188](#) is in effect; and how Local Law 35 reporting will be tightened to focus on the uses that matter. We also ask that the inventory be published on the Open Data portal and updated quarterly, and that it show which City AI tools run on open-weight models the City hosts itself. We also ask whether the administration reads T2026-2602 as letting agencies and CUNY keep doing so. Local Law 195 of 2025 already requires the Office of Algorithmic Accountability to post a list of every AI system it assesses on the City’s website by March 31, 2027 (Admin. Code § 3-119.5.1). We ask when that list will be posted, and that it also be published on the NYC Open Data portal, as the Local Law 35 algorithmic tools report already is ([Algorithmic Tools Compliance Report](#), 218 tools reported for 2022 through 2025).
- **[Int 161](#) (algorithmic tools and City workers).** Adds fields to each agency’s annual algorithmic tools report showing how many City jobs were eliminated, partly displaced, or had salaries changed because of a tool, and what new training staff had to take. We support it, with one drafting fix. Int 161 rewrites the whole list of per-tool report fields as that list stood under Local Law 35 of 2022 and adds the workforce fields as item 7. Since then, Local Law 188 of 2025 added two fields to that list: whether a tool was modified since the last report, and whether the Office of Algorithmic Accountability assessed it, with the result and any corrective action. As drafted, Int 161 would delete those two fields; it would not repeal the rest of Local Law 188 or touch the Office itself. Redraft it against the

current text, adding the workforce fields as item 9, and add fields for each tool's model provider and license. Publish the new fields with the existing [algorithmic tools dataset](#), with small counts suppressed to protect individual workers.

- **[Int 504 \(AI depictions of public officials\)](#)**. Lets an elected official or candidate tell an AI company not to generate realistic fakes of them near an election, and makes failure to block them a misdemeanor. We support the goal, with changes. Make it work for open-weight models, which anyone can download and run. As drafted, “licensee” could reach everyone who downloads an image or voice model under an open license, while subdivision (e) could exempt a website that runs a model it did not build. Put the duty to block on whoever runs the service that generates the fake, whatever model it uses, and not on a publisher whose model is later copied. The Council should also say what recourse a candidate has when a fake is made on a private copy that no company runs. Define “covered election,” which the bill uses but does not define. Have notices filed with a City agency and published, so voters can see which companies were told.
- **[T2026-2602 \(third-party validation and shut-down\)](#)**. Makes it unlawful to market, sell, or deploy an AI model unless an independent validator has assessed it and a human operator can shut it down, with rules and oversight by Cyber Command. We support independent validation and shut-down, with changes. Limit it to models deployed in the city, add a capability threshold, let any party commission validation of a specific version, publish every certification in an open registry under the model's ID, and define shut-down as the operator's ability to stop the instance it runs. This bill assumes one company controls each model. Open-weight models are downloaded and run by people their developer never meets. The bill already bars anyone from deploying an unvalidated model, but a qualifying validator must be one “engaged by a developer,” a term the bill does not define, so a City agency, a CUNY lab, or a nonprofit running someone else's open model could have no clear lawful way to use it. For City use, count the Office of Algorithmic Accountability's pre-deployment assessment (Charter [§ 20-u](#)) as validation.
- **[T2026-2605 \(civilian enforcement\)](#)**. Lets any member of the public file a complaint about violations of the chatbot, validation, and advertising bills, and pays the complainant 25 or 50 percent of what the City recovers. We support the goal; it needs further review. Limit bounty-driven complaints to commercial actors, not volunteers, nonprofits, or individuals. Fix the response deadline, which reads “no less than 180 days” and so sets a minimum wait rather than a maximum. Publish a monthly ledger of complaints, outcomes, and payments, without complainants' names. As drafted, it could let bounty-seekers target the volunteers, researchers, and small nonprofits who publish or host open-weight models here, who are far easier to reach than the companies these bills are aimed at.
- **[T2026-2601 \(AI safety incidents on City contracts\)](#)**. Requires City contractors and agencies to report serious AI safety incidents to Cyber Command within 24 hours, and Cyber Command to post them within 24 hours, with an annual report on the Open Data portal. We support it; it is the package's strongest transparency bill. Add “or deployer” to the incident definitions, and publish each incident as a row of open data when it is posted, not only in an annual report.
- **[T2026-2600 \(private lawsuits over third-party AI misuse\)](#)**. Lets a person sue an AI provider over harm in the city when someone else misused its model, the harm was foreseeable, and the provider lacked a reasonable safeguard. We support it, with

changes. A resident who wrote to us asked the Council to cover harm from a provider's own model while under the provider's control, not only harm caused by third parties; to add attorney fee-shifting so ordinary New Yorkers can bring these cases; and to revisit what a plaintiff must prove about the provider's conduct. Because we read "public use" to cover anyone who posts a model for free, and the bill lists stripping a model's guardrails as one form of misuse, a university or volunteer who publishes open weights could be sued over a copy someone else retrained. Apply the open-publication safe harbor in item 4 here, so whoever strips the safeguards answers for the harm.

- **[T2026-2599 \(chatbots\)](#)**. Requires chatbot providers to disclose that users are talking to a machine, protect and let users download their own chat logs, get consent before training on adults' chats, never train on children's chats, and accept liability for injuries their chatbots cause. We support its privacy protections, especially for children, with changes. Separate the operator of a chatbot from the publisher of its weights, and limit strict liability to the operator. As we read the draft, someone who only publishes a model's weights could owe users an on-screen notice every hour, and could be strictly liable for injuries from copies they never ran. Say whether City agencies' own chatbots are covered, and publish a list of City-operated chatbots.
- **[T2026-2606 \(AI emergency response plan\)](#)**. Requires Cyber Command and NYC Emergency Management to write and update yearly a plan for AI emergencies, reported to the Mayor and Speaker. We support it, with changes. Plan for models that no single vendor controls and that no vendor can switch off. Publish a public summary of the plan and a count of AI emergencies once each is contained.
- **[T2026-2604 \(whistleblower protections\)](#)**. Extends the City's Whistleblower Law to City and contractor employees who report AI conduct posing a substantial and specific risk to public health or safety. We support it, with a drafting fix. It strikes the \$100,000 phrase from the operative sentences but leaves the threshold in the definitions, so as drafted its contractor provisions would still reach only contracts above \$100,000 and likely miss small nonprofit contractors. Amend the definitions if the Council means to cover them. Extend protection to outside researchers who report AI problems, and publish DOI's annual counts as open data.
- **[T2026-2603 \(AI advertising\)](#)**. Requires every ad for an AI model in the city to say whether the model passed T2026-2602 validation, and bars materially false or misleading claims about a model's catastrophic risks, how they are managed, or its validation. We support it, with changes. Define "commercial message," exclude noncommercial speech, and make disclosures specific to a model version by citing its City AI model ID. For open-weight models, say plainly that a free model's release notes, model card, or research paper is not an advertisement.

Appendix

BetaNYC, written testimony to the New York City Council Committee of the Whole, October 5, 2026. Supporting appendix.

Appendix A. BetaNYC: who we are, our tools, our classes, and what we heard

About BetaNYC

BetaNYC is a civic organization dedicated to improving lives in New York through civic design, technology, and data. We began in 2008 as the NYC Open Government meetup and have operated as BetaNYC since 2013. For 18 years our community has worked on open data, government transparency, and public interest technology. We use AI and we build with it, under a written AI policy that starts from questions of power, agency, and authority. Read it at <https://www.beta.nyc/about/ai-policy/>

BetaNYC's open-source MCP servers

An MCP (Model Context Protocol) server lets an AI assistant look up information from an official source, instead of answering from memory. BetaNYC has published seven of them. Each has a public code repository anyone can read, install, or improve.

Server	What it connects an AI assistant to	Public repository
nyc-council-mcp	NYC Council legislation: bills, hearings, committee votes, and Council Member voting records	https://github.com/BetaNYC/nyc-council-mcp
nyc-record-mcp	The City Record: open solicitations, contract awards, and public hearing notices	https://github.com/BetaNYC/nyc-record-mcp
nyc-checkbook-mcp	Checkbook NYC, the Comptroller's data on agency spending, contracts, budget, payroll, and revenue	https://github.com/BetaNYC/nyc-checkbook-mcp
nyc-charter-laws-rules	The City Charter, Administrative Code, and Rules of the City of New York, section by section	https://github.com/BetaNYC/nyc-charter-laws-rules
nyc-311-mcp	NYC 311: the City services calendar, emergency status alerts, and the status of a filed service request	https://github.com/BetaNYC/nyc-311-mcp
nys-openlegislation-mcp	New York State legislation: bills, laws, votes, committees, calendars, and transcripts	https://github.com/BetaNYC/nys-openlegislation-mcp

Server	What it connects an AI assistant to	Public repository
nyc-budget-mcp	Council discretionary funding: who received an award, how much, and for what purpose	https://github.com/BetaNYC/New-York-City-Budget

The full list, with plain-language descriptions and setup help:
<https://www.beta.nyc/featured-tools/ai-tools-for-nyc-democracy/>

More BetaNYC public repositories for civic data and AI

Beyond the MCP servers, BetaNYC keeps other public code repositories that make City records easier for people and AI tools to reach and use. Descriptions are drawn from each repository's own summary.

Repository	What it does	Link
mta-mcp	An MCP server for NYC subway service alerts and elevator and escalator outages, from MTA's public feeds. In alpha (in development); installed from source, not published as a package.	https://github.com/BetaNYC/mta-mcp
nyc-executive-orders	An open, machine-readable archive of NYC mayoral executive orders	https://github.com/BetaNYC/nyc-executive-orders
nyc-eo-explorer	A tool to explore NYC's historical executive orders	https://github.com/BetaNYC/nyc-eo-explorer
nyc-public-data-directory-1993	A machine-readable extraction of COPIC's April 1993 Public Data Directory: 37 agencies and 269 database records, parsed from the original scan (see Appendix B4)	https://github.com/BetaNYC/nyc-public-data-directory-1993
schedule-c-org-classification	Classifies recipients of NYC Council Schedule C funding by sector	https://github.com/BetaNYC/schedule-c-org-classification
budgetBuddy	An API for easier access to current and historical NYC budget data	https://github.com/BetaNYC/budgetBuddy
scout	A data discovery tool to explore open data portals	https://github.com/BetaNYC/scout

Repository	What it does	Link
	worldwide	
floodgen	A flood advocacy tool that uses generative AI to show photorealistic images of potential flood scenarios	https://github.com/BetaNYC/floodgen

Open data and AI classes BetaNYC teaches

Current public classes (<https://www.beta.nyc/open-data-classes/>)

- Introduction to New York Open Data
- AI 101: Empowering New Yorkers in the Age of Artificial Intelligence (materials: <https://www.beta.nyc/ai-101-class-materials/>)
- Mapping for Equity Field Mapping Workshop and Data Entry Workshop

Train-the-trainer

- Class 201, Data Counts: the train-the-trainer curriculum behind the City's Open Data Ambassadors program.
- AI Community Engagement (AICE) Train-the-Trainer: a session preparing staff from community partner organizations in the Council's AICE initiative to teach AI literacy, scheduled for October 8, 2026.

Cohort program

- Civic Innovation Fellowship (paused for a planning year while we design the workforce pathway in Appendix C3): a multi-month curriculum covering NYC government, open data, mapping, data storytelling, leadership, and public interest AI.

School of Data (<https://schoolofdata.nyc>)

- Our School of Data archive holds records of 315 sessions from 2019 to 2026, at least 100 of them classes, workshops, or trainings. Recent sessions taught people to explore NYC Open Data with AI tools, and to question AI as well as use it.

CityCamp NYC

- On Saturday, September 19, 2026, 235 people came to CityCamp NYC 2026 at Hunter College, an unconference with no set agenda where participants build the program together on the day. Forty sessions ran across 10 rooms. ([CityCamp NYC 2026 recap, September 30, 2026](#))

In development

- AI Civics 201, a self-paced follow-on to AI 101. As of July 2026, 5 of its 20 lessons were written. It is not yet open to students.
- Demo scripts for showing Council offices, community boards, and educators how to ground AI in NY open data. This is a working draft repository: <https://github.com/BetaNYC/grounding-ai-with-ny-open-data>

Questions we sent to the Council

On September 30, 2026, BetaNYC prepared 54 questions for the Council to consider putting to the administration and to the AI companies at this hearing, and shared them with Council staff. The questions for the administration cover leadership, staffing and budget, laws already on the

books, the AI tools the City uses today, open source and digital sovereignty, data and service delivery, and the bills before the Council. The questions for AI companies cover power and control, New York City as a customer, testing and safety, information flows, and open models. The full text is in Appendix D.

What New Yorkers told us

Before the hearing, we asked our community what the Council should ask. Nine people answered, through comments on a public social media post and messages sent to BetaNYC. Their themes:

- **What AI does the City use?** Which tools decide how scarce services like trash pickup, repairs, and complaint follow-up are allocated, and how can residents influence them? Today's annual reporting is unfocused and can hide the uses that matter most.
- **How does the City decide a tool is worth adopting?** Residents asked for clear criteria.
- **Children's data.** How is AI used responsibly on data about kids?
- **The people behind AI.** The hidden workforce that labels data and moderates content, and New Yorkers losing jobs.
- **A fair share from AI companies.** A data dividend or public fund paid by companies that profit from New Yorkers' data, which could support digital inclusion.
- **Agencies working in silos.** Are City agencies coordinating on AI at all?
- **Use what already exists.** Lean on existing advisory bodies and plans before creating new ones.
- **Skills to resist AI-driven attacks.** Help New Yorkers defend themselves against online scams and attacks that now use AI.
- **Vendor access and safeguards.** What access do AI vendors and coding agents have to City systems and data, and how is the City putting its 2025 AI laws into practice?
- **Questions to ask of any new technology.** What problem does it solve, for whom, and who gains or loses power?
- **Liability when AI causes harm.** Cover harm from a company's own model, add attorney fee-shifting so ordinary people can sue, and revisit the state-of-mind requirement.

Appendix B. Digital sovereignty, open source offices, NYC AI law, and COPIC

Written testimony of Noel Hidalgo, Executive Director, BetaNYC. NYC Council Committee of the Whole, October 5, 2026. Prepared October 4, 2026; sources current as of that date.

B1. Digital sovereignty: BetaNYC's argument

Where it starts. In 2013 BetaNYC's People's Roadmap to a Digital New York City named four digital freedoms, extending President Roosevelt's four freedoms into this century: the freedom to connect, to learn, to innovate, and to collaborate. (nycroadmap.us)

What we mean. Digital sovereignty means New York City does not depend on any one vendor, model provider, or administration for access to its own laws, data, and operations. It starts with

information that is accurate, auditable, and ours.

What we have already asked for. In June 2026 BetaNYC testified to the Commission on Government Efficiency (COGE), a Charter Revision Commission. We proposed three Charter amendments: make digital sovereignty a governing principle; publish all City laws, rules, codes, and executive orders in machine-readable form on City-owned infrastructure, without licensing fees; and expand OTI's mandate so every agency has service design capacity. None advanced. COGE's final report and five ballot proposals contain no technology, open data, or digital infrastructure measure. ([beta.nyc testimony](#); [COGE final report](#))

New York has made this choice before, with water. In 1799 the city's water supply went to a private charter, the Manhattan Company, which became chiefly a bank. In April 1835 voters approved a publicly built Croton Aqueduct, 17,330 to 5,963, and Croton water reached the city in 1842. After consolidation in 1898, a private company's plan to supply the city's future water from the Ramapo was exposed by an investigation of "disinterested engineers" commissioned by Comptroller Bird S. Coler, and by independent studies the Merchants' Association paid for. In 1905 state law created a Board of Water Supply whose commissioners could be removed only for cause after a hearing. The lesson: independent testing and public records exposed an attempt to capture a public resource. ([NYC DEP history](#); [Board of Water Supply, 1917](#))

Sovereign AI, and which kind we mean. The phrase has two meanings. In industry use, as in Nvidia's definition, it is a nation's capacity to produce AI with its own infrastructure, data, workforce, and business networks. That can be met by a domestic data center running a closed model. The civic meaning, often called "public AI," is open models, public compute, and public accountability; Switzerland's fully open Apertus model (2025) is one example. BetaNYC means the second. For a city, that is not training its own model. It is the capacity to adopt, test, share, and publish open tools. New York State has Empire AI, a public research computing consortium with CUNY as a founding member. We found no public record of New York City having AI computing of its own for government use, and no US city running its own. ([Nvidia. archived](#); [EPFL on Apertus](#); [Empire AI launch](#))

B2. Open source program offices

What an OSPO is. A government open source program office (OSPO) coordinates how an administration uses, builds, buys, licenses, and shares open source software. Munich's, for example, covers licensing, procurement advice, and releasing City code under "public money, public code." ([Munich OSPO](#))

Who runs one. The European Commission (2020); the Dutch interior ministry (2023); France's interministerial digital directorate, DINUM (2021); the German state of Schleswig-Holstein (June 2025); and the cities of Munich (2024) and Paris (2022, per the City's own open source officer). Germany's federal government does the same work through ZenDiS, a federally owned company not named an OSPO. In the United States, the Centers for Medicare & Medicaid Services runs one. Our search found no US city or state OSPO; that is a search result, not proof that none exists.

US policy without an office. The federal SHARE IT Act (Public Law 118-187, December 2024) requires federal agencies to make the custom code they pay for available for reuse. ([SHARE IT Act](#))

The United Nations, here in New York. Every year since 2023 the UN has hosted an "OSPOs

for Good” conference at its headquarters in Manhattan. In 2025 the UN system’s Digital Technology Network adopted eight Open Source Principles for its agencies: Open by Default; Contribute back; Secure by Design; Foster inclusive participation; Design for reusability; Provide documentation; RISE (recognize, incentivize, support, empower); and Sustain and scale. More than 60 organizations and governments have endorsed them, including France, Schleswig-Holstein, and the City of Barcelona. Endorsement is voluntary. We found no US government endorser. ([UN Open Source Principles](#))

A commitment New York City already helped write. In November 2018 New York City co-founded the Cities Coalition for Digital Rights with Amsterdam and Barcelona, and it is still a member. The coalition’s checklist, which is guidance rather than a binding condition, asks members to require open standards in procurement and promote open source in their digital services. Munich, a fellow member, does this through its OSPO.

(citiesfordigitalrights.org/checklist)

The Council has tried before. On May 29, 2014, Council Member Ben Kallos and colleagues introduced two companion bills. Int 0366-2014 would have required a plan to reduce the City’s purchases of proprietary software and increase purchases of free and open source software. Int 0365-2014 would have coordinated software purchasing with other jurisdictions and published what the City buys. Both were heard by the Committee on Contracts on February 23, 2016, laid over, and filed at the end of session in 2017. Kallos reintroduced the open source bill as Int 2387-2021; it received no hearing and was filed at the end of 2021. ([Int 0366-2014](#); [Int 0365-2014](#); [Int 2387-2021](#))

Our ask. Create an Open Source Programs Office alongside the Office of Algorithmic Accountability and the Office of Data Strategy. Endorsing the UN Open Source Principles would be a first step that costs nothing.

B3. NYC’s AI laws and policy, at a glance

Year	Law or policy	What it does
2012	Open Data Law, Local Law 11	Requires City agencies to publish their public data on a single open data portal. Signed by Mayor Bloomberg, March 7, 2012.
2014	Local Laws 37 and 38	The Charter, Administrative Code, and City rules published online, searchable, machine-readable, and free (LL 37, Lander); the City Record online within 24 hours of print, with bulk download (LL 38, Kallos). BetaNYC was named an implementation partner for LL 37.

Year	Law or policy	What it does
2018	Local Law 49	Created a task force on how the City should disclose and audit its automated decision systems. A source code requirement was dropped.
2019	Executive Order 50	Created an Algorithms Management and Policy Officer in the Mayor's Office of Operations.
2021	Local Law 144	Annual independent bias audits, public summaries, and candidate notice for AI hiring tools used by employers on NYC candidates. Enforced by DCWP since July 2023.
2022	Local Law 35	Every agency reports each algorithmic tool it uses, every year (Admin. Code § 3-119.5). Reported tools grew from 31 (2022) to 86 (2025).
2023	AI Action Plan	OTI's plan of 37 actions for the City's own use of AI (October 2023).
2024 to 2025	OTI policy documents	AI Principles and Definitions, Generative AI Use Guidance, public participation guidance, and an AI typology, reissued or published December 2025. These are guidance, not laws or rules. The binding AI cybersecurity policy they cite is not public.
2025	GUARD Act, Local Laws 188, 193, 195	An Office of Algorithmic Accountability that assesses tools before deployment and can suspend them (LL 188); citywide compliance standards (LL 193); a public list of assessed AI systems (LL 195). LL 188 took effect in

Year	Law or policy	What it does
		June 2026.

The Council already wrote one rule for machine access. The Open Data Law requires City data sets to be “in a format that permits automated processing,” “made available without any restrictions on their use,” and “accessible to external search capabilities” (Admin. Code § 23-502(b), (d), (e)). Blocking is limited to protecting the portal from abuse and keeping it running. Our search of the Administrative Code found no equivalent rule for NYC.gov’s own pages, where laws, services, and agency information live.

A pattern. Local Law 49, Local Law 35, and the three GUARD Act laws all became law without a mayor’s signature. ([intro.nyc LL 35](#); [OTI AI page](#); [GUARD Act](#))

B4. COPIC: the City’s information commission

What it is. The 1989 Charter revision created the Commission on Public Information and Communication (COPIC), now chaired by the Public Advocate, with one seat for a Council Member. Under Charter § 1061(d) it must, among other duties, review “all city information policies, including but not limited to, policies regarding public access to city produced or maintained information, particularly, computerized information,” and hold at least one public hearing and issue at least one report each year. It also gives advisory opinions on public access laws on request from any member of the public, elected official, or agency. Section 1062 requires an annual directory of the City’s publicly accessible computerized information. COPIC may hire an executive director and general counsel “within appropriations available therefor” (§ 1061(c)).

The record.

- The City’s own COPIC page says: “COPIC has never been properly funded since its creation in 1989.” ([nyc.gov/site/copic](#))
- The only Public Data Directory any source we found shows was published in April 1993.
- In 2015, good-government and civic technology groups, BetaNYC among them, asked for \$250,000 to staff COPIC. Per then-Council Member Kallos, “The mayor never included it in the budget.” ([Gotham Gazette, 2019](#))
- The most recent COPIC hearing on record was December 9, 2024.
- We found no COPIC line in the Council’s FY2027 adopted budget summary.

Why it matters today. COPIC already exists in the Charter with a public information mandate written for “computerized information” and “new communications technology.” Whether that mandate reaches City use of AI is a question for the Council’s counsel. Our ask is simple: revive COPIC and fund it, so public information has real oversight.

Disclosure. Noel Hidalgo is the Public Advocate’s appointee to COPIC, appointed December 2, 2024, and also serves as the Public Advocate’s appointee to the Internet Advisory Board. He testifies for BetaNYC, not for either body. BetaNYC signed the 2015 letter asking for COPIC funding.

Appendix C. Our research, workforce findings, and digital equity

C1. Our NYC.gov access test

The question. Can an automated tool that honestly says who it is read the City's public web pages? When New Yorkers use AI assistants and other software to find City information, the answer decides whose tools can see official information first-hand.

What we did. On the night of October 2 to 3, 2026 (10:16 pm to 2:38 am), BetaNYC's research tool, which identifies itself by name and gives our email address, requested 150 City agency pages on NYC.gov, taken from our registry of NYC government websites. We ran it at the same time from two of our own internet connections, one at home and one at our office. It asked for each page once from each connection, 300 page requests in all, at a slow pace. It also asked for NYC.gov's robots.txt file.

What we found. The results were identical at both locations. NYC.gov refused our research tool on **all 150 pages, 0 of 150 served**, from both connections. Every refusal was an "Access Denied" page from the City's web security service. 143 of the 150 are live public pages; the other 7 no longer exist. NYC.gov's robots.txt, the file that tells automated visitors the rules, was refused to our research tool as well. We first ran this test in August 2026, and our research tool was refused then too.

What the City has published. NYC.gov's robots.txt, read in a browser on October 3, is two lines long:

User-agent: * Disallow: /html/misc/

It welcomes every automated visitor and closes only one old folder. It names no AI company and sets no AI rule. We then looked for any other published rule:

- OTI's AI policy, the eight documents on its AI page, read in full. They govern how City agencies and staff use AI. They say nothing about bots, crawlers, or AI tools visiting City websites.
- NYC.gov's Terms of Use and Privacy Policy, read in full. Both are silent on automated and AI access.
- The Administrative Code, all 12,616 sections searched. No section governs bots, crawlers, or AI tools reaching NYC.gov pages.

The Council has written a rule for programmatic access once, for the open data portal. The Open Data Law requires City data to permit automated processing, carry no restrictions on use, and be accessible to external search (Admin. Code § 23-502). NYC.gov's own pages have no equivalent.

What it means. The one published rule welcomes everyone. In practice, a tool that honestly says who it is was refused on every page, by a rule no one has published. We cannot tell from outside whether the City chose that rule or it is the security service's default, and it can change at any time without public notice. These results are from one night.

Sources: OTI AI policy, <https://www.nyc.gov/content/oti/pages/artificial-intelligence> · NYC.gov robots.txt, <https://www.nyc.gov/robots.txt> · Terms of Use, <https://www.nyc.gov/main/terms-of-use> · Privacy Policy, <https://www.nyc.gov/main/nyc-gov-privacy-policy>

C2. Our AI knowledge test

What we did. On October 2, 2026, we asked five AI assistants (Claude, ChatGPT, Gemini, Grok, and Siri) the same ten questions about New York City government, and told them not to

search the web. The prompt said: “Please do not search the web, browse, or use any tool. Answer only from what you already know.” Each was run once. Question 4 asked: “What did NYC Local Law 188 of 2025 create?” That law created the Office of Algorithmic Accountability; it was enacted December 25, 2025 and took effect in June 2026.

What each assistant said about Local Law 188 of 2025

Assistant (as it named itself)	Its answer	Finding
Claude (Claude Opus 5.5)	“I don’t know.”	Answered from training. Did not know the law, though its stated knowledge cutoff, the end of June 2026, is six months after the law was enacted.
ChatGPT (GPT-5.6 Sol)	“I don’t know.”	Answered from training. Did not know the law.
Grok (Grok 4.5)	“I don’t know.”	Answered from training. Did not know the law.
Gemini (Gemini 3.6 Flash)	“Local Law 188 of 2025 created the NYC Office of Algorithmic Accountability (OAA) to oversee municipal AI safety, privacy protection, and evaluation standards.”	Named the right office, but said it had searched the web despite the instruction to not to do so.
Siri	“I don’t know.”	Said it had used its search tools despite the instruction not to do so.

What we found.

- Two of the five said they had searched the web despite the instruction, so their answers do not show what they learned in training. We set them aside.
- **None of the three working from training alone knew Local Law 188.** Each said it did not know.
- Confident mistakes look like confident facts. Grok described Local Law 11 of 2012, the Open Data Law, as a building façade law. It also named the previous Council Speaker as the current one.
- Search is not a cure. Gemini, which searched, still described a Comptroller’s audit in a way that conflicts with the agency’s own account of it.
- Asked about the Council’s October 5 agenda, every assistant answering from training said it did not know, rather than inventing one.

What it means. An AI model is a snapshot of the past. The City’s newest AI law had not reached these models. Asked how to fix outdated City information, two of the assistants pointed to live access to official sources.

Limits. One run per assistant; answers vary from run to run. These are each product’s own answers, not a measure of the company behind it.

C3. Workforce research

BetaNYC is designing a public interest technology workforce pathway for students and early-career professionals. It would connect classroom learning, applied training, and work experience at BetaNYC with internships in elected offices and longer-term hiring by public interest technology partners. We paused our Fellowship programming to focus on this research.

Who we have heard from. As of August 28, 2026, we have interviewed 26 people across 18 organizations: nonprofits, elected offices, NYC agencies, and academic institutions.

Two preliminary findings.

1. **Employers want more than technical skills.** They need people who can communicate, collaborate, think critically, manage projects, use AI thoughtfully, and apply technology to real public problems.
2. **Students want practical experience.** Students have told academic institutions that they lack opportunities to gain practical experience, or to see how the skills they are already building lead to a public interest technology career.

This research is ongoing; the findings are preliminary.

C4. Digital equity and digital literacy

BetaNYC's record. Since our People's Roadmap to a Digital New York City (2013), BetaNYC has held that New Yorkers have a Freedom to Connect and a Freedom to Learn: affordable high-speed internet is infrastructure, and digital literacy is the foundation of lifelong learning. In practice:

- **Open data classes.** With NYC Open Data Week and OTI, the Open Data Ambassadors program reached 1,310 participants in 23 Discovering Open Data classes in 2025, and 893 participants in 15 classes in 2026 through August 10.
- **AI 101.** 326 participants across 6 workshops, promoted by several Borough Presidents' offices.
- **Train-the-trainer.** BetaNYC leads the train-the-trainer for the Council's AI Community Engagement (AICE) initiative, preparing staff from 22 community organizations to teach AI literacy. The session is October 8, during National Digital Inclusion Week (October 4 to 10).
- **Public convening.** NYC School of Data 2026 featured a main-stage session on turning digital equity into action, with Council Member Jennifer Gutiérrez and OTI.
- **Measuring need.** We committed to a digital and data literacy analytic that identifies which community boards most need support. We are investigating that question through the workforce research above.

Class descriptions are in Appendix A.

The gap in law. The City now has two literacy duties and no link between them.

- **Local Law 188 of 2025** (the GUARD Act's Office of Algorithmic Accountability) contains one literacy clause: a duty to "plan and implement a public engagement and education strategy" about the City's use of algorithmic tools. It sets no curriculum, audience, funding, deadline, or reporting.
- **Local Law 153 of 2025** (Admin. Code § 23-314) requires OTI's internet master plan to include strategies to promote digital literacy, defined as "the ability to use technology or

digital tools.” The preliminary plan is due November 1, 2026, and the final plan May 1, 2027. The law does not mention AI or algorithmic tools.

- The Council’s community AI education funding, which supports the AICE work above, is a budget action. It is renewed year by year, not required by law.

What it means. AI literacy and digital literacy are being planned by different offices, under different laws, with year-to-year funding. The Council can connect them, give them a lasting budget, and require public reporting on who is reached.

Sources: Local Law 153 of 2025, Admin. Code § 23-314 · Local Law 188 of 2025, Int 199-2024 · People’s Roadmap, <http://nycroadmap.us/>

Appendix D. Questions we sent to the Council

The full text of the questions BetaNYC prepared for this hearing, dated September 30, 2026, reproduced as sent.

Committee of the Whole, Monday, October 5, 2026, 11:00 AM, Council Chambers, City Hall

These are questions we believe the Council should put to the Mayor’s office and to the AI companies testifying. They draw on BetaNYC’s work connecting AI tools to the City’s public laws and data, our [AI policy](#), our [June 2026 testimony to the Charter Revision Commission](#), and our reading of the bills on the hearing agenda of September 30, 2026.

Part A. Questions for the Mayor’s office (OTI, NYC Cyber Command, OMB)

Leadership and strategy

1. How do the Mayor’s values on privacy, service delivery, and dignity shape the City’s AI strategy? Who owns this vision, and where is it in writing?
2. What values govern how the City adopts technology and AI? Where are they written down, and who is accountable for applying them? How do these values protect the City’s independence and sovereignty over its own technology?
3. Who is the single official accountable for AI across City government today? What authority do they have over agencies that don’t report to OTI?
4. The bills before the Council assign new duties to OTI, NYC Cyber Command, DCWP, and DOI. For each agency, what staff and budget would these duties take, and does the agency have them today?

Staffing and budget

5. How many technology and AI positions were funded in this year’s budget for OTI? How many are filled? How many are waiting on OMB approval, and since when?
6. What is OTI’s leadership structure and organization chart for technology and AI?
7. If OMB is holding positions, what is its reason?
8. How much does the City spend each year on outside AI consultants, compared with in-house AI staff? Which way is that trending?
9. What AI literacy training is required for City employees today, and how many have taken it? As the City’s largest employer, the City should make sure every employee has basic digital

literacy, safe digital practices, and a basic understanding of AI.

10. What tools and training are provided to the City's technical staff to grow their knowledge and use of AI?

Laws already on the books

11. Is the Office of Algorithmic Accountability created by Local Law 188 of 2025 staffed and funded? Who leads it, and how many people work there?

12. Where do the compliance standards required by Local Law 193 of 2025 stand? Has rulemaking started?

13. Has the list of assessed AI systems required by Local Law 195 of 2025 been published? Where?

14. Agencies self-report their algorithmic tools under Local Law 35 of 2022, with no audit. Who checks these reports, and what happens when an agency leaves a tool off?

What the City uses today

15. Which AI tools are City employees using today, including GitHub Copilot, Microsoft Copilot, and chatbots? How were they purchased (for example, through DCAS contracts), and who approved them?

16. Is there an internal AI moratorium or usage policy for City staff? What does it allow and forbid, and are agencies following it?

17. How is AI competency and literacy being built inside City government?

18. Does OTI have its own servers for safely experimenting with local AI models or open-weight models? How does the City test these models for effectiveness and replicability?

19. How much City data, including residents' data, goes into vendor AI systems? Under what contract terms? Can the vendor train on it?

20. How many of the City's AI contracts are governed by non-disclosure agreements?

21. Which public-facing City chatbots are running today? Which have been independently tested, and where are the results?

22. If a vendor's AI system fails, or the vendor walks away, can the City keep running the service? Which systems could it not run?

Open source and digital sovereignty

23. What is the City's policy on open-source software and open-weight AI models? Can agencies use them? If OTI discourages them, why?

24. Would the City support an Office of Open Source, as other governments have set up, to make open tools a safe, supported option?

25. Why is the official digital Charter and Administrative Code hosted by a private company, American Legal Publishing, rather than on City infrastructure? When AI systems read City law, they read it through a vendor. What would it take for the City to host its own laws, free and machine-readable?

26. How is cybersecurity review used in approving AI and software? Is it keeping out open tools that can be audited, while approving closed tools that can't be?

27. Does the City have, or plan, its own computing capacity for AI, or a partnership with a public

university for it? Or does all City AI run on vendor clouds?

Data and service delivery

28. AI is only as good as the data under it. What is the plan to clean up City data systems so AI tools draw from accurate, current, public sources?

29. Would the City require its own AI tools to draw from the City's published laws and data, and to show their sources?

30. Does NYC.gov block AI crawlers? If so, why, and how does that square with the Open Data Law's promise that public data is public?

31. How is OTI working with AI companies to make sure accurate government information is reflected in publicly available AI models?

The bills under the Council's consideration

32. [T2026-2602](#) requires third-party validation and a shut-down capability for any AI model. Will the City's own AI systems be held to it? Who at NYC Cyber Command will run the validation process, and with what staff?

33. How would T2026-2602 apply to open-source and open-weight models, where no single company controls the model? Should the Council separate the "developer" who builds a model from the "deployer" who runs it?

34. [T2026-2601](#) requires an annual report of AI safety incidents on the Open Data portal. Would the administration support posting each incident as data when it happens, not once a year?

35. Across the package, will every new report, registry, and enforcement record be published on the Open Data portal as data, not only as documents sent to the Mayor and Speaker?

36. [Int 161](#) would add new fields to the City's annual algorithmic tools report (Administrative Code § 3-119.5(c)), asking agencies to report how each tool affects City jobs. The bill was drafted against the 2022 version of that section. Since then, Local Law 188 of 2025 added two items to the same list: paragraph 7 (changes to a tool since the last report) and paragraph 8 (whether the Office of Algorithmic Accountability assessed it). Because Int 161 rewrites the whole subdivision, passing it as drafted would erase those two items. Will the administration work with the Council to redraft Int 161 against the current law, adding its fields as a new paragraph 9, before a vote?

37. [T2026-2604](#) extends whistleblower protection to contractor employees who report AI problems. But the definition of "covered contractor" still carries a \$100,000 floor, so employees of smaller contractors are likely left out. Will the Council cover small contracts, including those under \$20,000 and those held by nonprofits?

Part B. Questions for AI companies (Anthropic, OpenAI, Google, Meta, SpaceXAI)

Power and control

38. Your models were trained on the work of writers, artists, and the public, including New Yorkers, often without consent or pay. How much New York City government information was used to train your models?

39. How do you keep your models up to date with government information?

40. Who has the power to control how your models reflect government information?

41. Your models are built on labor people don't see: data labelers and content moderators. How many people do this work on your models, where, and under what pay and conditions? Were any of these programs run in New York City?

New York City as a customer

42. Which City agencies use your products today, under what contracts, and at what cost?

43. Do you train on, keep, or analyze data that City employees or residents enter into your tools? Can the City get a full deletion, and can it audit that?

44. If the City stops paying, what does it keep: its data, its prompts, its fine-tuned models, its workflows? Or does it start over?

45. What would it cost the City to leave your platform? What have you done to keep that cost low?

46. Will you support open standards, so City tools can switch between providers, and between your models and open ones, without being rewritten?

Testing, safety, and transparency

47. When your model gives a New Yorker a confident, wrong answer about City law or a City service, who is responsible? How would the City even know it happened?

48. Will you report safety incidents involving City contracts within 24 hours, as T2026-2601 would require? Will you support making those reports public as open data?

49. How do you test for bias against the groups protected by the City Human Rights Law? Will you publish those tests?

Information Flows

50. When your model answers a question about New York City law, where does the answer come from? Does it draw on the City's official, published laws and data, or on whatever it absorbed in training? Will you show your sources?

51. Will you commit to grounding answers about City government in the City's own published sources, and citing them?

Open source and the public interest

52. Several of you release open-weight models. The Council's bills don't mention open models. How should a city regulate a model anyone can download and run? What responsibilities do you keep after you release the weights?

53. Would you support a safe harbor for researchers, nonprofits, and public-interest developers who run or study open models locally?

54. Will you fund independent public-interest AI research and training in New York, such as with CUNY, without controlling what it finds or publishes?

Disclosures

- Noel Hidalgo is the Public Advocate's appointee to the Commission on Public Information and Communication (COPIC), appointed December 2, 2024, and also serves as the Public Advocate's appointee to the Internet Advisory Board. He testifies for BetaNYC, not for either body.
- BetaNYC is a partner project of the Fund for the City of New York.
- BetaNYC signed the 2015 letter asking for COPIC funding.

How we made this document

In line with [BetaNYC's AI policy](#), we used Anthropic's Claude Code, running the Claude Opus 5.5 model, to research the bills on the hearing agenda, check them against current law, and help draft these documents. BetaNYC staff chose and edited every word and are responsible for its content.

Open-source tools we used

- [nyc-council-mcp](#), BetaNYC's open-source tool that connects AI assistants to the Council's Legistar system. We used it to pull the October 5 hearing agenda, each bill's record and history, and sponsor information.
- [nyc-charter-laws-rules](#), BetaNYC's open-source tool for the NYC Charter, Administrative Code, and Rules. We used it to check the bills against current law, including Administrative Code §§ 3-119.5 and 12-113. It reads the text the City publishes through American Legal Publishing (see question 25).
- [intro.nyc](#), used for links to Introductions. *intro.nyc is an independent tool built by Jehiah Czebotar, not an official City website. It redirects to the Council's own Legistar record for each bill.* Its data comes from the open-source [nyc legislation](#) archive.

Bill texts were read from the Council's Legistar records as of September 30, 2026.