Trusted Data and Technology Institute
TECHNOLOGY WORTHY OF TRUST

RESEARCH PAPER 2026-10-01 · AI GOVERNANCE & PUBLIC ACCOUNTABILITY

Incidents Without Inventories: Do Government AI Failures Appear in Public AI Disclosures?

An independent test linking U.S. government incidents in the AI Incident Database to federal, state and local AI use-case inventories

Author
Joseph Atalla, PhD

Published
October 2026

ABSTRACT

Governments now publish inventories of the AI systems they use, and the public separately documents the AI systems that cause harm. If disclosure works, the second list should be visible in the first. We test that proposition. From the full AI Incident Database (AIID) export of 12 October 2026 (1,734 incidents), we identify 113 incidents in which a U.S. government body deployed the AI system, code each by level, technology and type of use, and search the 2023, 2024 and 2025 federal AI use-case inventories for every federal system involved. Of 13 civilian federal systems linked to incidents, 6 appear by name in at least one inventory, 3 appear only partially, and 4 never appear. Two-thirds of all incidents, however, involve local police, sheriffs and school districts, and 72 of the 85 state, local and court incidents (85%) occurred where the deploying body has no public AI inventory duty even today. A further 11 federal incidents arise from staff use of general-purpose generative AI or AI-made images in official work, a pattern inventories are not designed to see. We propose six fixes that connect incident reporting to disclosure.

113	4 of 13	85%	2.3×
AI incidents in which a U.S. federal, state, local or court body deployed the system (6.5% of all AIID incidents)	Civilian federal AI systems with recorded incidents never appear in any federal inventory, 2023–2025	of 85 state, local and court incidents involve a body with no duty to publish an AI inventory, even as of October 2026	Overrepresentation of discrimination harms in government incidents versus all other incidents (29% vs 13%)

About TDTI. The Trusted Data and Technology Institute is a California nonprofit public benefit corporation that advances data and technology quality, integrity, security and responsible use through research, education and practical collaboration, with a focus on trusted data, responsible AI evaluation, cybersecurity assurance and public-sector technology. Contact: info@trusteddti.com · trusteddti.com · Murrieta, California

Suggested citation. Atalla, J. (2026). *Incidents Without Inventories: Do Government AI Failures Appear in Public AI Disclosures?* TDTI Research Paper 2026-10-01. Trusted Data and Technology Institute.

Executive summary

AI use-case inventories are the government's own account of the AI it uses. Incident databases are the public's account of the AI that went wrong. The two have never been compared. This paper asks a simple question with practical consequences: when a government AI system harms someone badly enough to be documented, had the government already told the public that the system existed and how risky it was?

KEY FINDINGS

1 Government AI incidents are mostly local, and mostly policing. Of 113 U.S. government-deployer incidents, 71 (63%) involve local bodies and 64 (57%) involve police, sheriffs or immigration enforcement. Facial recognition (27) and license-plate readers (18) lead all technologies.

2 Federal disclosure is real but incomplete. Of 13 civilian federal systems with incidents, 6 are listed by name in at least one inventory and all 6 carry a rights-, safety- or high-impact flag. Four (the IRS audit-selection model, USCIS's ATLAS screening, USGS ShakeAlert and a DOGE contract-review tool used at VA) never appear.

3 Disclosure is unstable over time and inconsistent across components. CBP One was listed as rights- and safety-impacting in 2024 and is absent by that name in 2025. In 2025, the Bureau of Prisons rated the PATTERN risk tool "not high-impact" while the Office of Justice Programs rated the same tool "high-impact."

4 Most incidents happen where no disclosure is required. Of 85 state, local and court incidents, only 4 (5%) involved a body that was required to inventory its AI at the time, all in New York City. Another 9 gained such a duty later. The remaining 72 (85%) involve police departments, sheriffs, school districts, state agencies and courts that, as of October 2026, have no public AI inventory obligation.

5 A new class of incident sits outside inventory design. Eleven federal incidents since 2025 come from staff using general-purpose generative AI to write reports, briefs and use-of-force narratives, or from AI-made images in official communications. Inventories record tools, not these uses.

What should change. OMB should require agencies to reconcile their inventories against publicly reported incidents and to keep retired or renamed systems visible (R1–R2), and should require one risk determination per system across components (R3). Agencies should govern and log general-purpose generative AI used in official outputs (R4). States should extend disclosure to the local police and school systems where most incidents occur (R5). Incident databases should record whether a government system was disclosed (R6).

1 Why compare incidents with inventories

Federal agencies have inventoried their AI use cases since Executive Order 13960 (2020) and the Advancing American AI Act, with OMB consolidating the results annually.[1–3] OMB Memorandum M-24-10 required agencies to identify "rights-impacting" and "safety-impacting" AI and apply minimum risk practices; M-25-21 (April 2025) replaced those categories with "high-impact" AI.[4,5] Some states and cities have followed. New York City has required agencies to report their algorithmic tools annually since 2022, and Vermont, Connecticut, Texas, California and New York State require inventories of state-agency automated decision systems or AI, although Texas reports go to state bodies rather than the public.[16,19–23]

Inventories are only as useful as they are complete, and completeness is hard to test because no one outside an agency knows the full list of its systems. Incidents provide an outside sample. The AI Incident Database (AIID), maintained by the Responsible AI Collaborative, has catalogued real-world AI harms and near-harms from news and other public reports since 2020.[8,9] An AI system that has drawn enough public attention to become an incident is, almost by definition, one the public has an interest in knowing about. Whether such systems appear in inventories is a direct, external check on disclosure.

Research questions

RQ1. What does the incident record show about where and how U.S. governments' AI causes harm? **RQ2.** Were the federal systems involved in incidents disclosed in the 2023–2025 federal inventories, and with what risk designation? **RQ3.** How

many state and local incidents occurred where the deploying body is subject to any public AI inventory requirement? **RQ4.** Which kinds of incidents fall outside what inventories are built to capture?

Prior work and contribution

GAO found in December 2023 that only 5 of 20 agencies provided comprehensive inventory data.[6] An independent 2026 audit tested the evidence behind 2025 high-impact determinations but did not use incident data.[7] Researchers have mapped AIID incidents to risk taxonomies and to regulatory provisions,[10,11] and commercial vendors map incidents to use-case categories, but we found no study that matches incidents to the specific inventory entries governments publish. To our knowledge, this is the first such test at the federal level, and the first count of government AI incidents by jurisdiction against disclosure obligations. We searched the web, GAO and arXiv in October 2026; a formal systematic review was not performed.

2 Data and methods

Source	Coverage	Records	Used for
AI Incident Database, full Excel export[9]	Incidents 1–1747, exported 12 Oct 2026; MIT, CSET and GMF taxonomy fields	1,734	Deployer, date, description, risk domain
OMB 2023 consolidated inventory (as published in the 2024 repository)[1]	Raw consolidated file	710	Federal disclosure, 2023
OMB 2024 Federal AI Use Case Inventory (v2)[1]	41 agency submissions	2,133	Disclosure; rights/safety-impacting flag
OMB 2025 Federal Agency AI Use Case Inventory[2]	Individually reported use cases (plus 900 consolidated commercial entries)	3,611	Disclosure; high-impact flag and justification
State and local inventory laws and reports[16–25]	Statutes, ordinances and published reports	—	Coding disclosure obligations by jurisdiction

Identifying government incidents. We reviewed every distinct entity in the AIID "deployer" field and hand-coded 128 entity names as U.S. federal, state, local or court bodies. An incident is included when a U.S. government body is listed as a deployer and used the AI system in its official function, including an employee's misuse of an agency system (for example, an officer querying a license-plate reader for personal reasons). A title and description sweep added five incidents, including two COMPAS incidents in which AIID lists the developer rather than the courts that used the tool; these are flagged. We excluded incidents where government was only a victim or target (deepfakes of officials, scams impersonating the IRS), incidents abroad, one story AIID itself describes as apocryphal, and one in which employees circulated a deepfake of a colleague. Public universities and public hospitals were excluded because AIID does not distinguish them reliably from private institutions. The result is 113 incidents.

Coding. Each incident was coded by level, by technology (seven classes, keyword rules with manual review), and by type of use: **A**, an agency-operated AI system; **B**, staff use of a general-purpose generative AI tool in official work product; and **C**, AI-generated or altered media in official communications. Five federal incidents involve defense or national-security systems, which are outside the civilian inventories.[4,5] Risk domains are AIID's MIT AI Risk Repository classifications.[10]

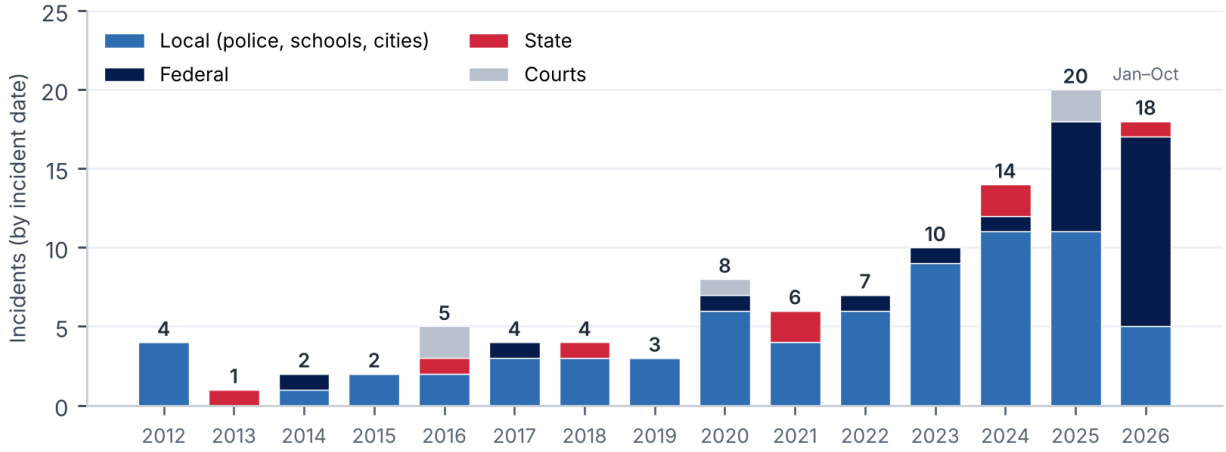
Matching. For each civilian federal type-A system, we searched all three inventories by system name, known aliases and function, within the responsible agency, and recorded each year as *listed* (an entry names the system), *partial* (a related entry exists but describes a different stage or a predecessor), or *not found*. We recorded the 2024 rights/safety flag and the 2025 high-impact flag and justification. For state and local incidents, we coded whether the deploying body is subject to a public AI or algorithm inventory requirement, at the time of the incident and today. **Statistics.** Group differences use Fisher's exact test; small counts are reported as counts. All coding is released (Section 6).

3 Findings

3.1 The incident record: local, policing-heavy and growing

The 113 incidents span 2003 to October 2026, and 83 (73%) date from 2020 or later. Local bodies account for 71 incidents (63%), federal agencies 28 (25%), states 8 (7%) and courts 6 (5%). Police, sheriffs and immigration enforcement deployed the system in 64 (57%). Federal incidents are recent: 21 of 28 date from 2023 onward, and 12 from 2026 alone.

Figure 1. U.S. government-deployer AI incidents by year and level of government

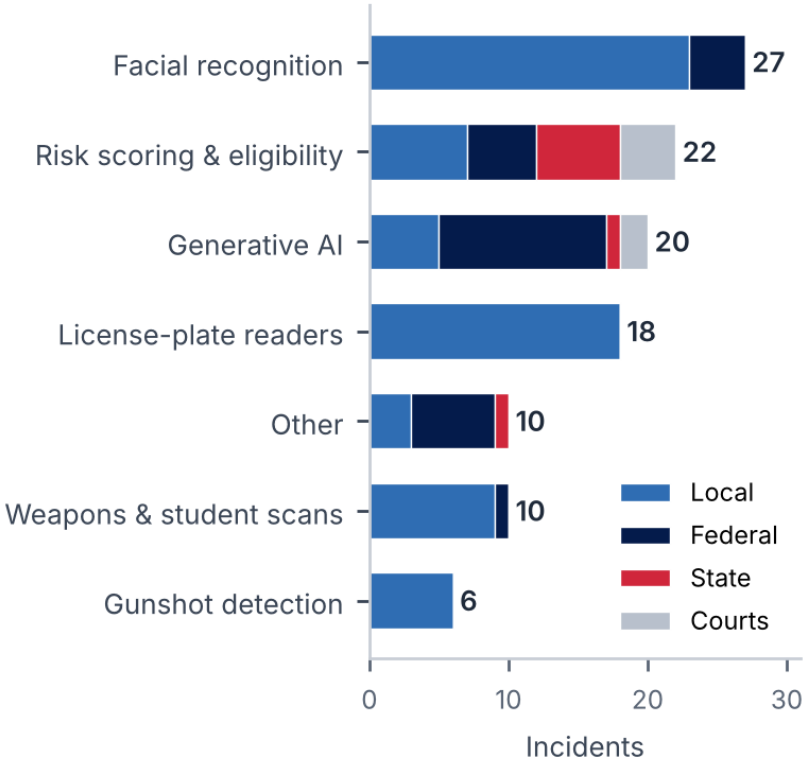


Source: AIID export of 12 Oct 2026; TDTI coding (n = 113; five incidents dated 2003–2009 not shown). Year is AIID's incident date, which may precede the first report. 2026 is partial.

The technology mix differs sharply by level. Facial recognition (27 incidents), license-plate readers (18) and gunshot detection (6) are almost entirely local. Of the 18 license-plate-reader incidents, 8 involve officers using the system to track former partners or other people for personal reasons; the other 10 are misreads or stale alerts that led to stops, several at gunpoint. Weapons-detection and student-monitoring tools account for 10 incidents, 8 of them in school districts. Risk-scoring and eligibility algorithms (22) appear at every level. Generative AI (20) is concentrated in federal incidents since 2025.

Government incidents also look different from other incidents. Using AIID's MIT risk domains, 29% of government incidents are classed as discrimination and toxicity, against 13% of all other incidents (Fisher's $p < 0.001$), and 34% as AI system failures, against 17% ($p < 0.001$). Only 4% are malicious misuse, against 35%. Government AI incidents are disproportionately about systems that work as designed but treat people unequally or fail them.

Figure 2. Incidents by technology and level



Source: AIID export; TDTI technology coding (n = 113).

3.2 Federal systems: six listed, three partial, four never disclosed

Of the 28 federal incidents, 13 involve a civilian agency-operated AI system that the inventories are meant to cover. Figure 3 shows what we found.

Figure 3. Were federal AI systems with incidents disclosed in the federal inventories?

AIID incident · year		2023 inventory	2024 inventory	2025 inventory
#461 · 2008	Audit selection (IRS)	Not found	Not found	Not found
#260 · 2014	ATLAS vetting (USCIS)	Not found	Not found	Not found
#531 · 2017	Body-scanner ATR (TSA)	Not found	Partial	Partial
#532 · 2020	Asylum translation (USCIS)	Not found	Partial	Partial
#154 · 2022	PATTERN risk tool (DOJ/BOP)	Not found	Listed	Listed
#484 · 2023	CBP One face app (CBP)	Partial	Listed	Not found
#699 · 2024	REACH VET (VA)	Not found	Listed	Listed
#1103 · 2025	Contract-review LLM (DOGE/VA)	n/a	n/a	Not found
#1303 · 2025	ShakeAlert (USGS)	Not found	Not found	Not found
#1353 · 2025	Mobile Fortify (ICE)	n/a	n/a	Listed
#1343 · 2026	Résumé screening (ICE)	n/a	n/a	Listed
#1362 · 2026	Field face recognition (CBP)	n/a	n/a	Listed
#1390 · 2026	Observer surveillance (DHS)	n/a	n/a	Partial

Source: AIID; OMB 2023, 2024 and 2025 inventories; TDTI matching. "n/a" means the system or incident postdates that inventory. "Partial" means a related entry exists (for example, TSA lists automated target recognition only as development work). The Mobile Fortify match for #1362 is inferred, because the incident report does not name the application.

Disclosed systems were flagged as risky. Every one of the six listed systems carried a rights- or safety-impacting flag in 2024 or a high-impact flag in 2025. Where disclosure happens, risk classification largely works. The ICE résumé-screening tool and the Mobile Fortify facial recognition app were both listed as deployed high-impact systems in the 2025 inventory, before the 2026 incidents involving them.

Four systems never appear. The IRS examination-selection model at the center of 2023 research finding that Black taxpayers were audited at 2.9 to 4.7 times the rate of others[12] is not among 203 IRS and Treasury entries across the three inventories. USCIS's ATLAS screening, reported in 2021 to flag naturalized citizens for possible denaturalization,[13] the USGS ShakeAlert detection pipeline and the large-language-model tool used to rate VA contracts for cancellation also do not appear. Each absence may have an explanation (ATLAS may be classed as rules-based; ShakeAlert's AI status is not established), but none is explained in the inventories themselves.

Case A: One tool, two risk verdicts

PATTERN, the Bureau of Prisons' recidivism-risk tool, is listed twice in 2025. The Bureau of Prisons entry says it is "not high-impact" because it "does not produce an output that serves as a principal basis for decisions." The Office of Justice Programs entry for the same tool says it is high-impact. Under the First Step Act, a prisoner's PATTERN risk level affects whether earned time credits can be applied toward release.[14] Both components rated it rights-impacting in 2024.

AIID #154. 2025 inventory rows 1269 and 1313; 2024 rows 1915 and 1959.

Case B: Disclosed, then gone

CBP One, whose facial-liveness check was reported in 2023 to fail more often for Black asylum seekers, was listed in the 2024 inventory as both rights- and safety-impacting, with an implementation date of 11 May 2023. The 2025 inventory has no entry named CBP One. DHS replaced the app with "CBP Home" in March 2025, and no 2025 entry carries either name; the closest is a CBP liveness tool called "CBP Link."[15] An incident record that points to a system the current inventory no longer names cannot be followed up from public data.

AIID #484. 2024 inventory row 58; 2025 inventory searched by name and function.

Disclosure often follows harm rather than preceding it. The VA's REACH VET suicide-risk model has been in operation since November 2016, according to its own 2024 entry, but first appears in the 2024 inventory, the same year the incident was reported. PATTERN, CBP One and REACH VET are absent from the 2023 consolidated file we examined. The small number of systems means these patterns are descriptive, not estimates of a rate; the 95% interval around 4 of 13 undisclosed runs from 13% to 58%.

3.3 Generative AI in official work: a gap by design

Eleven federal incidents are not about agency AI systems at all. Six involve staff using general-purpose generative AI in official work: fabricated citations in the HHS "Make America Healthy Again" report and in a Department of Justice brief, a Border Patrol use-of-force narrative that body-camera footage reportedly contradicted, DOGE's reported reliance on ChatGPT to cancel humanities grants, a National Weather Service forecast map with invented town names, and a special-operations analyst's chatbot query that reportedly misidentified ship cargo. Five involve AI-generated or altered images and video posted by the White House, the President, the State Department and the FBI. Two federal judges' chambers and one state judge produced rulings with AI-fabricated content in 2025.

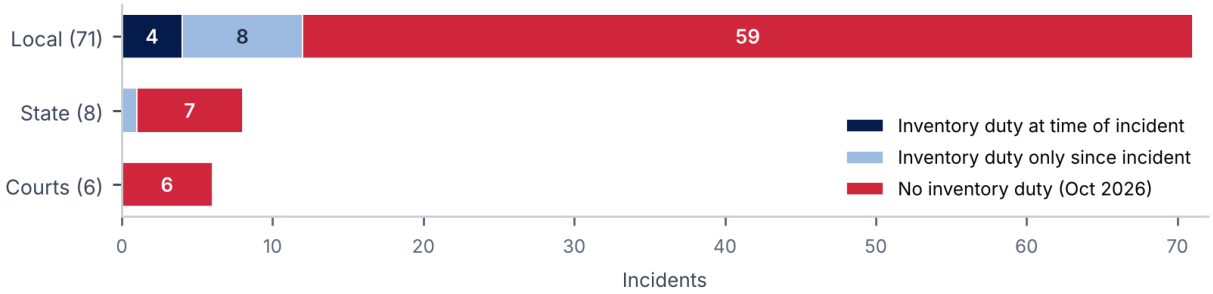
Inventories can list the enterprise tool: the 2024 DHS inventory, for example, lists "Commercial Generative AI for Text Generation." They are not designed to record that a tool was used to draft a particular report or post a particular image. These incidents are therefore not inventory failures, but they show that the inventory model, built around discrete systems, misses the fastest-growing source of federal AI incidents in our data.

3.4 State and local: most incidents occur outside any disclosure regime

We coded each of the 85 state, local and court incidents by whether the deploying body is subject to a public AI or algorithm inventory requirement. The state statutes we identified in Vermont, Connecticut, Texas, California and New York State apply to state executive agencies, not to cities, counties, police departments, school districts or courts.[19–23] A search

of the other states in our data (including Washington, Michigan, Oregon, Arkansas, Nevada, Alaska, Louisiana, Florida, Illinois and Kansas) found no enacted inventory statute.[25] The only local inventory requirement that applies is New York City's Local Law 35 of 2022.[16]

Figure 4. State, local and court incidents by the deploying body's inventory obligation



Source: AIID; TDTI coding of statutes and ordinances as of October 2026 (n = 85). The one state incident with a later duty is California's 2021 vaccine-allocation algorithm, now within AB 302. Surveillance-technology ordinances are not counted as inventory duties.

Only 4 incidents (5%) involved a body that was required to inventory its AI when the incident occurred, all New York City agencies after 2022. Nine more, including California's 2021 vaccine-allocation algorithm and earlier NYPD facial recognition cases, would fall under a duty that exists today. The remaining 72 (85%) involve bodies with no inventory duty even now. They include all 18 license-plate-reader incidents, all 6 gunshot-detection incidents, all 8 school-district surveillance incidents and 15 of 23 local facial recognition incidents.

Where disclosure exists, it shows both its value and its limits. New York City's report for calendar year 2023 lists NYPD facial recognition and the MyCity chatbot, the tools in 9 of the city's 12 incidents.[17] The subway weapons scanners piloted in 2024 postdate that report, and we could not confirm whether the 2025 report lists them. California's first AB 302 report found that state agencies reported no high-risk automated decision systems at all, a result critics attributed to self-reporting; a later report for program year 2025 lists eight.[20] A separate, older channel covers some police technology: 26 jurisdictions, including New York City under its POST Act, have adopted surveillance-technology ordinances requiring public use policies, but these are not AI inventories and cover few of the places in our data.[18,24]

3.5 Incident records rarely say a system was public

AIID's CSET taxonomy includes a "public sector deployment" field, but it is filled for only 200 of 1,734 incidents. Only 18 of our 113 government incidents carry it. No field in either record links an incident to an inventory entry or an inventory entry to a known incident, so every match in this paper had to be made by hand.

4 Discussion

Inventories work best where they are least needed. Federal inventories captured most of the high-profile federal systems in our sample and labeled them correctly. But the federal government accounts for a quarter of incidents. The systems that most often generate incidents, local facial recognition, license-plate readers, gunshot detection and school surveillance, are deployed by bodies that, in most of the country, have no obligation to disclose them at all.

The unit of disclosure is drifting from the unit of harm. Inventories describe systems. The newest federal incidents involve people using general-purpose tools in ways no system entry describes, and official communications that use AI-generated media. Governance for this class of use lives in acceptable-use policies, records retention and labeling, not inventories.

Incidents are a free audit signal. Each incident in this paper was public before we found it. A routine reconciliation, in which an agency or OMB checks each new government incident against the inventory and records the result, would have flagged every absence and inconsistency we report, at low cost.

5 Recommendations

R1 Reconcile inventories against incidents. OMB should require agencies, as part of annual inventory certification, to review publicly reported incidents involving their AI (including AIID entries) and state for each whether the system is inventoried, under which use-case ID, and if not, why.

R2 Keep systems visible across years. Inventories should carry stable use-case identifiers, record renamed, merged and retired systems with their predecessor IDs, and keep retired high-impact entries visible for at least three years.

R3 One system, one risk verdict. Where several components report the same system, OMB should require a single, reconciled high-impact determination, and publish the justification when components disagree.

R4 Govern generative AI in official outputs. Agencies should require disclosure when generative AI materially contributes to a published report, legal filing or enforcement record, and label AI-generated or altered media in official communications.

R5 Extend disclosure to where incidents happen. States should require public inventories, or surveillance-technology reports, from local law enforcement and school districts, starting with facial recognition, license-plate readers, gunshot detection and student-monitoring tools.

R6 Record disclosure status in incident databases. Incident databases should add fields for government level and inventory status (listed, not listed, use-case ID), so that the public record can be joined to official disclosures.

6 Limitations and reproducibility

Incidents are not a random sample. AIID reflects what the press and contributors report, so it overrepresents visible, controversial and recent uses, and English-language coverage. Counts describe the incident record, not the true rate of harm, and growth over time partly reflects AIID's own growth. **Allegations.** AIID entries summarize reported allegations; we did not verify the underlying events. **Coding judgment.** Inclusion, technology and matching decisions are the author's; the rules are stated above and every decision is in the released data. A second coder was not used. **Small federal sample.** Thirteen systems support description, not inference. **Inventory files.** The 2023 consolidated file in OMB's repository contains 710 entries and may not match every agency's separately published 2023 inventory. A system may be inventoried under a name we did not anticipate; we searched names, aliases and functions within each agency. **Law coding.** State and local obligations were coded from statutes and official reports available in October 2026 and are not legal advice.

Replication package. The AIID export, the three inventory files, all scripts, the coded incident set (us_gov_incidents_coded.csv), the federal matching table with inventory row references (federal_inventory_match.csv), and the jurisdiction coding accompany this paper.

References

1. Office of Management and Budget. *2024 Federal AI Use Case Inventory* (consolidated data, data dictionary, README as of 23 Jan 2025). Released 18 Dec 2024. https://github.com/ombegov/2024-Federal-AI-Use-Case-Inventory
2. OMB. *2025 Federal Agency AI Use Case Inventory* (3,611 use cases, 56 agencies). Released April 2026; see Nextgov, 16 Apr 2026. https://github.com/ombegov/2025-Federal-Agency-AI-Use-Case-Inventory
3. Executive Order 13960, *Promoting the Use of Trustworthy Artificial Intelligence in the Federal Government* (3 Dec 2020), and the Advancing American AI Act, as cited in reference 1.
4. OMB. *M-24-10, Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence*, 28 Mar 2024, §§2(c), 3(a), 6. https://www.whitehouse.gov/wp-content/uploads/2024/03/M-24-10-Advancing-Governance-Innovation-and-Risk-Management-for-Agency-Use-of-Artificial-Intelligence.pdf
5. OMB. *M-25-21, Accelerating Federal Use of AI through Innovation, Governance, and Public Trust*, 3 Apr 2025. https://www.whitehouse.gov/wp-content/uploads/2025/02/M-25-21-Accelerating-Federal-Use-of-AI-through-Innovation-Governance-and-Public-Trust.pdf
6. U.S. Government Accountability Office. *Artificial Intelligence: Agencies Have Begun Implementation but Need to Complete Key Requirements*, GAO-24-105980, 12 Dec 2023. https://www.gao.gov/products/GAO-24-105980
7. Shah, I. *fed-ai-inventory-audit*. GitHub, 2026 (not peer reviewed). https://github.com/ishaanshah101/fed-ai-inventory-audit
8. McGregor, S. (2021). Preventing Repeated Real World AI Failures by Cataloging Incidents: The AI Incident Database. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(17), 15458–15463. arXiv:2011.08512.
9. Responsible AI Collaborative. *AI Incident Database*, full Excel export (Incidents, Reports, Entities, taxonomies), 12 Oct 2026. Incident pages at https://incidentdatabase.ai/cite/[ID]
10. Slattery, P., Saeri, A. K., Grundy, E. A. C., Graham, J., et al. (2024). The AI Risk Repository. arXiv:2408.12622.
11. Mumtaz, U., Noor, R., and Mumtaz, S. (2026). Post-Deployment Accountability in AI Governance: A Cross-Regulatory Empirical Analysis of AI Incidents. arXiv:2605.16281.
12. Elzayn, H., Smith, E., Hertz, T., Ramesh, A., Fisher, R., Ho, D. E., and Goldin, J. (2023). *Measuring and Mitigating Racial Disparities in Tax Audits*. Stanford SIEPR working paper. https://siepr.stanford.edu/publications/working-paper/measuring-and-mitigating-racial-disparities-tax-audits
13. Biddle, S., and Saleh, M. Little-Known Federal Software Can Trigger Revocation of Citizenship. *The Intercept*, 25 Aug 2021. https://theintercept.com/2021/08/25/atlas-citizenship-denaturalization-homeland-security/
14. 18 U.S.C. § 3624(g); 28 C.F.R. § 523.44 (application of First Step Act time credits; minimum or low recidivism risk). https://www.ecfr.gov/current/title-28/chapter-V/subchapter-B/part-523/subpart-E
15. DHS Launches CBP Home App with Self-Deport Reporting Feature, as reported by Nextgov, 10 Mar 2025. https://www.nextgov.com/digital-government/2025/03/dhs-releases-new-mobile-app-pushing-migrants-self-deport/403619
16. New York City Local Law 35 of 2022 (Admin. Code § 3-119.5, annual reporting on algorithmic tools). https://intro.nyc/local-laws/2022-35
17. NYC Office of Technology and Innovation. *Algorithmic Tools Reporting*, calendar year 2023 (updated) and CY2025 report (March 2026). https://www.nyc.gov/assets/oti/downloads/pdf/reports/2023-algorithmic-tools-reporting-updated.pdf
18. New York City Local Law 65 of 2020 (Public Oversight of Surveillance Technology Act). https://intro.nyc/local-laws/2020-65
19. New York State LOADinG Act (S7543B/A9430B, 2024, amended 2025), State Technology Law § 103-f. https://www.nysenate.gov/legislation/laws/STT/103-F*2
20. California AB 302 (2023), Gov. Code § 11546.45.5; California Department of Technology, *High-Risk Automated Decision Systems Report* (first report) and *High-Risk ADS Report for Program Year 2025* (June 2026); CalMatters commentary, 21 Jul 2026. https://www.cdt.ca.gov/wp-content/uploads/2026/03/High-Risk-Automated-Decision-Systems-ADS-Report-2025.pdf
21. Vermont Act 132 (2022), 3 V.S.A. § 3305. https://legislature.vermont.gov/statutes/section/03/056/03305
22. Connecticut Public Act 23-16 (SB 1103, 2023). https://prdext3.cga.ct.gov/2023/ACT/PA/PDF/2023PA-00016-R00SB-01103-PA.PDF
23. Texas HB 2060 (88th Leg., 2023), Gov. Code § 2054.623. https://www.legis.state.tx.us/tlodocs/88R/billtext/html/HB02060F.HTM
24. American Civil Liberties Union. *Community Control Over Police Surveillance* (26 jurisdictions). https://www.aclu.org/community-control-over-police-surveillance
25. Center for Democracy & Technology. *Regulating Public Sector AI: Emerging Trends in State Legislation*, 10 Jan 2025. https://cdt.org/insights/regulating-public-sector-ai-emerging-trends-in-state-legislation/