Trusted Data and Technology Institute
TECHNOLOGY WORTHY OF TRUST

RESEARCH PAPER 2026-10-02 · AI GOVERNANCE & PUBLIC ACCOUNTABILITY

Following the Money: Can Federal AI Disclosures Be Traced to Federal Contracts?

An independent audit linking the 2024–2025 Federal AI Use Case Inventories to the Federal Procurement Data System

Author
Joseph Atalla, PhD

Published
October 2026

ABSTRACT

Federal agencies publish an annual inventory of their artificial intelligence (AI) use cases and separately report every contract action to the Federal Procurement Data System (FPDS). Oversight of contractor-built AI depends on connecting the two. We test that connection directly. Using the full 2024 (2,133 use cases) and 2025 (3,611 use cases) consolidated inventories published by the Office of Management and Budget (OMB), we classify every contract identifier agencies disclosed against the Federal Acquisition Regulation's identifier standard, then verify a random sample of 60 identifiers against FPDS. Only 54% of contractor-built AI use cases listed a usable identifier in 2024; we estimate that about 28% can be traced to a specific contract action, and only about 7% of verified contracts describe AI in their public record. Among the government's 351 rights- or safety-impacting use cases, 15% carried a usable identifier. In 2025, OMB removed the identifier question entirely, just as new acquisition rules began placing AI-specific obligations into contracts. We propose a five-tier *Contract Traceability* measure and six low-cost fixes.

15%	≈28%	7%	0
of 351 rights- or safety-impacting AI use cases (2024) listed a usable contract identifier	of contractor-built AI use cases can be traced to a specific contract action (95% CI 21–35%)	of verified contracts mention AI, machine learning or algorithms in their FPDS description	contract identifiers requested in the 2025 inventory, which asks only for a vendor name

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 · California

Suggested citation. Atalla, J. (2026). *Following the Money: Can Federal AI Disclosures Be Traced to Federal Contracts?* TDTI Research Paper 2026-10-02. Trusted Data and Technology Institute.

Executive summary

The federal government now discloses thousands of AI use cases each year, and it records billions of dollars in contract actions. Yet the public cannot reliably tell which contracts pay for which AI systems. This matters most for the systems that can affect people's rights or safety, which are often built or run by contractors. Without a link, auditors cannot check whether contracts contain the testing, monitoring and documentation terms that federal AI policy now requires, and researchers cannot measure what the government actually spends on AI.

KEY FINDINGS

1 Half of contractor-built AI cannot be traced at the first step. Of 693 AI use cases agencies said were developed with contractor resources in 2024, only 377 (54.4%) listed an identifier that could be looked up.

2 The riskiest systems are the least traceable, and the gap is concentrated in two departments. Only 52 of 351 rights- or safety-impacting use cases (14.8%) carried a usable identifier. The Department of Justice (124 such use cases) did not answer the contracting question for any of them, and Veterans Affairs supplied identifiers for 16 of 145.

3 A listed identifier is not the same as a traceable contract. In a random sample of 60 identifiers, 44 (73.3%) resolved in FPDS; 13 of those pointed to a government-wide vehicle rather than a specific order. We estimate about 195 of the 693 use cases (28%) trace to a specific contract action.

4 Contracts rarely say "AI." Only 3 of 44 verified contract records (6.8%) mention AI, machine learning or algorithms. Keyword-based estimates of federal AI spending therefore measure a small, visible fraction of AI-related contracting.

5 2025 moved backward. OMB's 2025 reporting guidance dropped the identifier question and asks only for vendor names. Even those were missing for half of live high-impact use cases (108 of 217 named), because Veterans Affairs named no vendor for any of its 215 high-impact entries.

What should change. OMB should restore a mandatory, validated contract-identifier field for every contractor-built use case and make it non-optional for high-impact AI (R1–R2). Agencies should tag AI contract actions with the inventory use-case ID so FPDS and the inventory can be joined (R3). Oversight bodies should treat keyword-based spending figures as lower bounds (R5). State and local governments building their own AI inventories, including New York and California, can adopt the same linkage from the start (R6).

1 Why traceability matters now

Federal agencies have been required to inventory their AI use cases since Executive Order 13960 (2020) and the Advancing American AI Act, with OMB consolidating the results each year.[1,2] OMB Memorandum M-24-10 (March 2024) added minimum risk practices for "rights-impacting" and "safety-impacting" AI.[3] In April 2025, M-25-21 replaced that framework with a single "high-impact" category and required agencies to document minimum practices within 365 days or stop using the system.[4]

The companion memorandum, M-25-22, moved AI governance into the contract itself. For solicitations issued 180 days after April 3, 2025, contracts for AI must let the agency regularly monitor and evaluate performance and must detail the vendor's testing and validation procedures; for high-impact uses, agencies must tell vendors what documentation will be required.[5] These obligations live in contract files. The public record of AI use lives in the inventory. Neither memorandum requires a link between them.

That link is the subject of this paper. If a reader cannot get from an inventory entry to the contract that built or runs the system, then contract-level safeguards cannot be verified from public data, spending cannot be attributed to uses, and the relationship between what agencies *say* they use and what they *buy* remains unknown.

Research questions

RQ1. What share of contractor-built federal AI use cases disclose a contract identifier that can be resolved in the public procurement record? **RQ2.** Does traceability differ for rights-, safety- or high-impact AI? **RQ3.** When a contract can be found, does its public record identify it as AI? **RQ4.** How did the 2025 reporting changes affect traceability?

Prior work and contribution

Three strands of work come close. GAO found in December 2023 that only 5 of 20 agencies reporting AI use cases provided comprehensive inventory data, and made 35 recommendations to 19 agencies.[6] An independent 2026 audit of the 2025 inventory assessed whether high-impact determinations and evidence claims hold up, and profiled vendor concentration, but did not examine contracts or spending.[7] On the spending side, the Center for Strategic and International Studies (CSIS) identified about 2,255 AI awards worth roughly $4.1 billion (FY2019–FY2025) using keyword searches of award descriptions, which its authors describe as a lower bound; that study did not use the inventories.[8] To our knowledge, this paper is the first to join the two records, the first to verify inventory-disclosed contract identifiers against FPDS, and the first to measure how often AI contracts are described as AI. We searched the web, GAO, and arXiv for prior linkage studies in October 2026; a formal systematic review was not performed.

2 Data and methods

Source	Coverage	Records	Used for
OMB 2024 Federal AI Use Case Inventory (v2 consolidated file)[1]	41 agency submissions; snapshot of 23 Jan 2025	2,133	Contract identifiers, development method, impact category
OMB 2025 Federal Agency AI Use Case Inventory[9]	56 submissions; repository commit 06c7ebe (14 May 2026)	3,611	Vendor names, contracting usage, high-impact status
OMB 2024 and 2025 reporting guidance and data dictionaries[1,9]	Field definitions and conditional requirements	—	What agencies were asked to report
Federal Procurement Data System, public ATOM feed[10]	All contract actions reported by civilian agencies	60 queried	Resolving identifiers; award type, agency, vendor, description
FAR 4.1603, uniform PIID structure[11]	Regulatory standard	—	Validating identifier format

Population. RQ1–RQ3 use the 693 use cases in the 2024 inventory whose development method was "contracting resources" or "both contracting and in-house." The 2024 data dictionary required a Procurement Instrument Identifier (PIID) for these entries unless the use case was a listed commercial product. Impact status was taken from the inventory's own rights/safety field.

Identifier classification. We extracted every alphanumeric token of 8 or more characters containing at least four digits, removed hyphens and spaces, and classified each entry into one of four groups: a token matching the FAR uniform PIID pattern (13–17 characters: six-character agency code, two-digit fiscal year, instrument-type letter, and serial number); another alphanumeric identifier, which captures valid pre-2017 and agency-specific formats; numeric references only; or no identifier. "Usable" means one of the first two groups.

Verification sample. From 289 unique usable identifiers, we drew a simple random sample of 60 (seed 20261014) and queried each in the FPDS public feed on 14 October 2026. For each hit we recorded award type, contracting department, vendor, first signature date, description of requirement, and product and service code (PSC). Two derived measures follow: whether the record is a specific action (order, call, definitive contract, purchase order, other transaction) or only a contract vehicle (schedule, GWAC, IDIQ base, BPA), and whether its description mentions AI explicitly, mentions analytics or automation, or is generic.

Statistics. Proportions carry Wilson 95% confidence intervals. Group differences use Pearson's chi-square test or Fisher's exact test. Population estimates scale sample proportions to the 377 usable identifiers. All code, coded data and the sample are released with this paper (Section 6).

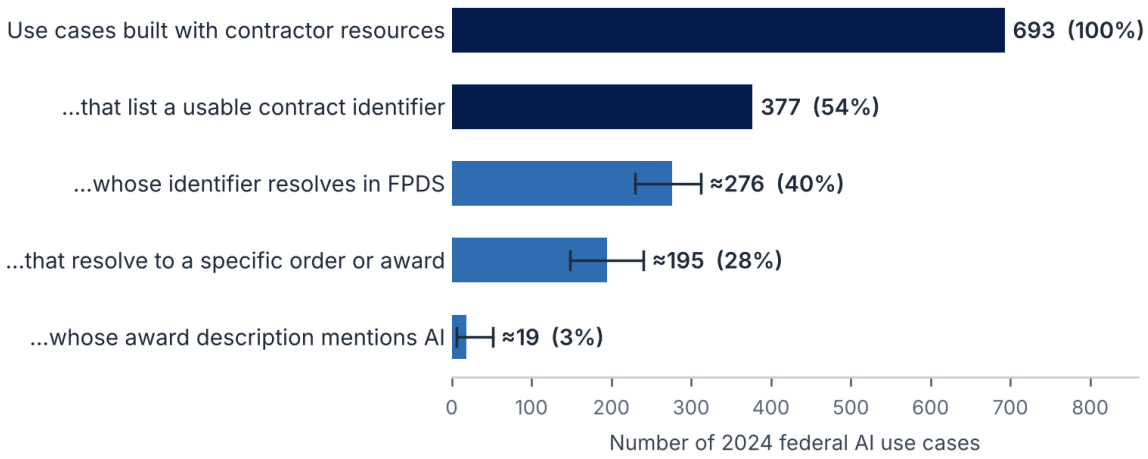
Contract Traceability Tiers (proposed). **T0** no identifier **T1** identifier does not resolve **T2** resolves to a vehicle only **T3** resolves to a specific action **T4** specific action whose public record identifies AI

3 Findings

3.1 Only about one in four contractor-built AI use cases reaches a specific contract

Of the 693 contractor-built use cases, 377 (54.4%) listed a usable identifier: 298 matched the uniform PIID format and 79 used another alphanumeric format. The rest were blank (225), contained non-identifier text such as "ITSS Contracts," "Not Reported" or a vendor name (76), or gave only numeric references (15). Verification then removed more. Identifiers in the uniform format resolved far more often (36 of 42, 85.7%) than other formats (8 of 18, 44.4%; Fisher's exact p = 0.003). Four sampled identifiers carried assistance-award codes (cooperative agreements and fixed-amount awards), which by design are not reported to FPDS.

Figure 1. From disclosure to contract: the 2024 traceability cascade



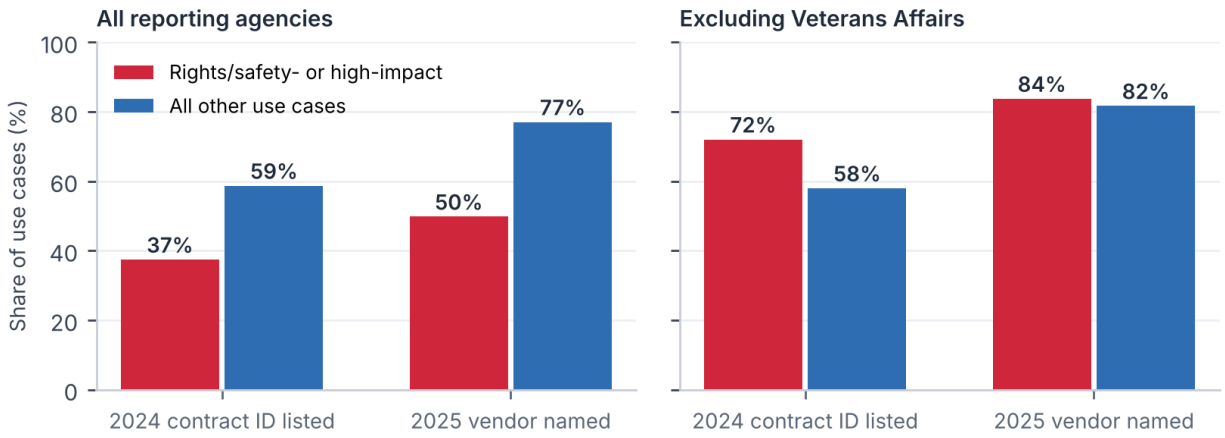
Source: OMB 2024 inventory (n = 693 contractor-built use cases); FPDS verification of a random sample of 60 identifiers. Navy bars are counts; blue bars are estimates scaled from the sample, with 95% Wilson intervals.

Scaling the sample to the 377 usable identifiers gives about 276 that resolve in FPDS, about 195 (28% of 693; 95% CI 21–35%) that resolve to a specific contract action, and about 19 (95% CI 6–52) whose public record identifies AI. In tier terms, roughly 46% of contractor-built AI sits at T0, 15% at T1, 12% at T2, 25% at T3 and 3% at T4.

3.2 Rights- and safety-impacting AI is the least traceable, for agency-specific reasons

Among contractor-built use cases, those flagged rights- or safety-impacting listed a usable identifier 37.4% of the time (52 of 139), compared with 58.7% for all others (325 of 554; $\chi^2 = 19.4$, $p < 0.001$). Across all 351 rights- or safety-impacting use cases, only 52 (14.8%) carried a usable identifier, because 184 (52.4%) did not state how the system was developed at all.

Figure 2. Higher-risk AI is less traceable, and one department drives the gap



Source: OMB 2024 and 2025 inventories. 2024: contractor-built use cases (n = 693; 578 excluding VA). 2025: deployed or pilot use cases that used vendors (n = 939; 810 excluding VA). "Higher risk" means rights/safety-impacting (2024) or high-impact (2025).

The gap is not a government-wide pattern. Excluding the Department of Veterans Affairs, which reported 145 of the 351 rights- or safety-impacting use cases, the relationship reverses: 72% of higher-risk contractor-built use cases listed an identifier, compared with 58% of others. A second department sits entirely outside the measure: the Department of Justice reported 124 rights- or safety-impacting use cases (35% of the government total) and left the development-method question blank for all 240 of its entries, so none can be assessed. Accountability for most of the government's higher-risk AI therefore depends on the reporting practices of two departments.

Department	Rights/safety-impacting use cases	Development method not stated	Usable contract identifier	Share traceable at step 1
Veterans Affairs	145	36	16	11%
Justice	124	124	0	0%
Homeland Security	34	8	15	44%
U.S. Agency for International Development	17	9	8	47%
All other agencies (12)	31	7	13	42%
Total	351	184	52	15%

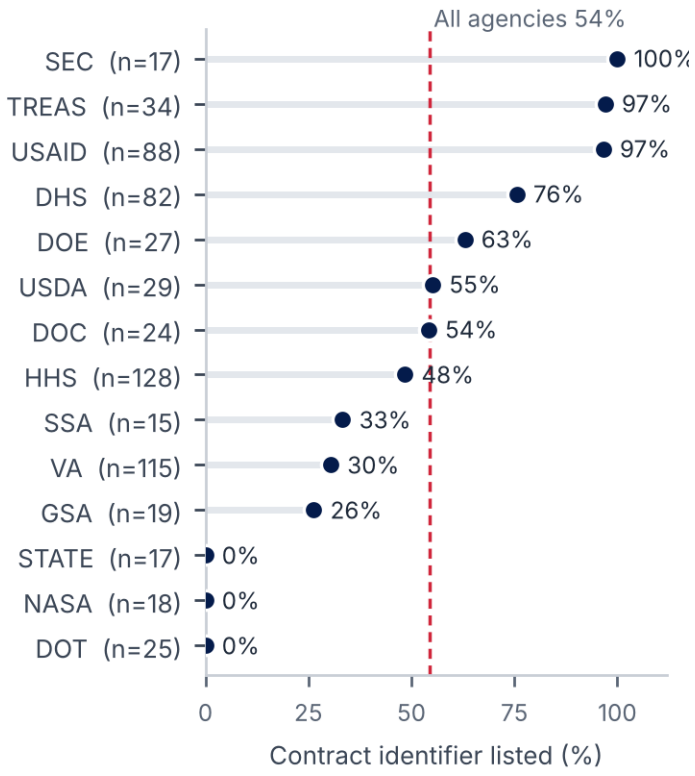
Table 1. Rights- and safety-impacting use cases in the 2024 inventory by department. Source: OMB 2024 inventory, TDTI analysis.

3.3 Agencies differ more than risk levels do

Among agencies with at least 15 contractor-built use cases, the share listing a usable identifier ranged from 0% (Transportation, NASA, State) to 97–100% (USAID, Treasury, SEC). The Department of Labor, Housing and Urban Development and others with fewer entries also reached 100%. This spread is far larger than the difference by risk category, which suggests that traceability is a matter of agency reporting practice, not of the systems themselves.

Because the 2024 field was free text with no format check, some agencies entered vehicle names ("SEAS IT," "ITSSC"), program names, or statements that the agency "does not use PIIDs." Each is an answer that a simple validation rule at submission would have caught.

Figure 3. Usable identifier rate by agency, 2024



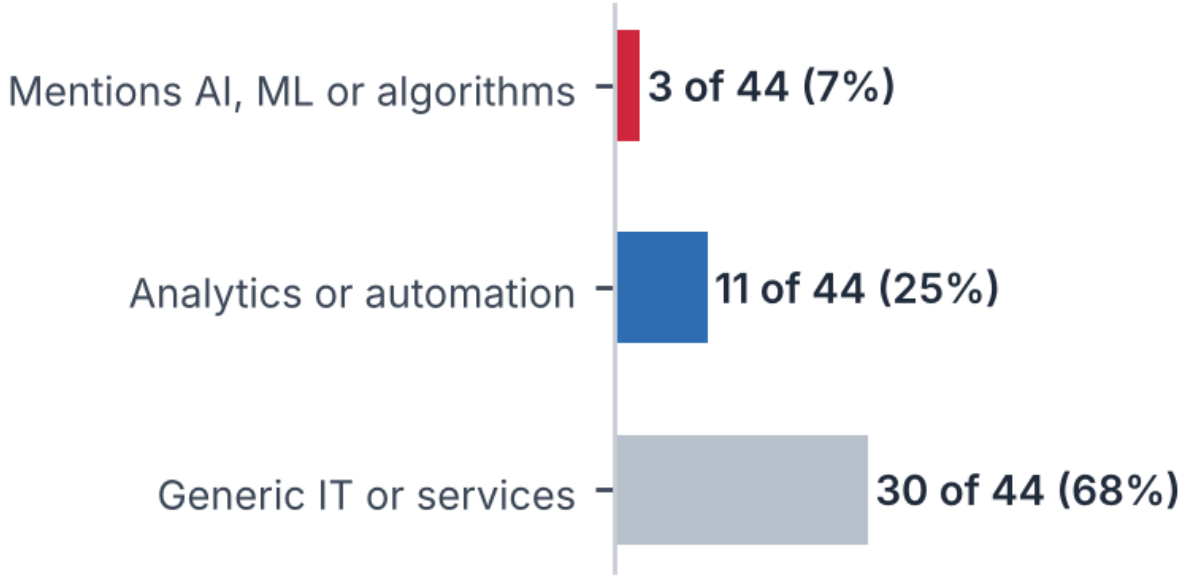
Source: OMB 2024 inventory; agencies with 15 or more contractor-built use cases. Dashed line: all agencies (54%).

3.4 Contracts rarely say "AI"

Of the 44 verified contract records, 30 (68%) used generic language such as "Federal Supply Schedule contract," "IT and telecom" or "cloud hosting services." Eleven (25%) referred to analytics, data science or automation. Only three (6.8%; 95% CI 2.3–18.2%) mentioned AI, machine learning or algorithms. These three were a Treasury purchase order for generative AI support, a VA order naming the ReflexAI training product, and a TSA contract for automatic target recognition algorithm development and evaluation.

This has a direct consequence for spending research. Any method that finds AI contracts by searching descriptions for AI terms would have found about 1 in 15 of the contracts that agencies themselves linked to AI.

Figure 4. How verified contracts describe themselves



Source: FPDS description of requirement for 44 resolved identifiers from the random sample.

The CSIS estimate of about $4.1 billion in AI obligations is explicitly a lower bound,[8] and our result suggests that lower bound may sit far below the true figure for civilian agencies. We do not estimate total AI spending; the sample is too small and many linked contracts are broad IT vehicles whose obligations cannot be allocated to one use case.

Case A: Electronic monitoring in immigration supervision

ICE listed "Biometric Check-in for ATD-ISAP (SmartLINK)" as rights- and safety-impacting in 2024 and gave identifier 70CDCR20D00000011. In FPDS this is an indefinite-delivery contract with B.I. Incorporated signed in March 2020, described as "Intensive Supervision Appearance Program (ISAP IV)," classified as PSC R408 and NAICS 561210 (facilities support services). Nothing in the record indicates biometrics or AI.

Tier T2. Sources: OMB 2024 inventory; FPDS.

Case B: Algorithm evaluation written into the contract

TSA listed "Automated Target Recognition (ATR) Developments for Standard Screening" as rights- and safety-impacting and gave identifier 70T04023C7573N001. FPDS shows a September 2023 definitive contract with Noblis for "automatic target recognition (ATR) algorithm engineering development, evaluation, and deployment." It is the model this paper argues for: the use case, the contract and the evaluation scope are publicly linked.

Tier T4. Sources: OMB 2024 inventory; FPDS.

3.5 The 2025 inventory removed the link

OMB's 2025 guidance shortened the reporting template. The contract-identifier question was dropped; agencies now report only whether a system was purchased, built under contract or built in-house, plus a free-text vendor name for pilot and deployed systems.[9] A vendor name cannot be joined to FPDS without knowing which of that vendor's contracts is involved; Microsoft alone was named in 103 of the 664 live entries that named any vendor.

Even vendor names were thin where they matter most. Of 939 live use cases that used vendors, 664 (70.7%) named one. For high-impact use cases the figure was 108 of 217 (49.8%), compared with 77.0% for others ($\chi^2 = 58.5$, $p < 0.001$). Again the gap reflects one department: Veterans Affairs named no vendor for any of its 129 vendor-involved live entries and reported 215 of the 445 high-impact use cases. Excluding VA, high-impact and other use cases named vendors at nearly identical rates (83.7% and 81.6%; $p = 0.66$). Across all 445 high-impact use cases, 113 (25.4%) name a vendor and none can be linked to a contract.

4 Discussion

A governance gap created by sequence. The 2024 inventory asked for the right data element but did not validate it. The 2025 inventory removed it at the same time M-25-22 began placing AI-specific testing, monitoring and documentation duties into new contracts. The result is that the safeguards most relevant to high-impact AI now sit in a record the public cannot connect to the inventory.

Traceability is cheap to fix. The FAR already standardizes the identifier.[11] Format validation would have rejected every non-identifier entry we found, and a lookup against FPDS, which we performed with public tools, would have flagged unresolvable identifiers before publication. Identifiers that followed the uniform format resolved 86% of the time.

Spending figures need a second source. If only about 7% of AI-linked contracts say "AI," keyword searches mostly find contracts whose AI content is already obvious. Linking from the inventory side catches the others: data-science task orders, cloud hosting, case-management upgrades and supervision programs that carry AI components without naming them.

Lessons for state and local inventories. New York publishes an AI systems inventory with agency, vendor and product fields, and California's 2025 report on high-risk automated decision systems lists eight systems across four agencies.[12,13] Neither, as published, links to a contract. Both can add a contract number field at little cost while their schemas are young.

5 Recommendations

R1 Restore the identifier. OMB should reinstate a contract-identifier field in the 2026 inventory for every use case built or operated with contractor resources, capturing both the order PIID and its parent vehicle.

R2 Make it mandatory and validated for high-impact AI. The field should be required for every high-impact use case, checked against the FAR 4.1603 format at submission, and resolved against FPDS before OMB publishes the consolidated file.

R3 Tag contracts with the use-case ID. Agencies should cite the inventory use-case identifier in the description of requirement for new AI contract actions, making the join possible from the procurement side as well.

R4 Publish a traceability score. OMB or agency inspectors general should report the share of contractor-built and high-impact use cases at each tier (T0–T4) as a standing measure of inventory quality.

R5 Treat keyword totals as floors. Oversight bodies and researchers should present keyword-based AI spending estimates as lower bounds and supplement them with inventory-linked contracts.

R6 Build the link into state and local inventories. States and localities designing AI inventories should include a contract or purchase-order number for every vendor-supplied system from the first release.

6 Limitations and reproducibility

Self-reported data. Inventories are agency self-reports. Department of Defense and intelligence community uses are excluded by policy,[4] and CSIS attributes 83% of identified AI obligations to Defense,[8] so these findings describe civilian AI only. **Sample size.** The verification sample is 60 identifiers; interval estimates are wide and are reported throughout. **FPDS access.** Records were retrieved from the FPDS public feed through an automated text-extraction tool, and fields were transcribed from its output; one description is a paraphrase. Identifiers that did not resolve may exist in records not shown in the public feed. **Classification rules.** Identifier parsing and description coding are rule-based and were applied by one analyst; no second-coder agreement was measured. **Scope.** We assess whether a link exists, not whether contracts contain required testing or monitoring terms; that is the natural next study once links exist.

Replication package. The scripts, the coded 2024 and 2025 extracts, the random sample, and every FPDS result accompany this paper and are available from TDTI. Source inventories are public on GitHub.[1,9]

References

1. Office of Management and Budget. *2024 Federal AI Use Case Inventory* (consolidated data, data dictionary and README). GitHub, snapshot of 23 Jan 2025. https://github.com/ombegov/2024-Federal-AI-Use-Case-Inventory
2. 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.
3. OMB. *M-24-10, Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence*. March 2024. 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
4. 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
5. OMB. *M-25-22, Driving Efficient Acquisition of Artificial Intelligence in Government*. 3 Apr 2025. Copy at https://digitalgovernmenthub.org/wp-content/uploads/2025/04/M-25-22-Driving-Efficient-Acquisition-of-Artificial-Intelligence-in-Government.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: An audit of the 2025 Federal AI Use Case Inventory*. GitHub, 2026 (not peer reviewed). https://github.com/ishaanshah101/fed-ai-inventory-audit
8. Atalan, Y., Tiersten-Nyman, E., and Jensen, B. *Understanding the U.S. Federal Government's AI Spending*. CSIS, 22 Sep 2026. https://www.csis.org/analysis/understanding-us-federal-governments-ai-spending
9. OMB. *2025 Federal Agency AI Use Case Inventory*, including *Guidance on 2025 Agency Artificial Intelligence Reporting* and data dictionary. GitHub, commit 06c7ebe. https://github.com/ombegov/2025-Federal-Agency-AI-Use-Case-Inventory
10. General Services Administration. Federal Procurement Data System, public ATOM feed, queried 14 Oct 2026. https://www.fpds.gov/ezsearch/FEEDS/ATOM?FEEDNAME=PUBLIC
11. Federal Acquisition Regulation 4.1603, *Procedures* (uniform PIID). https://www.acquisition.gov/far/4.1603
12. New York State. *NYS AI Systems Inventory: Beginning September 2025*. data.ny.gov, listed at https://catalog.data.gov/dataset/nys-ai-systems-inventory-beginning-september-2025
13. California Department of Technology. *High-Risk ADS Report for Program Year 2025*. June 2026. https://www.cdt.ca.gov/wp-content/uploads/2026/06/High-Risk-ADS-Report-for-Program-Year-2025.pdf