

Generative AI Acceptable Use Policy

A sample structure for state, local, and federal agencies. Adapt the bracketed items to your jurisdiction and legal review.

HOW TO USE THIS DOCUMENT

This is a structural sample, not legal advice and not a policy you should adopt unchanged. It shows the sections a defensible agency AI policy needs and the decisions each section forces. Items in **blue** are decisions your agency makes. Every policy of this kind should be reviewed by your counsel and your records officer before adoption.

1. PURPOSE AND SCOPE

This policy governs how **[Agency]** employees, contractors, interns, and volunteers may use generative artificial intelligence tools in the course of agency business. It applies to agency-provided tools and to any generative AI tool accessed on agency devices, on agency accounts, or in the performance of agency work, including tools the agency has not procured.

This policy does not govern embedded machine learning features in existing systems of record, which are covered under **[existing IT policy reference]**.

2. DEFINITIONS

- ▼ **Generative AI tool.** Any system that produces text, images, audio, video, or code in response to a prompt.
- ▼ **Agency data.** Any information created, received, or maintained by the agency, regardless of classification.
- ▼ **Approved tool.** A generative AI tool that has completed the intake process in Section 8 and appears on the agency's approved tool list.
- ▼ **Consequential use.** Use whose output informs a decision affecting a member of the public's rights, benefits, services, employment, or legal status.

3. PERMITTED USES

Staff may use approved tools, with the oversight required by the risk tier of the task, for:

- ▼ Drafting, summarizing, editing, and reformatting documents that a staff member reviews before use
- ▼ Research and background synthesis, where outputs are independently verified before being relied on
- ▼ Brainstorming, option generation, and first drafts of internal material
- ▼ Code authoring and review, subject to **[agency secure development standard]**
- ▼ Analysis of non-sensitive agency data the staff member is already authorized to access

4. PROHIBITED USES

Staff may not:

- ✦ Enter data classified at or above [classification threshold] into any tool not approved for that classification
- ✦ Enter personally identifiable information, protected health information, law enforcement sensitive information, or sealed records into any tool, regardless of approval status, unless the tool is explicitly approved for that data category
- ✦ Use any generative AI tool as the sole basis for a consequential decision as defined in Section 2
- ✦ Present AI-generated content as the independent work of a named staff member in any record, filing, or public communication
- ✦ Use unapproved tools for agency work, including free or personal-account versions of tools the agency has approved under an enterprise agreement
- ✦ Disable, circumvent, or misreport the logging and monitoring controls applied to approved tools

5. DATA HANDLING

Before entering any agency data into a generative AI tool, staff must determine its classification under [agency data classification standard] and confirm the tool is approved for that classification. Where no classification exists, the data is treated as the most restrictive category until classified.

Staff must assume that data entered into a tool may be retained by the vendor unless the agency's agreement with that vendor states otherwise in writing. The approved tool list records the retention terms for each tool.

6. HUMAN OVERSIGHT

A named staff member is accountable for every output used in agency business. Review requirements scale with the risk tier assigned under the agency's AI Risk Tier Framework:

- ✦ **Low risk.** Spot review before use. No documentation required.
- ✦ **Moderate risk.** Full review by the accountable staff member before use. Reviewer recorded.
- ✦ **High risk.** Full review plus supervisory concurrence before use, and a documented basis for the decision independent of the AI output.
- ✦ **Prohibited.** Not permitted under this policy regardless of review.

Review means substantive verification of accuracy, not a reading for tone. A reviewer who cannot independently verify a factual claim in an output must remove it.

7. RECORDS RETENTION AND PUBLIC DISCLOSURE

Prompts and outputs created in the course of agency business may constitute records under [applicable records statute] and may be subject to disclosure under [applicable public records law]. Staff must assume that anything typed into a generative AI tool in the course of agency business is discoverable.

The agency's records officer determines which categories of prompt and output are records, their retention schedule, and the system of record in which they are retained. That determination is published at [location] and is reviewed at the cadence in Section 11.

8. TOOL INTAKE AND APPROVAL

No generative AI tool may be used for agency business before completing intake. Intake requires, at minimum:

- ✦ A named business owner and a described use case
- ✦ A risk tier assignment under the agency's framework
- ✦ Confirmation of the data classifications the tool is approved to receive

- Review of the vendor's data retention, training, and subprocessor terms
- Security review under [agency security review process]
- Accessibility review under Section 508 or the applicable state equivalent
- A decision record naming the approver and the date

Requests are decided within [number] business days. Denials state a reason and the conditions under which the tool could be reconsidered.

9. INCIDENT REPORTING

Staff must report to [contact] within [number] hours any instance in which: restricted data was entered into a tool not approved for it; an AI output containing a material error was relied on or released; a tool behaved in a way that suggests a security or privacy failure; or an unapproved tool was used for agency business.

Good-faith reports do not result in disciplinary action. Failure to report does.

10. ROLES AND RESPONSIBILITIES

ROLE	RESPONSIBILITY
[Policy owner]	Owns this policy, decides exceptions, reports adoption and incidents to leadership
Records officer	Determines record status and retention for prompts and outputs
Security	Reviews tools at intake, monitors approved tools, investigates incidents
Privacy officer	Reviews proposed uses touching personal information
Business unit leads	Assign risk tiers, name accountable reviewers, confirm staff training
All staff	Know the tier of the work they do, review before use, report incidents

11. TRAINING, REVIEW, AND EXCEPTIONS

All staff with access to an approved tool complete baseline training before access is granted, and refresher training annually. Staff performing work at high risk tier complete role-specific training before performing that work.

This policy is reviewed [cadence] and whenever the approved tool list changes materially. Exceptions require written approval from the policy owner, a stated duration, and a compensating control.

THE THREE DECISIONS THIS POLICY ACTUALLY TURNS ON

Most agency AI policies stall in review because these are left open. Settle them first and the rest of the document writes itself: **which data classifications may enter which tools; who is accountable when an output is wrong; and whether prompts are records.** Everything above is downstream of those three.