

2024 Federal AI Use Case Inventory

A Critical Analysis for Social Impact Leaders

A Note on Scope and Selection

This analysis draws from the [2024 Federal Agency AI Use Case Inventory](#), published by the Office of Management and Budget (OMB) via GitHub and individual agency websites as of December 16, 2024.

Three agencies are examined:

- the Department of Health and Human Services (HHS),
- the Department of Homeland Security (DHS),
- the Department of Justice (DOJ).

From roughly 400 combined use cases across these three agencies, ten were selected based on four criteria applied in priority order:

1. relevance to social impact organizations,
2. controversy or civil liberties risk,
3. technical distinctiveness,
4. scale of population affected.

The selection deliberately weights enforcement-facing systems because those systems disproportionately affect the populations that social impact organizations serve.

Conclusions about the ten selected cases should not be generalized to the full 1,700-case federal inventory without qualification.

[INCOMPLETE ENTRY] flags are used throughout where the inventory entry lacks implementation detail, validation data, or demographic performance analysis that would be necessary to assess the system's actual risk profile.

Critical Analysis of 2024 Federal AI Use Case Inventory

The Ten Use Cases

Agency	Use Case	Component	Impact	Stage
DHS/ICE	Hurricane Score (DHS-2408)	ICE/ERO	Rights-impacting	Deployed
DHS/ICE	Biometric Check-in for ATD-ISAP (SmartLINK)	ICE/ERO	Rights-impacting	Deployed
DHS/CBP	AI for Autonomous Situational Awareness	CBP	Not assessed	Development
DHS/CBP	Automated Item of Interest Detection (ICAD)	CBP	Not assessed	Operation/Maintenance
DHS/HSI	Facial Recognition for Investigations of CSEA	ICE/HSI	Rights-impacting	Deployed
DHS/USCIS	ARGOS E-Verify Fraud Risk Scoring	USCIS	Rights-impacting	Deployed
HHS/CMS	Fraud Prevention System Models (CPI)	CMS	Not designated	Operation/Maintenance
HHS/FDA	Emerging Chemical Hazard Intelligence Platform (ECHIP/WILEE)	FDA	Not designated	Development
HHS/OIG	Grants Analytics Portal (GAP)	OIG	Not designated	Deployed
DOJ	eLitigation AI Tools (DOJ-wide cluster)	Multiple components	Rights-impacting	Deployed

Sources: [DHS Simplified AI Use Case Inventory](#); [HHS AI Use Cases FY2022 PDF](#) and [2024 public inventory](#); [DOJ AI Inventory](#); [consolidated OMB GitHub dataset](#).

What the Government Is Actually Doing with AI

The dominant pattern across these ten cases is not a story about AI expanding government capacity into new domains. It is a story about AI being used to accelerate decisions that were already being made, often about the same populations, using data that carries existing institutional biases.

Research on algorithmic systems in public administration consistently shows that models trained on historical enforcement data replicate the patterns in that data, including demographic disparities, into automated outputs ([Barocas & Selbst, 2016](#); [Eubanks, 2018](#)). The inventory does not disclose whether the agencies examined here have tested their systems against that known risk. In most cases, it does not say.

[HHS, DHS, and VA together account for roughly 50 percent of all publicly reported federal AI use cases in the 2024 inventory.](#) Within that concentration, the three agencies examined here have sorted their AI investments into two functionally distinct clusters: enforcement infrastructure and administrative efficiency.

DHS sits almost entirely in the first. HHS sits mostly in the second, with its most consequential cases at CMS and OIG clustering around fraud detection. DOJ spans both, with a 2024 inventory that [grew 1,507 percent from 2023, driven by a rigorous department-wide data call that pulled in components that had previously not reported at all.](#)

That growth figure is striking but requires one qualification: the inventory represents DOJ's own self-reporting, and the 1,507 percent increase likely reflects improved compliance with disclosure requirements more than a corresponding increase in actual AI deployment. The systems were already running.

What the inventory choices reveal about each agency's theory of AI is consistent across all three. AI is understood primarily as a tool for managing volume: more cases processed faster, more flags generated per analyst hour, more claims reviewed per investigator.

The [top three categories across the entire federal inventory are mission-enabling \(internal support\), health and medical, and government services including benefits delivery.](#) The enforcement weight in the DHS and DOJ cases examined here is not a distortion of that overall pattern. It reflects who has the most volume to manage and the most institutional pressure to process it faster.

The compliance data is the most important systemic finding in the inventory and it appears mid-document in most summaries. It belongs at the top. [Of the more than 1,700 AI use cases reported as of December 16, 2024, 227 were identified as rights-impacting and/or](#)

Critical Analysis of 2024 Federal AI Use Case Inventory

[safety-impacting. Of those, 206 received compliance extensions of up to one year to meet OMB's minimum risk management practices.](#)

That is a 91 percent non-compliance rate among rights-impacting systems as of the inventory deadline. The extensions are not waivers, but they document something important: federal agencies have been running AI systems that affect people's rights and cannot yet certify that the required safeguards are in place, across nearly every system in this category.

The DOJ and DHS governance failures analyzed below are instances of a pattern that runs across the federal government, not isolated agency failures.

Where This Intersects with Social Impact Work

The **CMS Fraud Prevention System Models (Use Case 7)** are the most directly instructive case for any organization that processes public health dollars. [CMS uses tree-based models and deep learning approaches applied to Medicare administrative and claims data to detect, prevent, and prioritize potential cases of fraud, waste, and abuse. The system assigns weights to conventional fraud identification models and identifies underperforming models for deactivation.](#)

The practical implication for a nonprofit FQHC, a behavioral health provider, or a community health center is that the same claims data used to reimburse them is also training data for a system designed to flag billing anomalies. The inventory entry does not disclose false positive rates, demographic performance, or how unusual but legitimate billing profiles are distinguished from fraud.

This is not a hypothetical concern. Research on algorithmic systems in public benefit administration documents that unusual billing and service patterns, generated by organizations serving high-need, high-complexity populations, register as risk signals in models trained on population averages ([Eubanks, 2018](#)).

Organizations serving patients with multiple chronic conditions, unhoused individuals, or populations requiring intensive care coordination generate claims patterns that deviate from means. Whether CMS's fraud detection system has been tested for differential false-positive rates across provider types is not disclosed in the inventory.

[INCOMPLETE ENTRY: No performance metrics, false positive rates, or demographic disparity analysis are publicly disclosed for the CMS fraud detection system.]

Critical Analysis of 2024 Federal AI Use Case Inventory

The **HHS OIG Grants Analytics Portal (Use Case 9)** operates by the same logic at the grantee level. [A deep neural network model extracts findings from single audit PDFs submitted by grant recipients, fuses those results with other datasets, and produces statistical risk assessments across HHS grant programs, enabling staff to evaluate programs with the most risk.](#)

For any organization receiving HHS funding, the immediate operational question is: what data feeds the risk score that determines which audits get opened?

Smaller organizations, rural providers, and immigrant-serving agencies frequently have structural audit findings that reflect legitimate programmatic complexity rather than mismanagement — including findings from predecessor organizations, findings under appeal, or findings related to inadequate financial management capacity rather than improper use of funds.

The OIG GAP shapes where oversight resources flow. Organizations that do not understand its logic cannot manage their exposure to it.

The **FDA Emerging Chemical Hazard Intelligence Platform (ECHIP/WILEE, Use Case 8)** belongs in a separate category of relevance. [ECHIP/WILEE uses AI to develop a horizon-scanning platform that identifies, aggregates, maps, and links information from internal and external data sources related to consumed foods, ingredients, and food chemicals. A model is trained to identify patterns of events that anticipate chemical hazard signal detection across time and multiple data sources.](#)

For public health organizations working in food security, environmental justice, or occupational health, this represents federal AI investment in something the social sector cannot replicate at scale: integrating global supply chain data, adverse event reporting, and toxicological literature into a single early-warning architecture.

The transferable lesson is the methodology, not the technology. Multi-source signal aggregation for early hazard detection has direct analogs in population health surveillance, housing instability prediction, and early childhood development programs.

The **DOJ eLitigation AI cluster (Use Case 10)** is the most underexamined item in this analysis for legal services organizations. [DOJ has consolidated reporting of eLitigation tools used across the department, with varying AI features, purposes, and contexts. The tools are used in rights-impacting settings, and DOJ has begun developing a department-wide governance policy.](#)

Legal aid organizations, public defenders, and immigrant legal services providers are now litigating against an opponent that has deployed AI to synthesize evidence, identify patterns in large document sets, and prepare cases faster.

Critical Analysis of 2024 Federal AI Use Case Inventory

[The DOJ Executive Office for United States Attorneys has deployed Palantir for "integration and analysis of case information" and to reduce time needed to update and maintain case management systems.](#)

The adversarial implications of this gap are real, but the picture is more complex than a simple asymmetry. Research published in 2025 on public defenders' actual AI adoption found that [AI adoption in public defense is constrained by costs, restrictive office norms, confidentiality risks, and unsatisfactory tool quality, with public defenders viewing AI as most useful for analyzing overwhelming volumes of digital evidence but facing systemic barriers to adoption.](#)

On the legal aid side, [a 2025 survey found that 74 percent of legal aid organizations are already using AI in their work, roughly double the adoption rate reported across the wider legal profession, with 90 percent of respondents saying that using AI to its full potential would enable them to serve more clients.](#)

The gap is not that one side has AI and the other does not. The gap is institutional: federal prosecutors operate within a funded, centrally managed AI governance structure while public defenders and legal aid attorneys adopt tools individually, without institutional support, and with genuine confidentiality constraints that their counterparts in federal prosecution do not face.

That structural difference is the documented finding. The outcome consequences remain, as of this writing, unquantified.

The **ICE ARGOS system at USCIS (Use Case 6)** carries the most direct implications for workforce development and economic inclusion organizations. [ARGOS is a USCIS system that analyzes public datasets to assess fraud risk by companies using E-Verify. The inventory documents known failure modes including data collection bias from search engine prioritization, lack of domain-specific accuracy across different industries, limited generalization to unseen data, and misinterpretation of sentiment.](#)

Any workforce organization that connects employers with immigrant workers should understand that their employer partners may be flagged by a risk model whose own inventory entry acknowledges multiple known failure modes. The documentation of those failure modes in the public inventory is a notable act of disclosure.

What the entry does not document is whether any mitigation actions have been taken, or what happens to an employer incorrectly flagged.

What Is Being Risked and Who Bears It

The **Hurricane Score (Use Case 1, DHS-2408)** is the single most consequential AI system in this analysis for organizations working with immigrant communities. [The Hurricane Score models the probability that a noncitizen released from ICE detention on Alternatives to Detention will fail to comply with the program. The model considers factors including violation history, length of time in the program, and whether the person has a travel document. ICE officers are directed to consider the score as one of many inputs in individualized decisions about a noncitizen's case.](#)

The "one of many inputs" framing is the key contested claim, and the inventory cannot resolve it.

Research on how human decision-makers interact with algorithmic risk scores in high-stakes government contexts consistently documents what [Alon-Barkat and Busuioc \(2023\)](#) identified as "selective adherence": decision-makers do not uniformly defer to or dismiss algorithmic advice, but they show stronger adherence when recommendations align with pre-existing group stereotypes, particularly when a score predicts high risk for members of already-stigmatized groups.

[Green and Chen \(2019, 2021\)](#) documented the same pattern in criminal risk assessment contexts, finding that participants adhered more strongly to algorithmic predictions of high risk for Black defendants than for comparable white defendants. The inventory provides no data on Hurricane Score's demographic performance across national origin, English proficiency, or family status, no description of officer deviation rates from the score's implied recommendation, and no disclosed validation study.

[INCOMPLETE ENTRY: No validation study, false positive rate, or demographic disparity analysis is disclosed in the public inventory. The analysis above searched for Hurricane Score documentation outside the inventory, including DHS Science and Technology Directorate publications and Congressional testimony, and found no publicly available validation study as of March 2026.]

The **SmartLINK biometric check-in system (Use Case 2)** presents a distinct risk architecture. [The Intensive Supervision Appearance Program Monitoring App uses facial verification technology that captures a participant's photo during enrollment and uses one-to-one matching during subsequent check-ins. If remote check-in fails, an officer manually reviews the check-in photo. The AI's output is described as not the primary basis for decisions affecting individual rights.](#)

Critical Analysis of 2024 Federal AI Use Case Inventory

The framing is technically careful but operationally incomplete. For someone on Alternatives to Detention, any compliance failure — including a false negative from a facial recognition system — initiates an escalation process that may result in detention.

The inventory does not disclose error rates under different environmental conditions, demographic performance variation, or the rate at which technical failures are attributed to participant non-compliance rather than system error. [INCOMPLETE ENTRY]

The **ICAD Automated Item of Interest Detection at CBP (Use Case 4)** presents surveillance risk that the inventory undersells through vague language. [The Matroid software processes and annotates images from field imaging equipment to determine whether they contain human subjects, using trained computer vision models to recognize objects, people, and events. The system is intended to expand to vehicles and subjects with long-arm rifles, while excluding items of little interest such as animals.](#)

The system is described as in operation and maintenance, meaning it is deployed. The inventory does not disclose accuracy rates for human detection under varying lighting or environmental conditions, nor the false positive rate for human identification or the rate of officer action following system classification.

For populations living and working in border communities, including US citizens, documented immigrants, and people seeking asylum, this is a deployed surveillance system with no disclosed performance floor. [INCOMPLETE ENTRY]

Use Case 5, the ICE Facial Recognition for Child Sexual Exploitation Investigations, requires specific analytical care because the stated purpose is widely supported and the risk is architectural rather than intentional. [Homeland Security Investigations' Child Exploitation Investigations Unit submits unidentified CSAM images to an AI-enabled facial recognition service that compares photos to publicly available online images to generate investigative leads.](#)

The stated function is victim identification in child abuse investigations, and nothing in the inventory suggests otherwise. The risk is structural: the same facial recognition capability used to compare abuse images against public sources can also be used, in a different query direction, to identify individuals of investigative interest from public social media imagery without individualized suspicion.

The inventory draws no operational boundary between these functions, and does not disclose what constraints, if any, exist on the use of the underlying technology outside the CSAM context. This is a plausible risk requiring disclosure, not a documented finding about actual expanded use.

What makes it warrant naming is the transparency failure that surrounds it: [the DHS Office of Inspector General found that 66 ICE use cases, including facial recognition systems, were](#)

Critical Analysis of 2024 Federal AI Use Case Inventory

[initially omitted from mandatory reporting requirements](#). An agency that omits its facial recognition systems from mandatory disclosure and then documents those systems incompletely when forced to report them has not established the basis for the public trust its inventory claims to support.

The **DOJ eLitigation cluster (Use Case 10)** closes the structural argument. [DOJ acknowledged that the nature and details of AI use in eLitigation tools vary, which may affect whether particular uses are rights-impacting, and that a department-wide governance policy is in development](#).

The governance policy was not in place when the tools were deployed. This pattern — deploying rights-impacting AI systems before completing governance frameworks — is not a DOJ anomaly. It is the documented federal norm: the 206/227 compliance extension figure establishes that nearly every rights-impacting system in the federal inventory was running before the required safeguards were certified.

The DOJ case is notable because the rights-impacting context is explicitly adversarial, involving individuals charged with crimes or facing civil liability against a department deploying AI to build cases against them, and because the governance process remains incomplete.

The **CMS fraud detection system (Use Case 7)** closes the loop on what is most at risk across all three agencies: the structural vulnerability of organizations and individuals who cannot see the models evaluating them.

A Medicaid provider receives a Targeted Probe and Educate letter. A grantee gets an audit notice. A noncitizen receives a detention recommendation informed by a risk score. A legal services attorney faces a prosecutor with AI-assisted case synthesis.

In none of these cases does the affected party have meaningful access to the model's logic, error rate, or training data. The inventory acknowledges these systems exist and, in some cases, names the risks. It does not give the people affected by those systems any mechanism to contest the inferences those systems generate.

Limitations

This analysis is bounded by what the inventory discloses. The 2024 inventory is a self-reported instrument, prepared by agencies with institutional interests in managing both disclosure and liability exposure. Use cases that are classified, operationally sensitive, or simply unreported are absent.

Critical Analysis of 2024 Federal AI Use Case Inventory

The [INCOMPLETE ENTRY] flags in this analysis reflect the state of the public inventory as of December 16, 2024; entries may have been updated since. The enforcement-weighted case selection reflects what is most relevant to social impact audiences but does not represent the full distribution of federal AI use.

Causal claims about how these systems affect outcomes require operational data the inventory does not contain, and in most cases have not been independently studied.

The question for social impact organizations is not whether to engage with this landscape. These systems are already running, and many were running before their governance frameworks were written.

The question is whether the organizations most affected by these systems — health providers, grantees, legal aid clients, and immigrants in detention and removal proceedings — have anyone at the table who can read these entries closely enough to know what is missing from them.

References

- Alon-Barkat, S., & Busuioc, M. (2023). Human–AI interactions in public sector decision making: "Automation bias" and "selective adherence" to algorithmic advice. *Journal of Public Administration Research and Theory*, 33(1), 153–169. <https://doi.org/10.1093/jopart/muac007>
- Barocas, S., & Selbst, A. D. (2016). Big data's disparate impact. *California Law Review*, 104, 671–732. <https://doi.org/10.15779/Z38BG31>
- Brennan Center for Justice. (2025). A start for AI transparency at DHS with room to grow. <https://www.brennancenter.org/our-work/analysis-opinion/start-ai-transparency-dhs-room-grow>
- Chien, C. V., & Kim, M. (2024). Generative AI and legal aid: Results from a field study and 100 use cases to bridge the access to justice gap. *Loyola of Los Angeles Law Review* (Forthcoming). <https://ssrn.com/abstract=4733061>
- Council on Criminal Justice. (2026, January 6). DOJ report on AI in criminal justice: Key takeaways. <https://counciloncj.org/doj-report-on-ai-in-criminal-justice-key-takeaways/>

Critical Analysis of 2024 Federal AI Use Case Inventory

- Department of Health and Human Services. (2022). Artificial intelligence use cases — FY2022. <https://www.hhs.gov/sites/default/files/hhs-artificial-intelligence-select-use-cases.pdf>
- Department of Homeland Security. (2024). AI at DHS: A deep dive into our use case inventory [Archived]. <https://www.dhs.gov/archive/news/2024/12/16/ai-dhs-deep-dive-our-use-case-inventory>
- Department of Homeland Security. (2024). United States Immigration and Customs Enforcement: AI use cases. <https://www.dhs.gov/ai/use-case-inventory/ice>
- Department of Homeland Security. (2024). United States Citizenship and Immigration Services: AI use cases. <https://www.dhs.gov/ai/use-case-inventory/uscis>
- Department of Justice. (2025, January 21). AI inventory. <https://www.justice.gov/ai/ai-inventory>
- Department of Justice. (2024, December 3). Artificial intelligence and criminal justice: Final report. <https://www.justice.gov/olp/media/1381796/dl>
- Eubanks, V. (2018). *Automating inequality: How high-tech tools profile, police, and punish the poor*. St. Martin's Press. <https://us.macmillan.com/books/9781250074317/automatinginequality>
- Everlaw, NLADA, & Paladin. (2025). The AI advantage: How technology can help bridge the justice gap. <https://www.everlaw.com/blog/everlaw-for-good/88-of-legal-aid-professionals-see-ai-as-key-for-access-to-justice/>
- FedScoop. (2026, February 12). DOJ ramps up AI for legal work, crime predictions, surveillance, inventory shows. <https://fedscoop.com/justice-department-artificial-intelligence-ai-surveillance-inventory-predictive-technology-algorithm-bias/>
- Gao, X., et al. (2025). How can AI augment access to justice? Public defenders' perspectives on AI adoption [Preprint]. arXiv:2510.22933. <https://arxiv.org/abs/2510.22933>
- Green, B., & Chen, Y. (2019). Disparate interactions: An algorithm-in-the-loop analysis of fairness in risk assessments. *Proceedings of the ACM Conference on Fairness, Accountability, and Transparency*, 90–99. <https://doi.org/10.1145/3287560.3287598>

Critical Analysis of 2024 Federal AI Use Case Inventory

- Green, B., & Chen, Y. (2021). Algorithmic risk assessments can alter human decision-making processes in high-stakes government contexts. *Proceedings of the ACM on Human-Computer Interaction*, 5, 1–33. <https://doi.org/10.1145/3479562>
- American Immigration Council. (2025, May). Invisible gatekeepers: DHS's growing use of AI in immigration decisions. <https://www.americanimmigrationcouncil.org/blog/invisible-gatekeepers-dhs-growing-use-of-ai-in-immigration-decisions/>
- Nextgov/FCW. (2026, January 29). Law enforcement is the leading DHS use case for AI. <https://www.nextgov.com/artificial-intelligence/2026/01/law-enforcement-leading-dhs-use-case-ai/411063/>
- Office of the Federal Chief Information Officer. (2025). Executive Order 13960 AI use case inventories reference. <https://www.cio.gov/policies-and-priorities/Executive-Order-13960-AI-Use-Case-Inventories-Reference>
- Office of Management and Budget. (2025). 2024 Federal AI Use Case Inventory [Data repository, December 16, 2024 snapshot]. GitHub. <https://github.com/ombegov/2024-Federal-AI-Use-Case-Inventory>