

ORION FORUM

Improving DOJ's AI Use Case Inventory: Four Targeted Reforms

OCTOBER 15, 2026

During a panel at the [DoDIIS Worldwide Conference](#) in August 2026, Federal Bureau of Investigation (FBI) chief AI officer Katie Noyes described 139 approved artificial intelligence (AI) use cases. The Department of Justice's (DOJ) public 2025 inventory, released in January 2026, listed 50 FBI use cases. An FBI spokesperson later [told](#) FedScoop that new approvals would appear in the next annual inventory. The 89-use-case difference highlights a reporting gap: comparing the two counts requires knowing what changed between them.

In an era of rapid [increases](#) in AI capability and performance and a [growing list](#) of AI failures and misuse, the gap matters for more than accurate counts. The public needs to know what newly approved uses do and what safeguards apply. Annual disclosure alone may not keep pace with changes in agency AI use.

To address this, the DOJ, which has over 40 components, including FBI, should publish a dated reconciliation supplement so oversight bodies and concerned citizens can better follow its AI uses from one inventory to the next. The Office of Management and Budget (OMB) should define the fields, and DOJ's chief AI officer should coordinate updates from the components. The supplement should include four elements:

- Use-specific explanations of impact classifications.
- Dates and reasons for all material changes.
- Predecessor and successor identifiers.
- Disclosure-status labels.

With these links in place, the concerned public could better distinguish routine administrative tools from the uses that shape investigations, impact constitutional rights, or inform consequential decisions. They can also better separate how the DOJ actually uses AI from how it reports that use.

Imagine a congressional committee staffer comparing two years of inventory lists. Across inventories, the staffer notices that several use cases have fallen off. Without time-consuming follow-ups with the agency, the staffer reviewing the list has no way of knowing if the missing entries were discontinued, renamed, or folded into an existing entry.

Impact classifications confront the same issue. Any evaluator examining if a use case received the proper risk classification would need a keen understanding of what the tool does, how it's used, and if it influences decisions. Oversight staff must then seek basic information about the tool before they can assess its classification.

OMB already has authority to set the inventory's public format. Section 7225 of the Advancing American AI Act [directs](#) OMB to require inventories of current and planned uses and determine how agencies publish them, subject to protections for privacy, sensitive law-enforcement information, and national security. The inventory requirement [excludes intelligence-community elements](#), including [the FBI's intelligence elements](#).

DOJ [explains](#) that it combines similar uses, removes duplicate or mislabeled entries, and changes listings after vendor mergers. Some responses were standardized while guidance evolved. But these general explanations do not identify what happened to each omitted entry. Fewer entries could reflect consolidation rather than less AI use.

The first element, use-specific explanations of impact classifications, would require DOJ to explain how each tool's actual use supports its classification. The Trump administration's OMB Memorandum M-25-21 automatically [presumes](#) certain AI applications are high-impact, triggering specified safeguards. An agency can rebut the presumed classification via a written determination explaining why a certain use does not meet the definition. For high-impact use cases, establishing human oversight alone does not remove a use from that category. Section 4(a)(iv) [requires](#) agencies to publish summaries of those determinations and any waivers of

safeguards, including their justifications, subject to disclosure restrictions.

The [2025 inventory](#) lists a Bureau of Prisons translation use (DOJ-0153) and a Drug Enforcement Administration (DEA) proposal to create IT security packages (DOJ-0357) as presumed high-impact use cases but determined otherwise. Even though the uses described different work, both provided the same justification: their outputs are not a principal basis for consequential decisions. Even taking both of these classifications as correct, neither justification explains how the tool's role supports that conclusion.

Borrowing a counterexample from the Department of Homeland Security (DHS), an Immigration and Customs Enforcement inventory entry for AI-Assisted Audio/Video Redaction ([DHS-2511](#)) is useful. The entry states specifically that the tool hides sensitive information and does not perform facial matching or identification. All redactions also undergo human review. These additional details help to directly tie the tool's classification to its role instead of just repeating the classification test.

The supplement would require DOJ to explain how the AI tool is used, which decisions it influences, and why a presumed high-impact use was classified otherwise.

The second element, dates and reasons for material changes, would note and record developments after publication. DOJ's [inventory page](#) now displays an October 7, 2026, update date but does not identify any corresponding changes to individual entries. DOJ's 2025 inventory, published in January 2026, lists both pre-deployment testing and an AI impact assessment as "In-progress" for the FBI's Facial Recognition Technology use (DOJ-0122). Now in October, that status does not establish whether the FBI completed those safeguards by OMB's April 3, 2026 [deadline](#) to implement and document minimum risk-management practices for high-impact AI. OMB's instructions [encourage](#) ongoing updates and require revisions, as appropriate, to reflect compliance with that deadline. In that spirit, this element would record the date and reason for each status change, including completion, a waiver, or discontinued AI functionality.

A tool's capabilities or role in agency decisions can change without increasing the inventory's total. M-25-21 [requires](#) impact assessments for high-impact AI to specify reassessment

procedures after significant system modifications. The supplement would date and describe those changes and any resulting revisions to classifications or safeguards.

The third element, predecessor and successor identifiers, would connect old and new entries when a use-case gets renamed, combined with another entry, or replaced. A historical Federal Emergency Management Agency (FEMA) inventory page demonstrates how. The page identifies FEMA OCFO GPT ([DHS-250](#)) as consolidated into the OCFO Response Augmentation Suite (DHS-2296). The old entry names its replacement and explains the consolidation, instead of leaving the person to reconstruct the change. Researchers have previously [proposed](#) persistent identifiers and cross-references across federal AI disclosures.

The fourth element, disclosure-status labels, would explain why use-case information is missing and fields are blank. Each blank field would carry a label, stating whether the information was withheld, did not apply, or was not reported. Each DOJ component would be separately responsible for explaining which use cases were withheld entirely, which were outside the reporting scope, and which were combined into other entries, providing counts when available and releasable. It would also identify the legal or policy basis for withholding or exclusion. OMB's existing [disclosure rules](#) would still apply, so protected information can still be withheld, but releasable details must be published.

Under the proposal, OMB would define the fields in its next instructions, and DOJ's chief AI officer would coordinate the first supplement with DOJ's next inventory. Components should report the following material changes for public release within 30 days: deployments, retirements, consolidations, impact reclassifications, and material changes in capabilities or decision-making roles, not routine software patches. The public supplement should retain existing use-case identifiers and justification fields. Updates would require staff time and disclosure review; the supplement would remain an agency account, not an independent audit.

DOJ's AI use-case inventory next year will provide the public a new total. Unless changes are made, the public and oversight bodies will struggle to make useful sense of year-over-year changes. Given the fast pace of AI capabilities and tools, they will be stuck reconstructing yesterday's inventory while trying to understand how DOJ uses AI today.

Disclaimer: Orion Policy Institute (OPI) is an independent, non-profit, tax-exempt think tank focusing on a broad range of issues at the local, national, and global levels. OPI does not take institutional policy positions. Accordingly, all views, positions, and conclusions represented herein should be understood to be solely those of the author(s) and do not necessarily reflect the views of OPI.