CITY OF LONG BEACH

Generative AI Guidance

V 1.2

Generative Artificial Intelligence (Gen AI) systems have become increasingly popular and prevalent in all areas of society today. The City recognizes the opportunity for a guided and responsible approach to using Gen AI that acknowledges the benefits to our efficiency while minimizing the risks around AI bias, privacy, and cybersecurity.

Gen AI can serve as an exploratory "personal assistant" tool that can enhance our productivity and support essential City tasks, and as such, **we encourage individual City staff use of Gen AI to enhance productivity**. However, in this Guidance, we provide important **guardrails intended to minimize risks** and share a **list of explicitly prohibited AI systems**. We advise users to protect City and personal information at all times, avoid over-reliance on the technology to limit the environmental impacts of Gen AI use, and not to use it to support critical processes until the City develops further policies.

Summary of Key Guidelines:

I. **Do not submit sensitive City information or private information of any individual or entity** to a Gen AI platform.
II. **Carefully review, verify, and fact-check** via multiple sources the content generated by Gen AI.
III. **Opt out of data collection** and sharing whenever possible.
IV. **We recommend that you create a new account for Gen AI use,** separate from your longbeach.gov email address and your existing personal email.
V. **If you are considering negotiating a contract with a Gen AI provider or using a Gen AI tool not mentioned in this Guidance, you must contact TID.**

For more details and a complete list of guidelines, please refer to "Guidelines" below.

Getting Started:

Currently, Long Beach does not recommend any particular Gen AI system for employees to use, nor does it have a contracted Enterprise Gen AI system in place. **If an email address is required to use a Gen AI system for work, we recommend that you first create a new email address just for professional use** (separate from your City of Long Beach email or your existing personal email), and use the new email to create an account with

one of the Gen AI tools commonly-used in a government context (ChatGPT, Google Gemini, Copilot, or Claude).

Be advised that any newly created email address, along with any address used in relation to a Gen AI system within the scope of your work, may be subject to PRA requests.

Gen AI System Integration:

Integration of Gen AI systems with the City's Microsoft environment (which includes OneDrive, SharePoint and Teams) **is currently not permitted** due to privacy and security risks that these systems pose to City data (e.g., they may share sensitive data with external entities without your consent or introduce vulnerabilities to your devices.). As a best practice, **we recommend that you do not use your longbeach.gov email address in setting up a Gen AI system account.** If you already have a Gen AI system account with your longbeach.gov email, you can continue using it.

If your department is considering negotiating a contract with a Gen AI provider to enable integration with the City's Microsoft environment, **you must work with your cyber and data security teams and inform TID at smartcity@longbeach.gov,** so we can assist you with resources and guidance.

Prohibited AI Systems:

The following AI systems are forbidden to use, due to serious concerns of transparency, data privacy and security: Deepseek, Otter.ai and Read AI, as well as other AI meeting assistants such as Atlas, Avoma, Colibri, Compas, Dubber, Fathom, Fireflies.ai, Gong, Grain, Meet Record, MeetGeek, and Sembly.

Next Steps:

Given the significant and evolving role of technology in our work, cybersecurity is a top priority. It is the responsibility of all City staff to keep our data secure. The City is currently onboarding Copilot by preparing our data environment and training our administrators. When Copilot and Copilot Chat become available for implementation, we will have a secure Enterprise Gen AI System that protects City data. In the meantime, staff must use Gen AI tools responsibly and in compliance with this Guidance.

How to Use this Document:

Refer to this document quarterly, as guidance will change with the technology, laws, and industry best practices. Key sections include:

- **Guidelines:** Practices intended to encourage responsible and innovative use of Generative AI but minimize risks to the City and to the public.

- **Sample use cases:** Several Gen AI prompts that you might use for your work, along with best practices for each.
- **Principles:** Core codes of conduct that reflect our Long Beach values and responsibilities as public servants.

This Guidance is intended to govern Gen AI only – AI systems that specifically generate content, not traditional AI systems at large. Refer to the "Definitions" section below for details.

Guidelines:

1. **Do not submit sensitive City information or private information to a Gen AI platform.** Everything you enter into a Gen AI system is subject to PRA requests. The information you provide may be viewable and usable by the Gen AI service provider and may be compromised if they experience a data breach. Know your data! Remember that the City Code of Conduct requires that we safeguard all information and data entrusted to us. Sensitive and private information includes but is not limited to the following:
    a. Personally identifiable information (PII), Protected Health Information (PHI), Tax IDs, passport numbers, driver's license numbers, addresses, contact details, financial information, credit card numbers, and more.
    b. Passwords, passcodes, and other confidential information.
    c. Sensitive business data:
        i. Any information that if exposed publicly would tarnish or harm the reputation of the City.
        ii. Any information that if exposed publicly would open the City to regulatory or legal action.

2. **To use a Gen AI system for work, we recommend you create a new email address just for professional use** (separate from your City of Long Beach email or your existing personal email) **and use the new email to create an account** with one of the commonly-used Gen AI providers (ChatGPT, Google Gemini, Copilot, or Claude). Note that the new account will be subject to PRA requests if used for work.
    a. We recommend that you do not use your longbeach.gov email address in setting up a Gen AI system account. This is because **integration of Gen AI systems with the City's Microsoft environment** (which includes Email, Calendar, OneDrive, SharePoint and Teams) **is currently not permitted.** If you already have a Gen AI system account with your longbeach.gov email, you can continue to use it.
    b. You are forbidden to use systems like Deepseek, Otter.ai, and Read AI, as well as other AI meeting assistants (such as Atlas, Avoma, Colibri, Compas,

Dubber, Fathom, Fireflies.ai, Gong, Grain, Meet Record, MeetGeek, and Sembly) due to serious concerns of transparency, data privacy, and security.

3. **Carefully review, verify, and fact-check via multiple sources the content generated by Gen AI** to eliminate instances of bias, offensive, inaccurate, or harmful material. Acknowledge that Gen AI systems can reflect cultural, social, and economic biases of source material. You are responsible for verifying that the information is accurate to mitigate biases and misinformation.

4. **Opt out of data collection and sharing whenever possible**. Gen AI retains the data you input and the content it generates as part of its training and self-learning processes. Sometimes, this data is shared with third parties. This means every interaction with Gen AI may be saved and utilized to improve its functionality over time. Many Gen AI systems allow you to opt out of data collection (or delete your chat history), meaning they will not keep any of the data you provide, and those data will not be used in the systems' models.
    a. Instructions on how to opt out can be found in the "How to Opt-Out of Data Collection" section.

5. **If your department is considering negotiating a contract** with a Gen AI provider to enable integration with the City's Microsoft environment – or if you are considering using a new Gen AI tool for your own professional use that has not been mentioned in this document – **you must work with your cyber and data security teams and inform TID at smartcity@longbeach.gov,** so we can assist you with resources and guidance.

6. **You are encouraged to cite or reference when you use Gen AI within your documents and communications so that others can learn from your use cases and are empowered to use Gen AI to enhance their own productivity.** While we realize it is impractical to cite every instance of Gen AI use, we ask that you reference Gen AI in instances where it proved particularly impactful to your work OR where minimal edits were made to the Gen AI output leading up to a final version of your work product (as is likely the case with generating images). Our goal is to build a positive culture around using Gen AI systems and knowing when Gen AI was involved with producing content enhances our collective learning as employees of the City of Long Beach.
    a. Examples of citations are below in the "How to Cite" section.
    b. Your citations should include which Gen AI system was used and confirmation that you have fact-checked everything.

7. **You are encouraged to share with TID when you use Gen AI for particularly impactful and/or creative use cases**. Our goal is to assemble such use cases from across the City to enhance peer-to-peer learning. TID may also use this information to implement further systems and policies dedicated to addressing emergent needs.
    a. Form: https://forms.office.com/g/PTAgN6q988

8. **City Departments may provide additional rules for their staff regarding Gen AI.** If your Department, Bureau, or Division has additional regulations for Gen AI, **please inform TID at smartcity@longbeach.gov,**so we can assist you with resources and guidance.

9. **We encourage you to participate in the City's Gen AI Staff Workgroup.** We'll meet virtually at least once every two months to stay up to date on AI topics, provide feedback on the City's AI work, share out impactful AI use cases, and understand departments' progress and needs. If you are interested in joining, please complete this form:
    a. Form: https://forms.office.com/g/PTAgN6q988

Sample Use Cases:

TID has compiled several staff-identified use cases for Gen AI in a Use Case Registry We encourage employees to complete the log form (https://forms.office.com/g/PTAgN6q988) to share when they use Gen AI for particularly impactful and/or creative use cases. Our goal is to assemble an actionable list of use cases that enhances peer-to-peer learning. Please keep in mind that these are just example use cases; they do not represent the only ways you can use Gen AI systems.

Principles:

These principles describe general codes of conduct that represent the City's values and align with our responsibilities to the residents we serve. They guide City employees in their use of Gen AI technology. **City employees should adhere to these principles and should champion accountability and compliance among their peers.**

The following principles reflect what we heard from City staff in TID's Gen AI survey. They build upon the City's Code of Conduct and Ethics, Smart City Initiative Guiding Principles, and Data Privacy Guidelines.

Safety & Security

- Every technology we employ impacts the safety and security of our overall environment, both physically and digitally. We must be mindful of safe data privacy practices for our constituents and minimize security risks presented by Generative AI. We commit to protecting personal data to the greatest extent possible by ensuring that data that is sensitive, proprietary, and non-public is never used in publicly available Gen AI apps.
- We must be aware that Generative AI applications and models can be compromised with malicious links and prompts if they are not kept up to date and combined with sound security practices.

Empowerment

- The use of Gen AI should support our work to deliver better, safer, more efficient, and equitable services and products to our residents.
- We trust our staff to do the right thing given the right tools and guidance. Staff will need to exercise judgement to ensure we get the benefits from Generative AI tools while avoiding negative impacts on the City and its constituents.

Exploratory Mindset

- We embrace the possibilities of technology and innovation and commit to responsibly exploring and evaluating creative AI technologies that can advance positive outcomes.
- We acknowledge we are still learning about the full scope of what Generative AI can provide us with, so let's learn together.

Transparency & Accountability

- We acknowledge that we do not have all the answers, nor can we foresee all consequences of adopting this technology. We recognize further research into this technology may reveal issues that require more restrictions on its use.
- We will be transparent when updates to this Guidance are implemented and will hold ourselves accountable to the possible impacts this technology may have.
- The City will practice data integrity and use content produced by Gen AI for stated purposes and purposes that serve the public good.
- Generative AI is a tool. We are responsible for the outcomes of our tools. For example, if autocorrect unintentionally changes a word, changing the meaning of something we wrote, we are still responsible for the text. Technology enables our work; it does not excuse our judgment or accountability.

Equity & Inclusion

- We are stewards of the public and we will use Generative AI to support and uplift our communities.
- Everything we do, regardless of the tools, is a reflection of the City and ourselves. We acknowledge that AI systems have the potential to perpetuate inequities, discrimination, harm, and bias against Long Beach residents, and the City will commit to using a racial equity lens to examine the burdens, benefits, and unintended consequences of Gen AI.
- We acknowledge that Generative AI introduces concerns about job security and job displacement, and we recognize that we must be resilient as a workforce in the face of emerging technologies.

How to Cite:

Citing the use of Gen AI is more than a best practice, it is a powerful way to communicate to your audience how Gen AI is improving your work. **Citing your use of Gen AI is easy and shows that you are an early adopter!**

Generative AI can be cited as a statement within your document or as a footnote, endnote, or header. Citations for content created by Generative AI should include the following:

- GEN AI SYSTEM: Name of the Generative AI system (ChatGPT, Google Gemini, Copilot, Midjourney, etc.)
- DD/MM/YYYY: Date when content was generated.
- Confirmation that you fact-checked information.
- *Optional*: State whether the content is unmodified from the Generative AI system.

In-text citation example unmodified: *"This document was drafted with support from (GEN AI SYSTEM) on DD/MM/YYYY. The content was fact-checked by City staff. This information is unmodified."*

In-text citation example modified: "*This report summary was drafted with support from (GEN AI SYSTEM) on DD/MM/YYYY. The content was edited for brevity and fact-checked by City staff."*

Footnote/endnote/header example: *(GEN AI SYSTEM, DD/MM/YY. Fact checked by City staff.)*

How to Opt-Out of Data Collection:

Some Gen AI systems offer the option to opt-out of data collection, and we encourage you to do so whenever possible. These options are typically found within a system's privacy settings or user settings.

One Gen AI system that allows users to opt-out is ChatGPT. Once a user opts out of data collection, their information is not used to train the AI system for other users. Below is a step-by-step guide on opting out of ChatGPT's data collection.

1. Select your name on the top right of the interface, then select “Settings”
2. Click “Data controls” in the settings menu and disable “Improve the model for everyone”

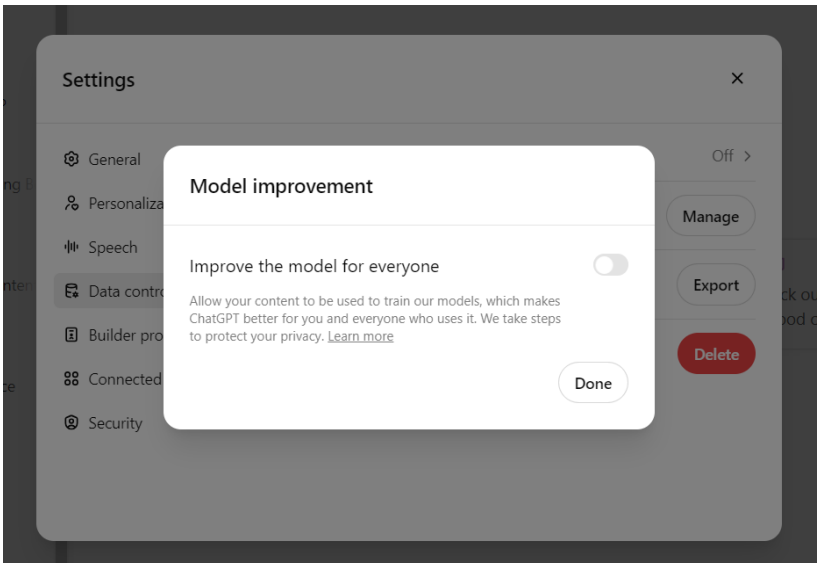

3. In addition, Chat GPT allows users to disable its “memory” so it does not remember details and preferences between your chats, so what you enter in one chat will not impact your next one. To do this, from the Settings menu, select “Personalization,” then disable the two “Reference saved memories” and “Reference chat history” toggles. Chat GPT conversations will only be retained for 30 days in this mode.

Assessing Risk in Gen AI Use Cases:

The risk presented by Gen AI tools varies by use case, with the risk spectrum ranging from mid-risk to high-risk to intolerable risk. San Jose’s Data Privacy Office developed the following framework.

Two key factors determine Gen AI risk:

1. **Risk of information breach**: the potential harm if the information exchanged with a Gen AI is released to an unintended audience. This can include entering personally identifiable information, sensitive records, or confidential business information into Gen AI. Additionally, any information entered into Gen AI may be subject to PRA requests. If you wouldn’t share the information in a public forum, don’t share it with a Gen AI system.
2. **Risk of adverse impact:** the potential harm of using the output for a decision, task, or service. This impact **may manifest itself** differently for different populations and should be considered from an equity lens, such as adverse impacts on people of a

certain race, age, gender identity, or disability status. Not only can Gen AI be biased, but it can also provide false information. In general, if Gen AI is used in relation to City processes that can alter an individual or community's rights, freedoms, or access to services, multiple users should thoroughly review it before any document is finalized or action is taken.

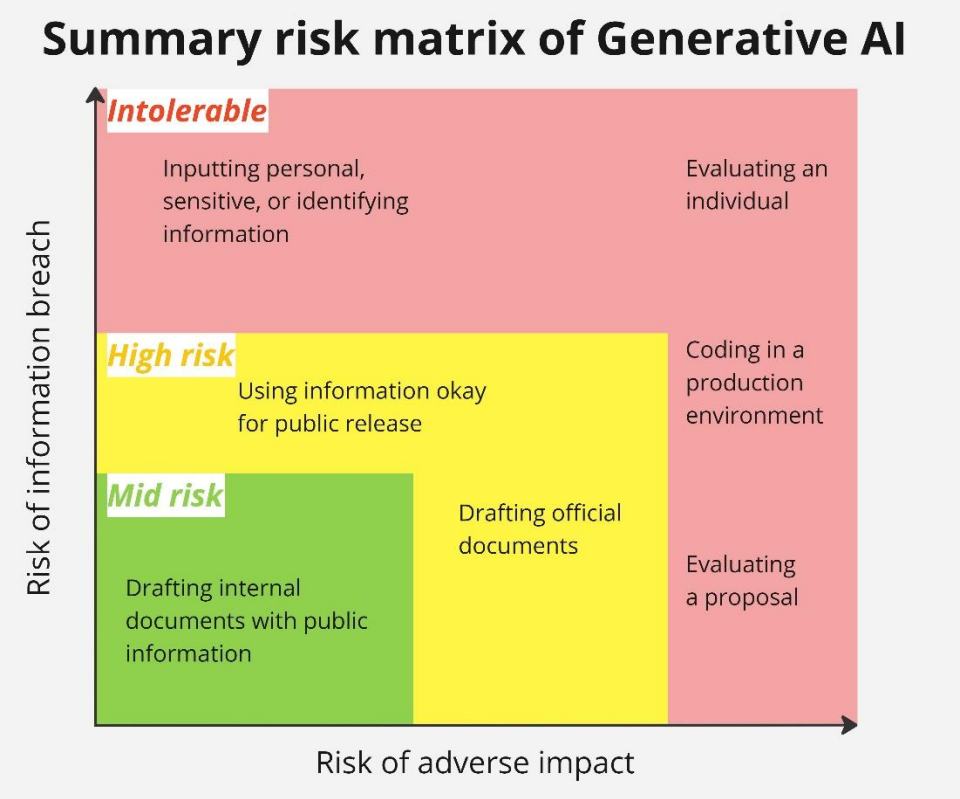

History:

Gen AI is a new branch of AI technology that can generate helpful content—such as emails, research summaries, reports, and images — at a user's request. Popular Generative AI systems include OpenAI's ChatGPT, OpenAI's Dall-E, Google Gemini, Midjourney, Microsoft O365 integration with OpenAI Copilot, Microsoft GitHub Copilot, and Stable Diffusion. These systems have the potential to support many City business functions and services, however their use also raises important questions, particularly around the sourcing of training data, ensuring proper attribution of generated content, and the handling of sensitive or public data. Further research into this technology may uncover issues that require more restrictions on its use. Please remember that Gen AI technology is already incorporated into many existing services and products we use, such as Office 365.

Through a survey distributed to all City of Long Beach staff in October 2023, the Technology and Innovation Department (TID) found interest in using such systems to conduct City business but low levels of actual use. Over 70 percent of respondents described work-related benefits that Gen AI presents. However, just 14 percent of staff respondents said that they had used Gen AI in the past for work-related purposes.

Definitions:

- **Generative Artificial Intelligence**, commonly referred to as "Generative AI" or Gen AI, is an "automated system" used to generate "content." Gen AI uses massive datasets to develop content when given a prompt. Common Gen AI systems include ChatGPT, Google Gemini, MS Copilot, and others.
- **Training data** is data gathered from a wide range of sources available on the internet, such as books, websites, articles, and other texts. The goal is to cover a broad human knowledge and language use spectrum. The Gen AI system uses this data to understand user prompts and generate "content."
- An "**automated system**" is any artificial intelligence system, software, or process that uses computation as part of a system to generate outputs and outcomes, make or aid decisions, inform policy implementation, collect data or observations, or otherwise interact with individuals and communities.
- **Traditional AI** is an artificial intelligence system trained for a specific, narrowly defined task. Commonly used AI systems include Google's search algorithm, recommendation engines like Netflix or Amazon, and writing assistants like Grammarly.
- "**Content**" includes text, emails, presentations, images, video, audio, architectural documents, diagrams, and other forms of media.
- **Large Language Models (LLMs)** are a type of AI model trained on immense amounts of broad data. They are capable of understanding and generating natural language and other types of content to perform a wide range of tasks with minimal fine-tuning.
- "**Personally Identifiable Information,**" or "**PII**" for short, is any data used to identify someone.
- The **California Public Records Act (PRA)** gives residents the right to access the public records that the government maintains. https://www.ftb.ca.gov/your-rights/california-public-records-act.html

Resources and Trainings:

- What Is a Prompt Injection Attack? https://www.ibm.com/think/topics/prompt-injection
- The Emerging Use of Generative AI in City Government (City of Long Beach): https://www.longbeach.gov/smartcity/news/the-emerging-use-of-generative-ai-in-city-government/
- An Intro to Generative AI (City of Long Beach Data Learning Community): https://longbeach.sharepoint.com/:f:/r/sites/Data-Collaborative/Shared%20Documents/General/Trainings/AI?csf=1&web=1&e=jM3T5X
- Government AI Coalition: https://www.sanjoseca.gov/your-government/departments-offices/information-technology/ai-reviews-algorithm-register/govai-coalition
- About Generative AI (Harvard University Information Technology): https://huit.harvard.edu/ai/about
- What is Generative AI? An AI Explains (World Economic Forum): https://www.weforum.org/agenda/2023/02/generative-ai-explain-algorithms-work/

Change Log:

Version	Notes	Date
V1.0	Interim Guidance published	Jan 8, 2024
V1.1	Guidance updated to reflect the evolving AI system landscape, reinforce AI systems' privacy and security risks, and explain the City's posture on common AI systems. Updated glossary to distinguish between Gen AI and traditional AI. Added language about permitted systems. Added Guideline #9.	Aug 13, 2024
V1.2	Added sections 'Summary of Key Guidelines, 'Getting Started,' 'Gen AI System Integration,' 'Prohibited Systems,' and 'Next steps.' Reordered guidelines. Added Guideline #2 and #5. Reframed the requirement to cite AI use as an encouragement to cite. Reframed the requirement to log instances of Gen AI use as an encouragement to log instances of particularly impactful / creative Gen AI use. Updated the 'How to Opt-Out of Data Collection' section with latest features of ChatGPT.	May 27, 2025