

AI Risk Tier Framework

A sample four-tier model for classifying agency AI use. Tiering is what makes an acceptable use policy enforceable instead of aspirational.

WHY TIERING COMES FIRST

An AI policy that applies the same rules to summarizing a meeting and to screening a benefits application will be ignored for the first and insufficient for the second. Tiering sorts uses by consequence, then attaches proportionate controls. Without it, a policy is a list of prohibitions staff route around.

THE CLASSIFICATION QUESTION

A use is tiered by one question, asked in order:

- ▼ **Does the output affect a member of the public's rights, benefits, services, employment, or legal status?** If yes, the floor is High.
- ▼ **Does the use involve restricted data?** Personal, health, law enforcement sensitive, or legally protected information raises the tier by one.
- ▼ **Can an error be caught before it causes harm?** If an error would reach the public or the record without a review step, raise the tier by one.
- ▼ **Is the output attributed to the agency?** Published or filed material is never Low.

Tiering is assigned to the **use**, not the tool. The same tool may be Low for one task and High for another, and the controls follow the task.

TIER	DEFINITION	REQUIRED CONTROLS	APPROVAL
1 — Low	Internal, non-attributed work. No restricted data. An error costs time, not harm. Drafting internal notes, reformatting, brainstorming.	Spot review before use. Approved tool only. No documentation.	None. Staff discretion.
2 — Moderate	Work that leaves the team or enters the record. No restricted data. Draft correspondence, internal analysis, public-facing copy before editorial review.	Full review by a named accountable staff member. Reviewer recorded. Factual claims independently verified.	Business unit lead, by category.
3 — High	Restricted data, or output informing a decision about a person, or material published under the agency's name. Case summarization, eligibility support, public notices.	Full review plus supervisory concurrence. Documented basis for the decision independent of the AI output. Prompt and output retained per schedule. Periodic sampling audit.	Policy owner, per use case, in writing.
4 — Prohibited	Uses the agency will not permit at any control level. Sole-basis determinations, biometric inference, uses barred by statute or by the agency's own commitments.	Not applicable. Use is barred.	Not available. Changes require policy amendment.

WORKED EXAMPLES

USE	TIER	WHY
Summarizing a recorded internal staff meeting	1	Internal, no restricted data, errors caught in use
Drafting a reply to a routine constituent email	2	Leaves the agency, attributed, but no restricted data and a review step exists
Summarizing case notes containing personal information	3	Restricted data, and output informs a decision about a person
Drafting a public notice published under the agency seal	3	Published and attributed; an error reaches the public as an official statement
Scoring applications to rank or screen applicants	4	Determines a person's access to a benefit or service
Generating code for an internal reporting script	2	No restricted data, but enters the agency's systems and needs review

APPLYING THE FRAMEWORK

Step 1. Inventory what is already happening

Staff are using these tools before the policy exists. An inventory that surfaces actual use, gathered without penalty, is more accurate than a survey of permitted use and gives the tiering real cases to work from.

Step 2. Tier the inventory, not the hypotheticals

Classify the uses you found. Most agencies find that the large majority fall into Tier 1 and Tier 2, which means the policy's practical job is to make those two easy rather than to police them.

Step 3. Attach controls and publish the tier table

Staff should be able to determine the tier of their own task in under a minute, without asking. If they cannot, the tiering is too complex and will be ignored.

Step 4. Review on a schedule and on change

Tiers move. A use that was Tier 2 becomes Tier 3 when the tool gains access to a new data source. Review when the approved tool list changes, not only on the calendar.

THE MOST COMMON FAILURE

Agencies tier the tool rather than the use. Once a tool is "approved," staff treat every task done in it as approved, including tasks the review never contemplated. Tier the use, publish the examples, and the distinction holds.

WHAT THIS FRAMEWORK DOES NOT COVER

This is a classification model. It does not substitute for a privacy impact assessment, a security authorization, a records determination, or an accessibility review, each of which has its own process and owner. The tier assignment tells those processes how much scrutiny a use warrants; it does not replace them.