

User disclosure mechanism

Agencies must disclose to users when they are interacting with AI-generated content.

Disclosure is “no surprises, clear expectations” and it should be user interface (UI)-proximate and role-accurate. For low-risk internal AI knowledge chatbots, disclosure should be a lightweight user interface label so staff understand they’re looking at AI-assisted content.

Disclosure is about expectation-setting, not technical explanation. If someone would reasonably feel misled without disclosure, disclosure is required.

Regular, Low-Risk Use Cases:

Examples: Chatbots, document review tools, voice/ meeting transcription, translation

What to Disclose: That AI was used as assistance, not authority.

Where:

- **In the tool UI** for AI products that surface AI-generated content to a user (Small label, tooltip, or footer note)
- **In the document**, if AI was used for significant content creation beyond editing and proofreading. For example, if AI was used for transcription of an entire meeting or translation of a document, that should be disclosed. If AI was used as a brainstorming partner before content was drafted, or if AI was used to suggest edits to a document primarily written by a human, disclosure is not required. In the middle? Use your best judgment.
- **Content that is external-facing** needs disclosure.

When:

- At the moment the user consumes the output (not later or elsewhere).

How: Use a plain-language pattern:

Use	Potential message
AI-generated content	“AI-assisted draft. Reviewed by staff.” or “Generated with AI and edited by [team/role].” “The script you will hear throughout this training is read using an AI-generated voice. The content itself was written and reviewed by [insert agency initials here] staff.”
Chatbot agent	“AI can make mistakes. Check citations to verify.”

Document review or processing agent	“AI can make mistakes. Please double-check all output.” “Generated with AI and edited by [team/role].”
-------------------------------------	--

What Not to Include:

- No model names.
- No training data claims.
- No legalistic explanations.

Governance check:

Ask: “Would a reasonable person feel misled if they later learned AI was involved?”

If yes → disclose.

For Higher-Risk Use Cases:

User disclosure requirements for higher-risk AI use cases, including the ability for a user to escalate to a human in certain cases, are in development.

NIST AI Risk Management Framework Mapping

- **GOVERN:** documenting and communicating AI risks and impacts more broadly (GOVERN 4.2) and setting risk management controls transparently (GOVERN 1.4).
- **MANAGE:** documenting residual risks to downstream users (MANAGE 1.4) and communicating incidents/errors to relevant AI actors, including affected communities (MANAGE 4.3).

Revision History

Version	Date	Change Log
1.0	June 12, 2026	Initial version