



Salt Lake City Corporation

Guide for Applying the Use of Generative Artificial Intelligence Policy to City Business

This document closely models the structure and language of the City of San Jose Generative AI Guidelines, and unless otherwise noted, has been modified where appropriate to conform to Salt Lake City's circumstances, and used with permission from the City of San Jose.

IMS Innovation Team: Version 5.1

Table of Contents

- OVERVIEW 2**
 - How to Use the Guide2
 - What is Generative AI and How Does it Work?3
 - Artificial Intelligence vs. Generative Artificial Intelligence3
- USE OF THE GENERATIVE ARTIFICIAL INTELLIGENCE POLICY AND APPLYING IT TO CITY BUSINESS RESPONSIBILITIES..... 5**
- PRIVACY 7**
 - Applying the Privacy Policy to City Business:.....8
- USE OF GENERATIVE AI..... 8**
 - Applying the Use of Generative AI Policy to City Business:9
- ASSESSING RISK IN USING GENERATIVE AI 11**
 - Be Aware of Targeted Cyber Attacks Using Generative AI 13
- APPENDIX A: Use of Generative Artificial Intelligence Policy..... 16**
- APPENDIX B: Learn More About Generative AI 20**
 - Generative AI 10120
 - Government Professionals and Generative AI20
 - Writing Prompts for Generative AI20
 - Use Cases for Generative AI21
- APPENDIX C: Examples of Generative AI Use Cases by Risk Level 22**
 - Examples of Mid-Risk Use Cases.....22
 - Examples of High-Risk Use Cases23
 - Examples of EXTREMELY High-Risk Use Cases.....24
- APPENDIX D: Glossary of Generative AI Terminology 26**

OVERVIEW

Generative artificial intelligence (AI) presents users with an opportunity to work better, faster, and smarter. This Guide is designed to help Salt Lake City employees, contractors, volunteers, or others conducting City business understand the [Use of Generative Artificial Intelligence](#) policy, and apply it when using generative AI for City business.

The Salt Lake City Information Management Services (IMS) Department is responsible for the management and monitoring of hardware, software, and web-based tools, services, and platforms used by City employees while conducting City business. By keeping the above guidance in mind when using generative AI tools, we can ensure the safe and responsible use of AI for the City.

Please send any questions, comments, or concerns around using generative AI to Innovation@slc.gov. The Innovation Team will work with you to ensure understanding and compliance of the City's [Use of Generative Artificial Intelligence policy](#).

Updated versions of the Guide can be found on the Innovation Tea's webpage: <https://www.slc.gov/ims/innovation-team/projects/>

How to Use the Guide

The Guide is designed to help you understand and follow the Use of Generative Artificial Intelligence policy. If you haven't done so already, take a minute to read the policy. It begins by helping you understand the difference between artificial intelligence (AI) and generative artificial intelligence (GenAI). Next you will find the policy sections Responsibilities, Privacy, and Use of Generative AI, with detailed information for applying the policy to your City work. Next is the risks associated with using generative AI and how to ensure you are using it safely while protecting City data. Lastly are the appendices. Appendix A contains the Use of Generative Artificial Intelligence policy. Appendix B provides resources for learning more about generative AI. Appendix C covers risk levels for different uses of generative AI, including prohibited uses. A glossary of terms is available in Appendix D.

A summary of the fundamental rules when using any generative AI are:

- Never submit personal or confidential information into a generative AI.
- Review, revise, test, and fact check any output from a generative AI.
- Be transparent when content is drafted using generative AI.
- Return to the policy and this document quarterly, as guidance on usage will change rapidly.
- The policy does not apply to personal purposes or business purposes unassociated with Salt Lake City.

GETTING STARTED WITH GENERATIVE AI FOR CITY USE

What is Generative AI and How Does it Work?

Simply stated, generative AI is an artificial intelligence system whose primary function is to generate content, which can take the form of code, text, images, and more.

First, the model is fed a massive amount of text. Second, the machine organizes the text based on words that most often go together. The resulting catalogue of connected words, phrases, and concepts allows it to generate predictions of how to finish a sentence, for example, based on just a few words. With the right amount of sample text—say, a broad swath of the internet—these text models can become quite accurate.

AI systems apply this same approach to music, books, poems, voices, videos, and anything else created on the internet.

It is important to note that caution is always necessary when using generative AI because it has also been used for several irresponsible applications including faking news headlines, leaking personal information, and enabling phishing cyber-attacks.

Generative AI has two inherent – and significant risks:

- The data is not always current, and tools may be drawing on sources that are months or years out of date.
- If the dataset is in any way biased, this bias will be transferred into the generated content.

Artificial Intelligence vs. Generative Artificial Intelligence

The Use of Generative Artificial Intelligence policy specifically applies to use of generative artificial intelligence. The main different between AI and generative AI is that AI focuses on replicating human intelligence, while generative AI focuses on generating new content. The below table provides detailed information to help distinguish the two.

ARTIFICIAL INTELLIGENCE (AI)	GENERATIVE ARTIFICIAL INTELLIGENCE (GENAI)
<p>Technical Definition</p> <p>Machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments.</p>	<p>Technical Definition</p> <p>Branch of AI technology that can generate content – such as stories, poetry, images, voice, music, business letters and other documents that present an argument, summarize information, or analyze information like survey results – at the request of a user.</p>
<p>Less Technical Definition</p> <p>AI is a computer system that can take a set of goals or instructions given by humans and use them to make predictions, offer recommendations, or make decisions. These decisions can affect things in the real world, like telling you the weather, helping to drive a car, or in virtual spaces, like determining the outcome of an online game.</p>	<p>Less Technical Definition</p> <p>GenAI is a branch of AI that can produce various types of content, such as stories, poetry, images, music, and documents like business letters or reports. It works by taking user requests and then generating the content, whether it's creating new ideas, summarizing information, or analyzing data, like survey results. Essentially, it's a tool that can assist in creating and organizing information based on what you need.</p>
<p>Examples of AI</p> <ul style="list-style-type: none"> • Digital Assists: Apple's Siri, Amazon's Alexa, Google Now • Navigation Apps: Google Maps, MapQuest, Waze, Apple Maps • Facial Recognition: unlocking phones, payment methods, banking • Recommendation algorithm: Amazon, Netflix, Spotify, YouTube, LinkedIn, Goodreads • Chatbots: Apple's iMessage, Facebook Messenger, WhatsApp and website chatbots 	<p>Examples of GenAI</p> <ul style="list-style-type: none"> • Text: OpenAI's ChatGPT, Microsoft's Co-Pilot, Google Gemini (Bard), AWS Claude • Image: OpenAI's DALL-E, Google Gemini • Audio: OpenAI's Jukebox • Code generation: GitHub Copilot by OpenAI's Codex

ARTIFICIAL INTELLIGENCE (AI)	GENERATIVE ARTIFICIAL INTELLIGENCE (GENAI)
Examples of Use of AI in City Work <ul style="list-style-type: none"> • ServiceNow • Workday • Autocorrect • Canva • SmartSheet • PowerPoint 	Examples of Use of GenAI in City Work <ul style="list-style-type: none"> • Summarize meeting notes • Extract main ideas from documents (like policies) • Analyze data (for example, producing summary statistics or graphs) • Answer questions about a particular topic appearing in documents or data • Draft content, develop outlines, improve grammar and spelling, transform content into plain language
Is AI included in the Use of Generative Artificial Intelligence policy?	Is generative AI included in the Use of Generative Artificial Intelligence policy?
NO	YES

NOTE: Generative AI relies on a user to “prompt” the AI to generate content. “Prompts” use natural language the user puts into a generative AI that gives the system direction. See APPENDIX B: Learn More About Generative AI for information about prompts.

USE OF THE GENERATIVE ARTIFICIAL INTELLIGENCE POLICY AND APPLYING IT TO CITY BUSINESS

RESPONSIBILITIES

[Policy 52.13.E Use of Generative Artificial Intelligence](#) states that usage of generative AI shall follow Salt Lake City’s AI principles:

- A. **Privacy:** Only submit information to generative AI tools that is ready for public disclosure. This includes any text, photos, videos, or voice recordings shared with the AI. Be mindful that the AI output may include unexpected personal information from another user, so users must remove any potentially private information before publishing.

- a. Anything that would not be released or shared with the public should not be inputted into the AI tools. This includes protected, private, or controlled records, and information such as draft Request for Proposal requirements that should not be public yet, vendor transactions, procurement approvals, or internal City decisions.
 - b. General rule: If the information exchanged with a generative AI system would be harmful to a person or community if made public, it is a high or intolerable risk. Services can be compromised and leak personal information, all information exchanged with generative AI has a reasonable risk of being compromised. If unsure, do not input into generative AI.
 - c. **Mid-risk information includes non-identifying and non-confidential information**. For example, a simple email response or instructive documents often contain only general information that would not present any risk if made public.
 - d. **High-risk information includes personally identifiable information** (e.g., full name, birth date, email address) **and confidential business information** that may have larger implications to City processes.
 - e. **Prohibited risk information includes highly sensitive and identifying information**. This includes data such as credit card numbers, bank account information, social security numbers, and other information that requires rigorous security measures and compliance standards before being processed.
- B. **Accuracy**: Salt Lake City maintains trust with its constituents by providing accurate information. Review and fact check all outputs received from a generative AI. Users should consult trustworthy sources to confirm that the facts and details in the AI-generated content are accurate. Be aware that many systems may only use information up to a certain date and cannot guarantee the content they generate is accurate.
- a. Trustworthy sources include City reports, records, master plans, codes, policies, and department/division webpages. City staff should also consult their manager, department director or division leadership to verify the accuracy of information.
- C. **Transparency**: The user shall be clear when they use generative AI. This includes citing that AI was used in creating a product.
- D. **Equity**: AI system responses are based on patterns and relationships learned from large datasets derived from existing human knowledge, which may contain errors

and may be historically biased across race, sex, gender identity, ability, and many other factors. Users of generative AI must be mindful of any assumptions generative AI may make based on past stereotypes. Users must edit or disregard biased output.

- E. **Accountability:** The person using AI is accountable for the content it generates. Use generative AI with a healthy dose of skepticism. The level of caution used should correspond to the risk level of the use case. It is always important to verify information provided by generative AI.
- F. **Beneficial:** Users should be open to responsibly incorporating generative AI into their work where it can make services better, more just, and more efficient.

PRIVACY

- A. **Opt Out:** Some services offer an option to opt out of data collection. This means the generative AI system will not keep the data you provide, and it will not be used in the system's models. Salt Lake City as a customer of AI products, shall opt out of data collection and model training whenever possible.
- B. **Verify the Copyright of All Generated Content:** Users shall verify the content they use from any generative AI systems does not infringe any copyright laws. If users are uncertain if content violates copyright, they should either edit the content to be original or not use it.
 - a. Visit the U.S. Copyright Office to learn if content is copyrighted:
<https://cocatalog.loc.gov/cgi-bin/Pwebrecon.cgi?DB=local&PAGE=First>
- C. **Ownership of Generated Content:** In most cases, the user owns the content they input into a generative AI service and the information they receive as an output. The user can use the content at their discretion, in accordance with Salt Lake City policy and any terms and conditions they agreed to with the generative AI software. However, many generative AI companies still retain the right to use both the input and output content for their own commercial purposes. For example, this could include a generative AI company using Salt Lake City data to train their models or distributing Salt Lake City output data for marketing campaigns. This emphasizes the importance that only information Salt Lake City is ready to make public should be entered into a generative AI system.

Applying the Privacy Policy to City Business:

- **Opt Out of Data Collection, if possible.** Some generative AI systems offer the option to opt out of data collection, and we encourage you to do so whenever possible. Once a user opts out of data collection, their information is no longer used to train the AI system. Opting out results in a user's chat history no longer impacting future conversations with the AI system and conversations get deleted after a period. How to opt out varies depending on the AI tool. Typically, you will start in "Settings" and from there click on "Data controls" or "Personalization" but it will depend on the tool. Within "Settings" menu, look for wording like "Data controls" or "Personalization". You will want to uncheck "Chat history & training" or click and turn "Memory" off.
- **Verify the Copyright of All Generated Content.** Users shall verify the content they use from any generative AI systems does not infringe any copyright laws. If users are uncertain if content violates copyright, they should either edit the content to be original or not use it. Visit the U.S. Copyright Office to learn if content is copyrighted: <https://cocatalog.loc.gov/cgi-bin/Pwebrecon.cgi?DB=local&PAGE=First>.
- **Ownership of Generated Content.** In most cases, the user owns the content they input into a generative AI tool and the information they receive as an output. The user can use the content at their discretion, in accordance with City policy and any terms and conditions they have agreed to. However, many generative AI companies still retain the right to use both the input and output content for their own commercial purposes.

For example, this could include a generative AI company using City data to train their models or distributing City output data for marketing campaigns. This emphasizes the importance that only information the City is ready to make public should be entered into a generative AI system.

USE OF GENERATIVE AI

Users are required to follow these rules while using generative AI for Salt Lake City work:

- A. Information you enter into generative AI systems could be subject to a Government Records Access and Management Act (GRAMA) request, a record requiring retention considerations, may be viewable and usable by the City, and may be leaked in a data breach. Do not submit any information to a generative AI platform that would not be available to the general public (such as confidential, protected or personally identifiable information).

- B. Thoroughly review, revise, and fact check any output from generative AI. Users are responsible for any material created with AI support.
- C. Users must cite the generative AI when a substantial portion of the content used in the final version comes from the generative AI.
- D. Users may not utilize City generative AI accounts for personal use and users may not utilize personal generative AI accounts for official City business.

Applying the Use of Generative AI Policy to City Business:

Usage of Generative AI may be Subject to GRAMA. Any retained conversations relating to City work may be subject to public records requests and must comply with the City's retention policies and the Government Records Access and Management Act (GRAMA) and other applicable public records laws for all City usage of generative AI (see [Utah Code 63G-2](#)). This means any prompts, outputs, or other information used in relation to a generative AI tool may be released publicly. Do not use any prompts that may include information not meant for public release.

Any information that includes personally identifying information about constituents, City staff, volunteers, consultants etc. could inadvertently be shared with others. Basically, if you wouldn't share it with other people or in a public forum avoid sharing the information in the prompt.

Check Generative AI Results. Review, revise, and fact check via multiple sources any output from a generative AI. Users are responsible for any material created with AI support. Many systems only use information up to a certain date, meaning some generated results are out of date when outputted by the system. When in doubt, staff should consult their supervisors.

Citing Generative AI. Citing the use of generative AI is easy and shows transparency. Cite or reference when a substantial portion of your final document comes from

Citing Generative AI

Example: "This document was drafted with support from ChatGPT. The content was edited and fact-checked by City staff. Sources for facts and figures are provided as they appear."

In text citation example unmodified: "This document was drafted with support from (insert name of 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 email was drafted with support from (insert name of GEN AI SYSTEM) on DD/MM/YYYY. The content was edited for brevity and fact checked by City staff."

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

generative AI. It is best practice for audiences to know when generative AI was involved with producing content, and we encourage you to lead by example and disclose this information whenever possible. Our goal is to build a positive culture around using generative AI systems and being transparent with its usage is the first step in challenging any stigmas that may exist. Generative AI is not a shortcut; it is a tool used to enhance our creativity, optimize our efficiency, and strengthen our effectiveness.

How to cite Generative AI as a footnote, endnote, header, or footer. A best practice for citing text-generated content is to include the following:

- Name of generative AI system used.
- Confirmation that the information was fact-checked.

How to embed citations for images and videos. Some generative AI tools can create images or video clips based on text prompts. The City needs to maintain its legitimacy as a trustworthy source when using video and images, which requires substantial precautions whenever using AI-generated visual content.

1. **Use only for illustrative purposes.** For historical events, use real images rather than generated images. For example, if you want a picture of a giraffe wearing a suit and tie for your presentation, generate it. If you are proposing a new visual diagram or abstract concept, you can also generate it. If you want a picture of the Mayor at City Hall, find the actual picture.
2. **Citations embedded into the image or video.** Images and videos can easily be taken out of their original context and misinterpreted as reality. To prevent a news article or other secondary source from treating an image as fact, all images and frames of a video should specify that they were generated using an AI system. The citation shall be included in the image itself and cannot be removed without editing or cropping the image.

Create an Account Specifically for City-related Work. If you choose to use generative AI for City-related work, you should have an account for all generative AI usage in your role at the City using a City email address. The purpose of this is to ensure proper retention of public records and avoid co-mingling of public and personal records. This account should not be used for any personal purpose. Users should use their City email address for City usage.

If users use personal devices or accounts to conduct City work, the records generated may still be subject to search and disclosure. The records generated may include both the content users input and the content users receive from the generative AI system.

Understand the Terms and Conditions. The City does not currently have agreements in place for common generative AI systems, such as ChatGPT or Bing AI. If you choose to use

generative AI for City work and agree to the terms and conditions of a system, you are responsible for complying with those terms and conditions.

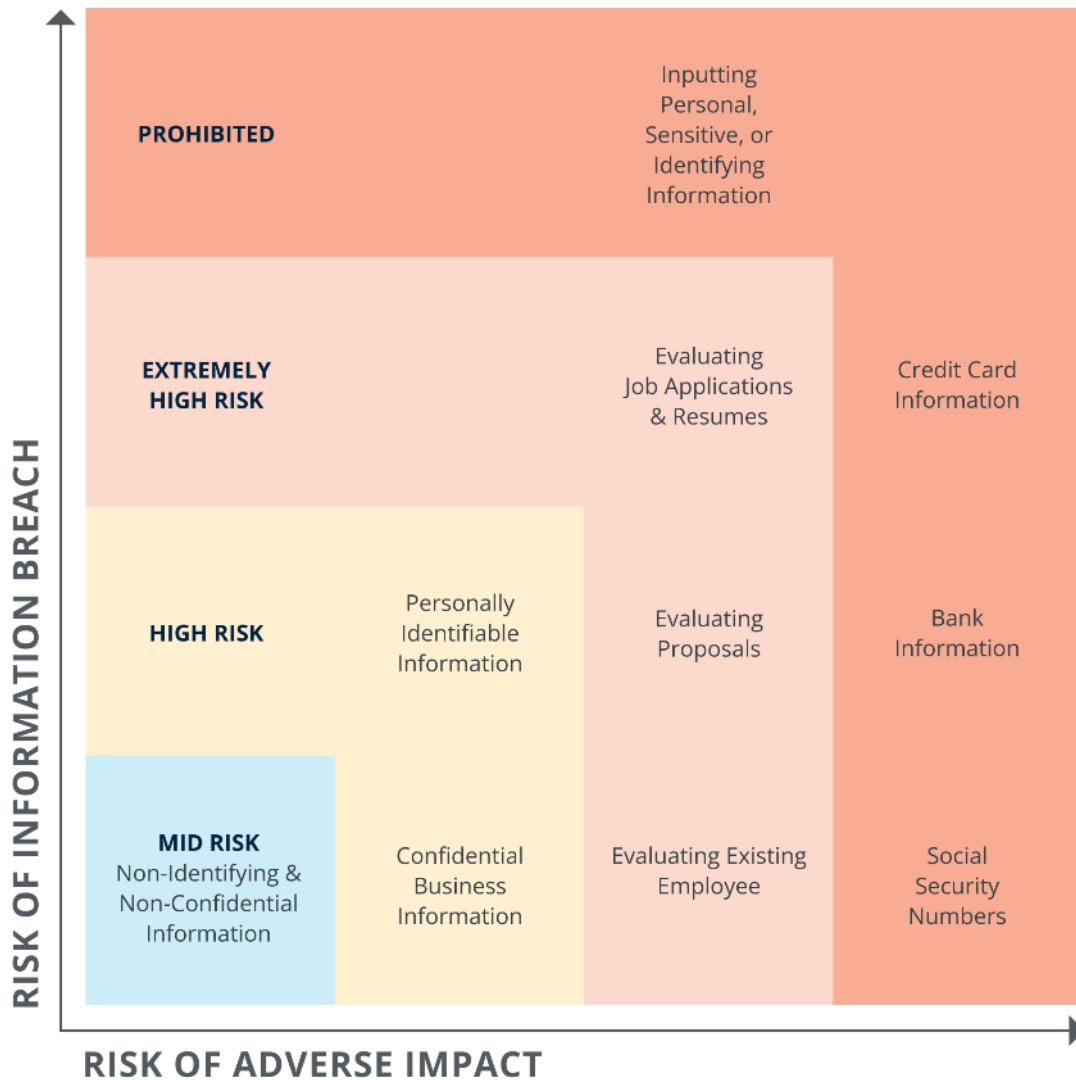
ASSESSING RISK IN USING GENERATIVE AI

The risk presented by generative AI tools varies by use case, with the risk spectrum ranging from mid-risk, high-risk, extremely high-risk, to prohibited risk.

Generative AI risk is determined by two key factors:

1. **Risk of information breach:** the potential harm if the information exchanged with a generative AI system is released to an unintended audience. This can include entering personally identifiable information, sensitive records, or confidential business information into a generative AI system. Additionally, any information entered in a generative AI system may be subject to GRAMA. If you wouldn't share the information in a public forum, don't share it with a generative AI.
2. **Risk of adverse impact:** the potential harm of using the output for a decision, task, or service. This impact can be different for different populations and should be considered from an equity lens, such as adverse impacts to people of a certain race, age, gender identity, or disability status. Not only can AI be biased, but it can also provide false information (hallucinations). In general, if generative AI is used in relation to City processes that can alter an individual or community's rights, freedoms, or access to services, it should be thoroughly reviewed by multiple users before any document is finalized or action is taken.

SUMMARY OF RISK MATRIX OF GENERATIVE AI



When Engaging in High-Risk & Extremely High-Risk Use Cases

Keep in mind the tone and specific language in the AI output. Generative AI is trained in a global context and may not use the vocabulary or tone consistent with the City and its values. **Simple examples include replacing “citizen” with “constituent” in documents, and capitalizing “City” when referring to Salt Lake City Corporation.** These documents, like any others, require thorough review before moving from draft to final product.

Cite verifiable sources for all facts and figures. Generative AI systems are not definitive sources. Facts should be accompanied by links or citations to sources that the public could

find. Generative AI systems can fabricate sources if asked, so do not rely on them for finding citations either. Find sources directly and confirm they are legitimate before using them.

Anything that would not be released or shared with the public should not be inputted into the AI. This includes information such as draft Request for Proposal requirements that should not be public yet, vendor transactions, procurement approvals, or internal City decisions.

Understanding Risk of Information Breach

- **General rule:** If the information exchanged with a generative AI system would be harmful to a person or community if made public, it is a high or prohibited risk. Services can be compromised and leak personal information, all information exchanged with generative AI has a reasonable risk of being compromised.
 - **Mid-risk information** includes non-identifying and non-confidential information. For example, a simple email response or instructive documents often contain only general information that would not present any risk if made public.
 - **High-risk information** includes personally identifiable information (e.g., full name, birth date, email address) and confidential business information that may have larger implications to City processes.
 - **Extremely high-risk information** includes evaluating job applicants or resumes, proposals or an existing employee.
 - **Prohibited risk information** includes highly sensitive and identifying information. This includes data such as credit card numbers, bank account information, social security numbers, and other information that requires rigorous security measures and compliance standards before being processed.

Be Aware of Targeted Cyber Attacks Using Generative AI

Although City staff are already familiar with handling cyber risks like phishing and malware, the advent of generative AI introduces heightened cybersecurity risks as the attacks can be more complex and personalized. Cyber threat actors may use generative AI in their attacks in the following ways:

- **Writing AI-powered, personalized phishing emails:** With the help of generative AI, phishing emails no longer have the tell-tale signs of a scam—such as poor spelling,

bad grammar, and lack of context. Plus, with generative AI like ChatGPT, threat actors can launch phishing attacks at unprecedented speed and scale.

- **Generating deep fake data:** Since it can create convincing imitations of human activities—like writing, speech, and images—generative AI can be used in fraudulent activities such as identity theft, financial fraud, and disinformation.
- **Cracking CAPTCHAs and password guessing:** Used by sites and networks to comb out bots seeking unauthorized access, CAPTCHA can now be bypassed by hackers. By utilizing AI, they can also fulfill other repetitive tasks such as password guessing and brute-force attacks.

Understanding Risk of Adverse Impact

- **General rule:** If you are using generative AI in relation to City processes that can alter an individual or community's rights, freedoms, or access to services, it is at least high-risk if not extremely high-risk and should be thoroughly reviewed before any document is finalized or action is taken. Additionally, any action that could reasonably lead to the City engaging in legal infringements on intellectual property is prohibited.
 - **Mid-risk impact** includes tasks associated with drafting internal messages, internal documentation, and idea generation. These tasks can be sped up with the support of generative AI but require many more steps before reaching a public impact.
 - **High-risk impact** includes tasks associated with official City documents or messaging. It also includes uses that require substantial editing and review before usage. These tasks require thorough review at the time of generation before using in any work context. Special care should be taken when a task may impact individuals differently across factors such as race, age, gender identity and disability (e.g., a memo about tree canopy inequity in neighborhoods).
 - **Extremely high-risk impact** includes tasks that undermine trust in the City through false statements or news; deny people due process such as in resource allocation, job evaluations, and purchasing decisions; or expose the City to substantial security or legal risks. Generative AI does not have reasoning behind the content it produces and cannot justify a decision.
 - **Prohibited risk impact** includes tasks associated with highly sensitive and identifying information. This includes data such as credit card numbers, bank account information, social security numbers, and other information that requires rigorous security measures and compliance standards before being processed.

CONCLUSION

Remember the fundamental rules when using any generative AI:

- 1. Never submit personal or confidential information into a generative AI system.**
- 2. Review, revise, test, and fact check any output from a generative AI.**
- 3. Be transparent when content is drafted using generative AI.**
- 4. Return to this document quarterly, as guidance on usage will change rapidly.**

By keeping the above guidance in mind when using generative AI tools, we can ensure the safe and responsible use of AI for the City.

If you or your department have any questions, comments, or concerns around using generative AI, please contact Innovation@slc.gov. The team will work with you and your team to ensure understanding and compliance of the City's [Use of Generative Artificial Intelligence](#) policy.

APPENDIX A: Use of Generative Artificial Intelligence Policy

52-13E-1: PURPOSE:

This article provides guidance regarding employees, contractors, volunteers, or other persons performing a role for Salt Lake City use of generative artificial intelligence (AI) for City business and/or when using City resources.

52-13E-2: DEFINITIONS:

ARTIFICIAL INTELLIGENCE(AI): Machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments. Artificial intelligence systems use machine- and human-based inputs to perceive real and virtual environments; abstract such perceptions into models through analysis in an automated manner; and use model inference to formulate options for information or action.

ARTIFICIAL INTELLIGENCE SYSTEM: Any system, software, sensor, or process that automatically generates outputs including, but not limited to, predictions, recommendations, or decisions that augment or replace human decision-making. This extends to software, hardware, algorithms, and data generated by these systems, used to automate large-scale processes or analyze large data sets.

GENERATIVE ARTIFICIAL INTELLIGENCE: Branch of AI technology that can generate content—such as generating business letters and other documents that present an argument, summarize information, or analyze information like survey results — at the request of a user. Salt Lake City recognizes the opportunity for a controlled and responsible approach that acknowledges the benefits to efficiency while minimizing the risks around AI bias, privacy, and cybersecurity.

USERS: Employees, contractors, volunteers, or other persons performing a role for Salt Lake City use of generative artificial intelligence for City business and/or when using City resources.

52-13E-3: SCOPE:

This article applies to all use of generative AI by Users.

52-13E-4: RESPONSIBILITIES:

Usage of generative AI shall follow Salt Lake City's AI principles:

A. Privacy: Only submit information to generative AI tools that is ready for public disclosure. This includes any text, photos, videos, or voice recordings shared with the AI. Be mindful that the AI output may include unexpected personal information

from another user, so users must remove any potentially private information before publishing.

a. Anything that would not be released or shared with the public should not be inputted into the AI tools. This includes protected, private, or controlled records, and information such as draft Request for Proposal requirements that should not be public yet, vendor transactions, procurement approvals, or internal City decisions.

b. General rule: If the information exchanged with a generative AI system would be harmful to a person or community if made public, it is a high or intolerable risk. Services can be compromised and leak personal information, all information exchanged with generative AI has a reasonable risk of being compromised. If unsure, do not input into generative AI.

c. Mid-risk information includes non-identifying and non-confidential information. For example, a simple email response or instructive documents often contain only general information that would not present any risk if made public.

d. High-risk information includes personally identifiable information (e.g., full name, birth date, email address) and confidential business information that may have larger implications to City processes.

e. Prohibited risk information includes highly sensitive and identifying information. This includes data such as credit card numbers, bank account information, social security numbers, and other information that requires rigorous security measures and compliance standards before being processed.

B. Accuracy: Salt Lake City maintains trust with its constituents by providing accurate information. Review and fact check all outputs received from a generative AI. Users should consult trustworthy sources to confirm that the facts and details in the AI-generated content are accurate. Be aware that many systems may only use information up to a certain date and cannot guarantee the content they generate is accurate.

a. Trustworthy sources include City reports, records, master plans, codes, policies, and department/division webpages. City staff should also consult their manager, department director or division leadership to verify the accuracy of information.

C. Transparency: The user shall be clear when they use generative AI. This includes citing that AI was used in creating a product.

D. Equity: AI system responses are based on patterns and relationships learned from large datasets derived from existing human knowledge, which may contain

errors and may be historically biased across race, sex, gender identity, ability, and many other factors. Users of generative AI must be mindful of any assumptions generative AI may make based on past stereotypes. Users must edit or disregard biased output.

E. Accountability: The person using AI is accountable for the content it generates. Use generative AI with a healthy dose of skepticism. The level of caution used should correspond to the risk level of the use case. It is always important to verify information provided by generative AI.

F. Beneficial: Users should be open to responsibly incorporating generative AI into their work where it can make services better, more just, and more efficient.

52-13E-5: PRIVACY

A. Opt Out: Some services offer an option to opt out of data collection. This means the generative AI system will not keep the data you provide, and it will not be used in the system's models. Salt Lake City as a customer of AI products, shall opt out of data collection and model training whenever possible.

B. Verify the Copyright of All Generated Content: Users shall verify the content they use from any generative AI systems does not infringe any copyright laws. If users are uncertain if content violates copyright, they should either edit the content to be original or not use it.

a. Visit the U.S. Copyright Office to learn if content is copyrighted:
<https://cocatalog.loc.gov/cgi-bin/Pwebrecon.cgi?DB=local&PAGE=First>

C. Ownership of Generated Content: In most cases, the user owns the content they input into a generative AI service and the information they receive as an output. The user can use the content at their discretion, in accordance with Salt Lake City policy and any terms and conditions they agreed to with the generative AI software. However, many generative AI companies still retain the right to use both the input and output content for their own commercial purposes. For example, this could include a generative AI company using Salt Lake City data to train their models or distributing Salt Lake City output data for marketing campaigns. This emphasizes the importance that only information Salt Lake City is ready to make public should be entered into a generative AI system.

52-13E-6: USE OF GENERATIVE AI

Users are required to follow these rules while using generative AI for Salt Lake City work:

A. Information you enter into generative AI systems could be subject to a Government Records Access and Management Act (GRAMA) request, is a record

requiring retention considerations, may be viewable and usable by the City, and may be leaked in a data breach. Do not submit any information to a generative AI platform that would not be available to the general public (such as confidential, protected or personally identifiable information).

B. Thoroughly review, revise, and fact check any output from generative AI. Users are responsible for any material created with AI support.

C. Users must cite the generative AI when a substantial portion of the content used in the final version comes from the generative AI.

D. Users may not utilize City generative AI accounts for personal use and users may not utilize personal generative AI accounts for official City business.

APPENDIX B: Learn More About Generative AI

Generative AI 101

- **Generative AI Explained in 2 Minutes** (Video 2:02 minutes) <https://youtu.be/rwF-X5STYks>
- **Introduction to Generative AI** (Video 22:07 minutes) <https://youtu.be/G2fqAlgmoPo?feature=shared>
- **What is Generative AI? An AI Explains** (Article): <https://www.weforum.org/agenda/2023/02/generative-ai-explain-algorithms-work/>
- **Percipio** (105 courses available) <https://slc.percipio.com/search?categories=Course&q=generative%20artificial%20intelligence>
- **Responsible AI for Public Professionals: Using Generative AI at Work** (7 module course) <https://innovate-us.org/course/responsible-ai-for-public-professionals-using-generative-ai-at-work>
- **Generative AI 101 Podcast. Host Emily Laird** (Podcast) <https://generativeai101.podbean.com/>
- **The Generative AI Podcast Host Tony Wan** (Podcast) <https://podcasts.apple.com/us/podcast/the-generative-ai-podcast/id1660801320>

Government Professionals and Generative AI

- **The Government is Using AI to Better Serve the Public** (Report): <https://ai.gov/ai-use-cases/>
- **AI for Government: Tips for Using ChatGPT in the Public Sector [Includes Prompts!]** (Article): <https://opengov.com/article/ai-for-government/>
- **A Guide to Responsible and Efficient Use of Generative Tools** (Video, available in English and Spanish, with subtitles, 13:04 minutes) <https://innovate-us.org/generative-ai-explainer>

Writing Prompts for Generative AI

- **Prompting 101: Writing prompts for AI** (Article): <https://www.godaddy.com/resources/skills/writing-prompts-for-ai>
- **Prompts for communicators using the new AI-powered Bing** (Article): <https://blogs.microsoft.com/blog/2023/03/16/prompts-for-communicators-using-the-new-ai-powered-bing/>
- **How to Craft the Perfect ChatGPT Prompt** (Article-requires LinkedIn log-in): <https://www.linkedin.com/pulse/how-craft-perfect-chatgpt-prompt-charlene-li-kowge/>

- **650+ Best Prompts for *ChatGPT*** (Article): <https://www.writingbeginner.com/best-prompts-for-chatgpt/>
- **Engineering prompts to use with *Copilot* for Microsoft 365** (Percipio Video): <https://slc.percipio.com/videos/1c405ff2-7290-4910-a0fe-9bfb797a69ce>
- **50 *ChatGPT* Prompts for State and Local Government** (Article): <https://www.govtech.com/artificial-intelligence/chatgpt-example-prompts-for-state-and-local-government>

Use Cases for Generative AI

- **Generative AI Use Cases in State and Local Government** (Video 26:53 minutes): <https://www.youtube.com/watch?v=kYtVALQo948>

APPENDIX C: Examples of Generative AI Use Cases by Risk Level

Examples of Mid-Risk Use Cases

1. **Drafting messages to staff and trusted partners:** Generative AI tools can help user's draft emails or other messages to staff and trusted partners. You can prompt generative AI tools to provide formal sounding language from general framing of the message. You can also have it draft emails in different tones by asking for a different tone.

You may be inclined to use generative AI to help with email replies. *Do not copy your current email thread into ChatGPT.* The email was sent to select people and may be confidential.

Be mindful about the purpose of the email, and if it is appropriate to use generative AI for drafting it.

2. **Framing written content not intended for official release.** Generative AI can be useful for creating an outline or structure for your written content. This can include an outline for a cover letter, long-form writing, project documentation, or speaker notes for a presentation. When the written content is not intended for official public release, it presents less risk than official City publications (like memos or policies). You can write a few key points you would like to detail, any themes you want to present, the kind of voice you would like, and how long you would like the content. Remember that information you input into a generative AI system may be subject to a GRAMA request.

Unless you have a generative AI trained to your context, the tool will provide generic language that does not apply to the City. For example, generative AI may use the word "citizens" rather than "constituents" when referring to the people we serve because it is not used to Salt Lake City's specific nomenclature. *As always, make sure to review, revise, and fact-check any output from generative AI.*

3. **Learning from a document.** You may copy a large public document into a generative AI tool and ask the AI questions about the document. The document or information you paste into the generative AI tool should already be public information.
4. **Learning about a new topic in a way that you can understand.** For example, "explain capital budgeting to me like I'm five." Verify anything you learn from generative AI before applying the knowledge in a City context.
5. **Helping you find the right word for a concept.** For example, "what is the word for the second-to-last episode in a series." Once the AI provides the word, search the word on Google (or elsewhere) to confirm it means what you think it means.

Examples of High-Risk Use Cases

1. **Drafting memos and or other public-facing City documents.** Generative AI tools can help user's draft memos and other public-facing documents more efficiently. Because the content is meant for public release, it is treated as high-risk and should be reviewed and edited multiple times before release.

The City expects users to produce their own research that informs memos, such as information related to policy changes and program changes. Memos, press releases, and other publications also have their own City-specific formats and standards to follow.

2. **Writing advertisements, social media posts.** Generative AI tools can help users draft promotional material. Because the content is meant for public release, it is a high-risk use case.
3. **Writing job postings or job descriptions.** If you provide a generative AI with a list of qualities you want and a role title, it can help you draft a formal-sounding job description. Because the content is meant for public release, and job requirements can have a substantial impact on who applies, it is a high-risk use case.

The City expects users to follow existing standards on the format and content of job postings based on classifications. Consult your Human Resources business partner for information on job classifications and postings.

Additionally, be mindful of the language used in the requirements, responsibilities, and tone used in the job posting. Check if the job description seems to use language stereotypically associated with a specific race or gender.

- Use gender-neutral language: Avoid using gender-specific pronouns (he, she) and job titles (fireman, firewoman). Instead, opt for inclusive terms such as "they" and "fire fighter."
 - Remove gender-coded words: Avoid using adjectives that may be associated with a specific gender, such as "aggressive" or "nurturing." Use neutral descriptors, like "results-driven" or "collaborative."
4. **Creating images or videos.** Some generative AI tools can create images or video clips based on text prompts. The City needs to maintain its legitimacy as a trustworthy source when using video and images, which requires substantial precautions whenever using AI-generated visual content.

Best Practices:

- **Use only for illustrative purposes.**
 - For historical events, use real images rather than generated.

- For example, if you want a picture of a giraffe wearing a suit and tie for your presentation, generate it. If you are proposing a new visual diagram or abstract concept, you can also generate it. If you want a picture of the Mayor at City Hall, find the actual picture.
 - **Require a citation to always be embedded into the image or video.**
 - Images and videos can easily be taken out of their original context and misinterpreted as reality. To prevent a news article or other secondary source from treating an image as fact, all images and frames of a video should specify they were generated using an AI system. The citation should be included in the image itself, and itself and cannot be removed without editing or cropping the image.
5. **Creating presentation slides.** If you provide generative AI with some public information, it can create a presentation for you. Presentations are automatically high risk because they go beyond text into images, which can present false information if the audience believes the image is real.
 6. **Creating diagrams.** For example, “create a flow chart of a tree turning into wood pulp and then into paper.” Replace pictures before publishing.
 7. **Drafting papers.** For example, “Here is my outline for my research paper, and my findings, draft a complete paper.”
 8. **Programming or coding:** Code generated by an AI may be outdated, copyrighted, have identified vulnerabilities, or rely on other code that no longer works. The generated code is not cited to a date (like a stack overflow post would be), so it is unclear when the code would have been good.

AI can help frame your coding problem, and help you draft pseudo-code to solve your problem conceptually. You can request code snippets to help define syntax and it can be useful for testing projects in a low-risk, non-production environment.

Examples of EXTREMELY High-Risk Use Cases

1. **Evaluations and Decisions:** Evaluating job applicants using generative AI can lead to biased application reviews. This evaluation issue also extends to other areas such as evaluating proposals or an existing employee. AI-based evaluations expose the City to public protest across many key City functions such as hiring and purchasing.

Additionally, generative AI shall not be used to determine highly sensitive decisions such as an individual's health plan, cost of bail, conviction of a crime, grades, or admissions to a program.

AI can help flag key words and identify phrases within a document (see the mid-risk use case). However, the actual evaluation must be made by a person.

2. **Language Translation:** Large Language Models (LLMs) are not yet demonstrably better for translation than other existing translation software. Existing translation software are generally AI systems, but are built for specifically translating text, compared to modern generative AI systems, which attempt to be a more general AI system for more problems.

Future generative AI systems may be substantially better than existing translation AI systems, but they will require an evaluation of their performance before the City should use them.

3. **Creating voice or other audio:** Replicating a person's voice with AI in any City document or recording would undermine the trust of staff and the constituents. Potential legal concerns also exist regarding replicating a person's voice. Use extreme caution when generating audio through AI.

APPENDIX D: Glossary of Generative AI Terminology

ALGORITHM	A series of logical steps through which an agent, typically a computer or software program, turns particular inputs into particular outputs.
ARTIFICIAL INTELLIGENCE (AI)	Machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments. Artificial intelligence systems use machine- and human-based inputs to perceive real and virtual environments; abstract such perceptions into models through analysis in an automated manner; and use model inference to formulate options for information or action.
ARTIFICIAL INTELLIGENCE SYSTEM	Any system, software, sensor, or process that automatically generates outputs including, but not limited to, predictions, recommendations, or decisions that augment or replace human decision-making. This extends to software, hardware, algorithms, and data generated by these systems, used to automate large-scale processes or analyze large data sets.
GENERATIVE ARTIFICIAL INTELLIGENCE	Branch of AI technology that can generate content—such as stories, poetry, images, voice, music, business letters and other documents that present an argument, summarize information, or analyze information like survey results — at the request of a user.
GRAMA	Government Records Access and Management Act or GRAMA is an open records request.
HALLUCINATIONS	Incorrect or misleading results that AI models generate. These errors can be caused by a variety of factors, including insufficient training data, incorrect assumptions made by the model, or biases in the data used to train the model.
MACHINE LEARNING	A branch of AI and computer science that focuses on the using data and algorithms to enable AI to imitate the way that humans learn, gradually improving its accuracy.

PERSONAL IDENTIFICATION INFORMATION (PII)	Any information connected to a specific individual that can be used to uncover that individual's identity, such as their social security number, full name, email address or phone number.
PROMPTS	Natural language the user enters into a generative AI tool that gives the system direction.
USE CASE	Describes how the user of a system interacts with the system to achieve a certain outcome. Use cases should be specific to scenarios that show how a user will interact with the system.
USERS	Employees, contractors, volunteers, or other persons performing a role for Salt Lake City use of generative AI for City business and/or when using City resources.