

Battlefield Validation Is Not Civil Authorization

Governance risks when AI validated in military, intelligence, and security settings migrates into civil administration.

By Alec F., Independent Researcher

Original June 2026 · Revised October 7, 2026 (version 1.2) · Research current to October 7, 2026

AI use. During the preparation of this work, the author used AI tools (Grok, Claude, DeepSeek, Gemini, and ChatGPT) to gather and check sources, draft and revise sections, and review the model legislation. The author directed the research, reviewed, edited, and verified the output, and takes full responsibility for the content.

How to read this edition. Numbered superscripts open the source. Key empirical claims carry evidentiary tags: [CONFIRMED] (publicly documented), [REPORTED] (an on-record claim by an official, agency, or vendor, or a single press account, not independently verified), [INFERRED] (a reasoned conclusion from documented facts), [OPAQUE] (material facts not located in public sources). Appendices E and F set out model legislative language. Notes before Section 01 list what changed in versions 1.2 and 1.1; a changelog at the end lists earlier revisions.

Version 1.2

What Changed in Version 1.2 (October 7, 2026)

Version 1.2 was prepared later on October 7, 2026. It changes only the state bill (Appendix F, F.1). No factual finding changed.

- **A narrow, certified continuation (§5(f)(5)–(7)).** When a system moves into Tier 2 only because of its security lineage, the agency may keep using it for non-adverse work while it prepares the full dossier, but only if the agency head certifies within 5 business days, backed by the vendor’s written attestation, that every capability that could feed an adverse decision is switched off at the license or configuration level, that any attempt to switch it on is logged for the State Auditor, and that switching it on needs a new dossier.
- **When continuation ends.** It is lost if the system flags or queues people for investigation, audit, denial, or extra review (a human signing off at the end does not change that), slows others down, attaches conditions to approvals, or if the certification proves false. The State Auditor can revoke it at any time. Without a timely certification, the system stops.
- **Protection for people.** If a system is stopped, benefits and approvals already granted stand, cases go back to ordinary processing, and the stop cannot be used as a reason to deny or delay anyone.

- **Conditional approvals count as adverse (§3(l)(7)).** Granting something subject to extra conditions, monitoring, reporting, or restrictions is now an adverse decision, and the favorable-only test in §3(m) now excludes systems that add conditions, cause missed deadlines or added scrutiny for others, or queue people for extra review.
- **Vendor accountability (§4(b), §10(b)).** New contracts must let the agency terminate, recover fees, and seek suspension or debarment if the vendor knowingly supplies false or materially incomplete dossier information, lineage disclosures, or continuation attestations. The optional vendor penalty now covers those attestations too.
- **New questions for counsel** on whether self-certification with Auditor revocation is enough and how “disabled at the license or configuration level” should be verified.

Version 1.1

What Changed in Version 1.1 (October 7, 2026)

Version 1.1 was prepared on October 7, 2026, after further review of the model bills and the brief’s evidence labels. No factual finding changed. The main aim was to make the safeguards the bills describe match the safeguards they actually require, and to make every summary match the bill text. In plain terms:

State bill (Appendix F, F.1)

- **Testing now has consequences.** Before testing, an independent evaluator sets accuracy and fairness criteria for the specific use. A system that fails them cannot be deployed until it is fixed and retested, and the agency cannot deploy over a failed evaluation. A later finding of material error or disparate impact requires suspension or a published remediation plan (§6(a)–(b), (f)).
- **Testing and the human-decision rule cover every adverse-decision system,** not only security-derived ones. Security-derived systems keep extra requirements: they are always Tier 2, their lineage must be disclosed, and they need a published comparison of the security and civil settings, legislative notice, and a shorter sunset (§6(e), (g)).
- **State Auditor review is mandatory.** The Auditor “shall review” (it previously “may review”) the dossier and evaluation of every Tier 2 system. Deployment waits for the later of the 60-day notice period or the Auditor’s findings, which are due within 90 days. If the Auditor misses that deadline, the agency may proceed only after saying so publicly, and the Auditor can still act afterward. The Auditor can bar or suspend a system that fails its criteria or lacks a complete dossier; whether a state’s auditor can hold that power is flagged for counsel (§6(c)–(d)).
- **One timing rule for moving into Tier 2.** Whether a system escalates through a change in use, is reclassified by the Auditor, or is found to have security lineage, its use in adverse decisions stops until the Tier 2 requirements are met, and a system already in use must file a full dossier and auditor submission within 60 days or come out of service. The earliest trigger governs (§5(f)).

- **Rules engines are covered.** Deterministic, rules-based systems used to make or materially support adverse decisions, like Michigan's MiDAS, now clearly count (§3(a)).
- **A test for "favorable-only" systems.** A system that only speeds approvals stays in Tier 1 only if people it does not select get unchanged processing (no added delay or staffing shift) and non-selection is never used as a flag (§3(m)).
- **Established versus presumed lineage.** The dossier must say whether security lineage is established or only presumed from the vendor's product line (§3(g)).
- **Faster appeals where harm is urgent.** Expedited appeals, with benefits continued and collection paused pending the decision (§7(c)–(d)).
- **Transitional inventory entries** for systems already in use, with a deadline for their dossiers (§9(a), §12(a)).

Both bills

- **Complaints need a real answer.** Within 90 days the agency must issue a written decision with findings, corrective action or reasons for denial, and the record relied on. An acknowledgment does not count; silence or a non-answer is a deemed denial that is final agency action (E.1 §6(c), F.1 §10(c)). E.1 complaints now also cover defective dossiers, firewall violations, and expired authorizations.
- **Federal waiver narrowed (E.1 §6(d)).** It reaches only the listed sections; reporting, Inspector General notice, complaints, firewalls, and the savings clause stay in force; one 180-day renewal, then no further waiver for the same system and purpose. The state waiver now preserves the same protections (F.1 §11(d)).
- **Federal renewal (E.1 §4).** Reauthorization requires an updated dossier and an independent review, as in E.2.
- **E.2 stands on its own** with its own definitions and dossier contents.
- **E.1's scope is stated plainly.** It is a lineage-plus-reliance gate, not a universal ban. Systems outside it remain covered by general rules for rights-impacting AI, by E.2, and, for state agencies, by F.1.

The brief

- **Evidence tags.** Section 09, the cover, and the closing note now use the same definitions. They distinguish a documented official act from a performance claim, and "not located in our search" from "not publicly available."
- **MiDAS.** Its notice and appeal procedures are described as defective in practice, not absent.
- **Why security lineage matters.** Sections 01 and 03 now give concrete reasons for scrutiny beyond the adverse-decision safeguards; shared vendors alone are not the argument.
- **Authorization language.** "Separate democratic authorization" now describes the bills' actual mechanism: a statute, regulation, or published agency authorization, with a dossier and independent review.

- **Error equation (Section 10)** is presented as an illustration of assumptions, not evidence.
- **A worked example** walks a hypothetical benefits fraud-scoring tool through the state bill (Appendix F).
- **New questions for counsel** on the reliance gate, complaint reviewability, the Auditor's power to block deployment, and rules-engine coverage (Appendices E and F).

Section 01: Executive Summary

Bottom line

Five findings

- Military AI command infrastructure is becoming permanent. NATO's Maven Smart System reached full operational capability and classified-network accreditation in June 2026,⁴ [CONFIRMED] and a March 9, 2026 Pentagon memo directed that Maven become a program of record by September 30, 2026.⁵ [CONFIRMED] Whether it did is unverified. [OPAQUE]
- The same vendor's platforms now span U.S. targeting, NATO command, Ukraine's wartime data environment, and DHS-wide enforcement software, including a \$1 billion DHS purchase agreement.^{3,5,7,37} [CONFIRMED]
- Courts have found or restrained unlawful flows of tax and Medicaid data into immigration enforcement.^{8,9} [CONFIRMED]
- The main U.S. federal AI governance memoranda and the EU AI Act carve out national-security systems,^{10,11,12} [CONFIRMED] so the hand-off from security to civil use is where scrutiny is weakest. [INFERRED]
- No documented instance of a battlefield-validated model entering a civil system was located, and the contract terms and lineage that would show one could not be found in public sources. [OPAQUE] Automated-administration failures such as Robodebt, SyRI, and MiDAS show what happens when systems are deployed before legal authority, transparency, and adequate notice and appeal procedures are secured.^{51,52,54} [CONFIRMED]

Five actions

- Require a civil authorization dossier before deployment (OMB to require; agencies to produce; model text in Appendix E, E.2, and Appendix F, F.1).
- Require separate legal authorization before a security-derived system is used in civil administration on the strength of its security record (Congress; OMB as interim guidance; the European Commission and EU Member States; model text in Appendix E, E.1, a lineage-plus-reliance gate rather than a universal ban).
- Mandate independent technical, operational, and democratic audits (Congress to fund and mandate; Inspectors General; GAO).
- Establish domain firewalls between military, intelligence, enforcement, and civil-benefits systems (Congress; agencies).
- Guarantee notice, explanation, and appeal rights for affected people (agencies; state legislatures; EU Member States).

All fifteen recommendations, with owners, are in Section 14.

AI systems developed or validated in military, intelligence, and security settings increasingly shape how governments process information, prioritize cases, identify risks, and coordinate

action. In Ukraine and NATO, AI-enabled data fusion has shown operational value for intelligence analysis, battlefield coordination, demining, logistics, and war-crimes evidence processing.^{1,2,3} The same features that make these systems valuable in conflict become dangerous when carried into ordinary civil administration: speed, suspicion, fusion, targeting, and decision compression.

Civil governance rests on a different foundation: legality, proportionality, due process, contestability, reversibility, equal treatment, public justification, and accountable human judgment. Systems optimized for wartime or security use do not automatically meet those standards.

Battlefield validation is not civil authorization.

Since this brief was first issued, the pathway it warned about has become more concrete. NATO's version of Palantir's Maven Smart System reached full operational capability and classified-network accreditation; a March 9, 2026 Pentagon memo directed that Maven become a program of record by September 30, 2026; DHS has opened a \$1 billion, department-wide purchasing agreement with the same vendor; and courts have found or restrained unlawful flows of tax and Medicaid data into immigration enforcement.^{4,5,7,8,9} [CONFIRMED] Whether Maven actually became a program of record by the deadline is unverified. [OPAQUE] Pentagon officials have also said, as reported by DefenseScoop, that Maven's U.S. user base doubled this year.⁶ [REPORTED] Meanwhile, the main U.S. federal AI governance memoranda and the EU AI Act both carve out national-security systems,^{10,11,12} [CONFIRMED] so the hand-off from security to civil use is where scrutiny is weakest. [INFERRED]

The model bills apply their baseline safeguards (legal authority, a published dossier, testing, and appeal) to any AI system used in decisions about people, whatever its origin. Security lineage adds a further layer of scrutiny for four reasons set out in Section 03: such systems were validated on different populations with different error costs; the documentation behind them may be classified or proprietary; a battlefield or classified-network record can be offered as procurement evidence for a civil task it never tested; and they are often built for the data fusion and sharing that civil purpose limits are meant to restrain. [INFERRED]

Analysis focal points

- Ukraine's wartime AI ecosystem and the military-AI data flywheel.
- NATO's Maven Smart System and AI-enabled command infrastructure.
- Domestic pathways into immigration, policing, welfare, and tax.
- China as a cautionary comparison for weak contestability and fusion.

Core translation principles

01 Reversibility

02 Contestability

03 Feedback integrity

04 Lawful authorization

05 Domain separation

Section 02: What Has Changed Since June 2026

The original brief was current to June 1, 2026. The developments below, current to October 7, 2026, bear directly on its argument. Only items that could be traced to an official document or credible reporting are included.

- **June 2026** [REPORTED]

Ukraine's Ministry of Defence reports more than 100 companies using the Palantir-based Brave1 Dataroom to train military AI on real combat data.

*The wartime data flywheel is scaling.*¹³

- **June 2026** [CONFIRMED]

ICE places a ~\$45.8 million Palantir order to modernize HSI case management into an "Enterprise Lakehouse" interoperable with CBP, DOJ, and FBI systems.

*Cross-domain fusion is now written into procurement requirements.*¹⁴

- **June 22, 2026** [CONFIRMED]

NATO's Maven Smart System reaches full technical operational capability and receives full security accreditation for NATO's classified network, clearing rollout to all Allied Command Operations subordinate headquarters.

*Allied AI command infrastructure is now fully operational.*⁴

- **June 23, 2026** [CONFIRMED]

Australia's Federal Court approves an additional A\$548.5 million Robodebt settlement (A\$475 million in compensation); total costs exceed A\$2.4 billion.

*Automated administrative harm has a long and costly tail.*¹⁵

- **July 15, 2026** [CONFIRMED]

Ukraine's Defence Minister Mykhailo Fedorov, who launched the Brave1 Dataroom and hosted Palantir in May, is removed in a cabinet reshuffle; Yevhenii Khmara is confirmed as minister on August 19.

Wartime data partnerships outlast the officials who negotiate them; their terms need to be durable and public.^{16,17}

- **July 17, 2026** [CONFIRMED]

Federal officials admit CMS shared Medicaid data with ICE beyond a court order; a state attorneys general motion, citing discovery documents, says ICE passed it to Palantir, and an ICE declaration says about six users still held copies.

*Purpose limitation and deletion are hard to enforce once data is fused.*⁹

- **July 27, 2026** [CONFIRMED]

The EU “AI Omnibus” enters into force: Annex III high-risk obligations (which cover law enforcement, migration and border control, and access to public benefits) now apply from December 2, 2027 instead of August 2, 2026.

The main rights-sensitive EU safeguards arrive later than planned.^{18,19,20}

- **August 2026** [CONFIRMED]

The Pentagon names a Maven program director in the CDAO (August 5).

*Military AI command infrastructure is becoming permanent.*²¹

- **Late August 2026** [REPORTED]

Financial press reports that Maven’s program-of-record designation took effect; no DoD confirmation was located, so whether it did is unverified.

*Permanence claims need an official record.*²²

- **August 2026** [CONFIRMED]

The UK Information Commissioner publishes audits of police facial recognition in five forces in England and Wales, making 107 recommendations on oversight, record-keeping, image sourcing, and bias.

*Even regulated police AI shows governance gaps that audits must catch.*²³

- **September 8, 2026** [CONFIRMED]

The D.C. Circuit upholds the block on the IRS–ICE taxpayer-address exchange, finding it “indisputably contravenes” federal tax-privacy law.

*Tax data is a live domestic migration pathway.*⁸

- **September 2026** [CONFIRMED]

Ukraine’s Ministry of Defence announces that it and Brave1 ran standardized trials of AI terminal-guidance modules (September 8).

Battlefield autonomy is advancing quickly.^{61,24}

- **September 2026** [REPORTED]

By the Ministry’s account, six of seven guidance modules passed, and successful AI-guided strikes rose tenfold since January; no underlying figures are published.

*Performance claims are outrunning independent verification.*⁶¹

- **September 22, 2026** [REPORTED]

DefenseScoop reports remarks by Deputy Under Secretary James Mazol (Maven users rose from ~50,000 to 100,000+ in 2026) and CDAO Cameron Stanley (Maven helped

strike 13,000 targets in 38 days during Operation Epic Fury); DoD has published no underlying data.

*This is the operating logic civil agencies may be pressed to import.*⁶

- **September–October 2026 [CONFIRMED]**

In a Maine lawsuit, the government acknowledges in a September 18 filing that an HSI agent's case-management entry on an observer of immigration arrests led to a March 2026 border stop; the court has not ruled.

*A real-world analogue of the risk-log problem.*¹⁴

Earlier developments the original omitted, now incorporated

- WIRED's March 2026 reporting on a Palantir-built IRS audit-selection pilot.
- Executive Order 14179 (January 2025), the basis for current federal AI guidance.
- Executive Order 14243 on eliminating federal "information silos" (March 2025).
- OMB memoranda M-25-21 and M-25-22 (April 2025), which exclude national security systems.
- America's AI Action Plan (July 2025).
- The November 2025 district-court ruling against IRS–ICE address sharing.
- Executive Order 14365 on state AI laws and the DOJ AI Litigation Task Force (December 2025).
- OMB M-26-04, which applies "Unbiased AI Principles" to federal procurement of large language models (December 2025).
- California's automated decision-making regulations for businesses, with obligations from January 1, 2027.
- DOJ's intervention against Colorado's AI Act and Colorado's replacement statute, SB26-189 (April–May 2026).
- Germany's 2023 Constitutional Court ruling against police automated data analysis.

Still unresolved. As of the September 2026 congressional tracker reviewed, no federal statute preempting state AI laws had been enacted; H.R. 5388 remained at the introduced stage.³⁶ No published Commerce Department evaluation of state AI laws, which EO 14365 required by March 2026, could be located for this revision.³⁰ [OPAQUE]

Section 03: Operational Success Does Not Confer Civil Legitimacy

AI deployed in conflict zones has shown real operational value in data fusion, intelligence analysis, battlefield coordination, evidence processing, and decision support. But a system that proves useful in war, intelligence, or security operations does not thereby become lawful, legitimate, or publicly authorized for immigration enforcement, policing, welfare administration, tax compliance, or other rights-sensitive civil domains.

The central risk is not that military AI exists. It is that systems built or validated for conflict, intelligence fusion, border enforcement, or national-security analysis may migrate into ordinary civil administration without separate legal authorization for the civil use.

In the model bills, that authorization is a public legal act, not a legislative vote on every system: a statute or regulation, or a written authorization by the agency head published in the Federal Register or state register, which cannot override a statute restricting the data or the use. It must come with a published dossier. At the federal level the dossier goes to the Inspector General and a summary to Congress; at the state level, systems that feed adverse decisions also need independent testing and State Auditor review, and security-derived systems need notice to the legislature (Appendices E and F).

Wartime and security environments reward speed, correlation, anomaly detection, fused visibility, operational prediction, and rapid prioritization across fragmented data streams. Civil governance requires different values: legality, proportionality, notice, due process, contestability, reversibility, equal treatment, public justification, and human accountability. A system can be operationally useful in war and still be democratically unsuitable for welfare eligibility, tax enforcement, immigration case management, policing, or benefits administration.

Wartime or security logic	Civil governance requirement
Speed and decision compression	Deliberation, notice, and due process
Predictive suspicion and targeting	Legality, proportionality, and equal treatment
Broad data fusion across streams	Purpose limitation and data minimization
Operational secrecy	Public justification and contestability
Mission-driven optimization	Rights-sensitive review and appeal
Command hierarchy	Democratic accountability

Why security lineage warrants additional scrutiny. Every system that feeds adverse decisions about people should meet the same baseline: legal authority, a published dossier, testing, and a way to contest the result. Security lineage adds four risks on top of that baseline. First, validation transfers poorly. A model tuned on combatants, aerial targets, or intelligence subjects was tested on a different population, with different base rates and a different balance of error costs; where missing a threat is judged worse than a false alarm, a system may be tuned toward false alarms, and in benefits or policing the false alarm is the harm. Second, the evidence is harder to check: training data, accreditation records, and test results may be classified or held by the vendor, so auditors and affected people may be unable to see the basis for trusting the system. Third, the security record can become the procurement argument, with battlefield or classified-network validation offered as proof of reliability for a civil task it never tested. Fourth, security platforms are often built to fuse and share data across sources, which opens surveillance and enforcement pathways that civil purpose-limitation rules were written to close. Shared vendors matter because they make these risks cheaper to import, not because a shared vendor is a risk in itself. [INFERRED]

Even an accurate system can be inappropriate if its purpose, data sources, institutional context, or review procedures conflict with civil-rights protections.

Section 04: Ukraine as a Military-AI Data Flywheel

Ukraine's wartime use of AI-enabled data fusion shows the operational value of rapidly integrating drone footage, battlefield reports, signals, imagery, logistics data, demining records, and war-crimes evidence. In this setting, speed and integration can save lives, preserve evidence, and improve command decisions.

Brave1 Dataroom. Launched on January 20–21, 2026 by Ukraine's Ministry of Defence with the Ministry of Digital Transformation, the Armed Forces, the Defence Intelligence Research Institute, and Palantir, the Dataroom is a secure environment, built on Palantir software, where Ukrainian defense developers train, test, and validate models on real combat data. The initial datasets are visual and thermal imagery of aerial targets such as Shahed-type drones, and access requires a mandatory security-compliance procedure.^{37,38} [CONFIRMED] By May 2026 then-Defence Minister Mykhailo Fedorov said more than 100 companies were training over 80 detection and interception models, and the Ministry repeated the 100-company figure in June; neither figure has been independently verified.^{13,39} [REPORTED] Officials have said the Dataroom may later serve as a channel for sharing battlefield-tested algorithms with allies.³⁸ [REPORTED] The result is a data flywheel: battlefield data improves models, improved models support operations, and operations generate more data.

Palantir's role. Palantir's work in Ukraine began after CEO Alex Karp met President Zelenskyy in Kyiv in June 2022. TIME reported in 2024 that the company supplied its software free of charge and that more than half a dozen Ukrainian agencies used it for targeting, war-crimes evidence, demining, refugee resettlement, and anti-corruption work.¹ [REPORTED] Palantir and Ukraine's Ministry of Economy later signed a formal demining partnership.⁴⁰ [CONFIRMED] In April 2023, Reuters reported that Palantir would help the Prosecutor General's office pool and analyze evidence related to the more than 78,000 war crimes reported since the invasion, initially without charge.² [CONFIRMED]

On May 12, 2026, Palantir CEO Alex Karp met Zelenskyy and Fedorov in Kyiv.^{41,42} [CONFIRMED] Zelenskyy said they discussed technology "in the context of combat operations and civilian needs"; Fedorov said the existing joint work already includes air-attack analysis, AI processing of large volumes of intelligence data, and the integration of Palantir technology into planning deep-strike operations inside Russia.^{39,41,42} [REPORTED] The Ministry's own readout says the parties "explored possible areas for further cooperation"; no new agreement was announced. [CONFIRMED] Fedorov left office two months later in a cabinet reshuffle.¹⁶ [CONFIRMED]

The work is moving quickly. On September 8, 2026, the Ministry of Defence announced that it and Brave1 had run standardized trials of AI terminal-guidance modules against moving ground targets; the Ukrainian outlet The Defender also reported the release.^{61,24} [CONFIRMED] By the Ministry's own account, six of seven manufacturers were recommended for procurement, and a human still makes the decision to strike.⁶¹ [REPORTED] The Ministry also reported a tenfold

increase since January in successful strikes using AI guidance; it has not published the underlying figures, and the claim has not been independently verified. [REPORTED]

Questions that outlast the battlefield

- Who owns derivative models trained on wartime data?
- What limits apply to vendor reuse or adaptation?
- Can wartime ontologies be repurposed for civil administration?
- How are data sovereignty and public control preserved?
- What prevents emergency infrastructure from becoming permanent civil infrastructure?

No public disclosure of the contractual terms on derivative-model ownership or vendor reuse of Dataroom data could be located. [OPAQUE] The concern is not that Ukraine should be denied operational tools during war. It is that wartime validation may later be treated as proof of suitability for civil governance, even though the legal and ethical standards are fundamentally different.

Section 05: NATO, Maven, and the Normalization of AI-Enabled Command Infrastructure

NATO's adoption of AI-enabled command tools reinforces the normalization of AI as decision-support infrastructure. On March 25, 2025, the NATO Communications and Information Agency and Palantir finalized acquisition of the Maven Smart System NATO for Allied Command Operations, which was expected to begin using it within 30 days. NATO called it one of the fastest procurements in its history, at six months from requirement to acquisition, and SHAPE described it as a sole-source buy.^{3,43} [CONFIRMED]

On June 22, 2026, MSS NATO reached full technical operational capability; NATO's Security Accreditation Board also accredited it for NATO's classified network, and NATO says the platform supports multiple AI models and will be extended across all ACO subordinate headquarters.⁴ [CONFIRMED]

In the United States, a March 9, 2026 memorandum from Deputy Secretary of Defense Steve Feinberg directed that Maven become a formal program of record by the end of fiscal year 2026. It moved system administration and oversight from the National Geospatial-Intelligence Agency to a CDAO Maven program office within 30 days, assigned authorizing-official duties to Research and Engineering, and moved future contracting to the Army Enterprise Agreement.^{5,44,45} [CONFIRMED] In August the Pentagon named a Maven program director, and FY2027 budget materials sought more than \$1.5 billion to expand access.²¹ [CONFIRMED] Financial press reported in late August that the program-of-record designation had taken effect.²² [REPORTED] Whether it did is unverified: no official DoD confirmation was located, and DoD officials on September 22 still described the memo as a direction. [OPAQUE] On September 22, DefenseScoop reported remarks by two Pentagon officials at a DefenseScoop-hosted conference: Deputy Under Secretary of Defense for Research and Engineering James Mazol said Maven's user base had grown from about 50,000 in January to more than 100,000, and Chief Digital and AI Officer Cameron Stanley said Maven helped U.S. forces strike 13,000

targets in 38 days during Operation Epic Fury, the 2026 U.S. strike campaign against Iran.⁶ DoD has published no data behind either figure, and neither has been independently verified; both rest on this single press account. [REPORTED]

This matters because military adoption creates downstream institutional pressure. Once AI-enabled command infrastructure is normal in defense, civilian agencies may seek similar tools for ranking, prioritization, anomaly detection, fraud detection, enforcement targeting, and case management. The same vendor's platforms now span U.S. targeting, NATO command, Ukraine's wartime data environment, and DHS-wide enforcement software.^{3,5,7,37} [CONFIRMED] Shared vendor lineage is not evidence that data moves between these domains, but it lowers the technical and contractual cost of migration. [INFERRED]

A fair objection is that no battlefield-validated model has been shown to have crossed into a civil system. That is correct, and the brief's claim should be read precisely. The pathways are documented: the same vendor platforms already run in military targeting, allied command, wartime data environments, and civil enforcement.^{3,5,7,37} [CONFIRMED] The harms of data fusion in civil enforcement are documented, from unlawful tax-data transfers to Medicaid data shared beyond a court order and case-management entries that the government says led to an observer's border stop.^{8,9,14} [CONFIRMED] No documented instance of a battlefield-validated model being transferred into civil administration was located for this brief. The contract terms, model lineage, and reuse rights that would show such a transfer could not be found in public sources. [OPAQUE] The brief's claim does not depend on whether a specific battlefield-validated model has crossed into a civil agency. It depends on whether the authorization gap exists and whether it should be closed: the gap is what makes the documented pathways dangerous, and the documented harms show what fusion looks like when it proceeds without authorization.

That migration should not happen through procurement convenience, vendor expansion, emergency carryover, or technological familiarity. Civil use requires its own legal basis, public justification, and accountability framework.

Section 06: Domestic Migration Pathways: Immigration, Policing, Welfare, and Tax

The domestic migration risk is clearest in enforcement domains such as immigration and policing, where AI-enabled data fusion can combine identity records, location signals, law-enforcement databases, prior incidents, financial information, travel records, and case files. In these settings a person may become visible to the state as a risk profile before having any meaningful chance to challenge the underlying data or inference.

Immigration enforcement. Palantir has been an ICE contractor since 2011 and has supported ICE's Investigative Case Management (ICM) system since 2014; ICE's FALCON Search & Analysis system ingests DHS, other-agency, and commercial data for investigators.^{46,47} [CONFIRMED] In April 2025, ICE added about \$30 million to an existing Palantir ICM contract to build ImmigrationOS, covering enforcement targeting and prioritization, near-real-time "self-

deportation” tracking, and removal logistics, with a prototype due September 25, 2025 and the contract running at least through September 2027.⁴⁶ [CONFIRMED] On September 25, 2025, ICE awarded a further ~\$29.9 million sole-source order for continued ImmigrationOS licenses and maintenance.⁴⁸ [CONFIRMED]

The footprint has since grown. In February 2026, DHS signed a five-year, \$1 billion blanket purchase agreement for Palantir software across the department.⁷ [CONFIRMED] WIRED also reported a Palantir tool, ELITE, that maps potential deportation targets using DHS and Department of Health and Human Services data.⁷ [REPORTED] In June 2026, ICE ordered a ~\$45.8 million modernization that merges case management and investigative analytics into an “Enterprise Lakehouse” built to interoperate with CBP, DOJ, and FBI systems.¹⁴ [CONFIRMED] No public documentation was located of which datasets feed ELITE or how its address-confidence or targeting outputs are validated. [OPAQUE]

Tax and health data. Executive Order 14243 (March 2025) directed agencies to remove barriers to interagency data access.²⁷ [CONFIRMED] Two rights-sensitive civil datasets then moved toward enforcement. In *Center for Taxpayer Rights v. IRS*, a federal district court held in November 2025 that the IRS’s August 2025 disclosure of roughly 47,000 taxpayer addresses to ICE was unlawful, and in September 2026 the D.C. Circuit upheld the block, noting ICE had sought addresses for nearly 1.3 million people through an automated procedure without individual review.^{8,29} [CONFIRMED] Separately, federal officials admitted in a California case that CMS had shared Medicaid data with ICE beyond what the court’s order allowed, and the judge paused CMS–ICE data sharing for enforcement.⁹ [CONFIRMED] A July 2026 motion by more than 20 state attorneys general, citing documents produced in discovery, says ICE then shared that data with Palantir; Palantir told NPR the dataset was purged at the government’s instruction.⁹ [REPORTED] Palantir’s footprint also extends into tax compliance itself: WIRED reported in March 2026 that the IRS paid Palantir \$1.8 million to improve a pilot “Selection and Analytic Platform” to identify the “highest-value” cases for audit, collection, and potential criminal investigation.²⁵ [CONFIRMED] No public disclosure was located of how that tool’s case-selection criteria are validated or whether they are disclosed to taxpayers. [OPAQUE]

Policing. The Los Angeles Police Department ended its LASER program in April 2019, after the Police Commission’s Inspector General found inconsistent criteria for designating “chronic offenders,” weak oversight, and insufficient data to measure effectiveness.^{49,50} [CONFIRMED] Systems designed to prioritize risk in this way can embed prior enforcement patterns into future enforcement attention. [INFERRED] Germany offers a constitutional parallel: in February 2023 its Federal Constitutional Court struck down Hesse’s and Hamburg’s laws authorizing automated police data analysis (Hesse’s system was built on the hessenDATA platform) because the powers lacked adequate thresholds.³⁵ [CONFIRMED] In August 2026, the UK Information Commissioner reported audits of facial-recognition use by five police forces in England and Wales, finding a “mixed picture” and making 107 recommendations on senior oversight, records of data sources and sharing, image retention, and accuracy and bias checks.²³ [CONFIRMED]

Welfare and tax administration. These domains may seem less coercive than policing or immigration, but automated administrative decisions can still cause severe harm. A person may

lose benefits, face repayment demands, be flagged for investigation, wait for support, or be pushed into a burdensome appeal because a system treated anomaly, correlation, missing documentation, or statistical deviation as evidence of risk.

Robodebt · Australia

Income averaging generated unlawful debt notices between 2015 and 2019. The 2023 Royal Commission called the scheme “crude and cruel” and found it was not legal; a further A\$548.5 million settlement was approved in June 2026.^{15,51} [CONFIRMED]

SyRI · Netherlands

On February 5, 2020, the District Court of The Hague held that the welfare-fraud risk-profiling legislation violated Article 8 of the European Convention on Human Rights for lack of transparency and verifiability.⁵² [CONFIRMED]

MiDAS · Michigan

From 2013 to 2015 the system automatically accused about 40,000 people of unemployment fraud; a review of 22,000 determinations found 93% did not involve fraud. A \$20 million class settlement received final approval in January 2024.^{53,54} [CONFIRMED]

None of these three systems was built or validated in a military, intelligence, or security setting, and none would have been caught by a rule aimed only at security-derived systems. They are included because they show what happens when automated administrative decisions are deployed before legal authorization, effective contestability, and review are in place. That is why this brief argues that authorization has to come first for any AI system in a rights-sensitive domain, with heightened scrutiny when security lineage is present.

Civil harm often appears not as physical force but as delay, denial, debt, investigation, documentation burden, and bureaucratic exhaustion.

Section 07: Domestic Risk Logs and Cross-Domain Fusion

A core danger in military-to-civil translation is building risk profiles from heterogeneous data. Internal or administrative risk logs look technical and neutral, but they can become the infrastructure through which people are classified, prioritized, or targeted.

Hypothetical illustration. This record was constructed by the author to show the structure of a fused risk log. It does not depict any real agency system, vendor product, record, or person.	
event_id	77a9c2-4b1f
model	risk_assessment_v4.2
data_sources	license_plate_readers; financial_records; prior_incident_reports
output	risk_tier = elevated; route_to = enforcement_queue
human_review	reviewer_id recorded; rationale field blank

In this illustration, three sources collected for different purposes are fused into a single score that routes a case to enforcement. The log records that a reviewer touched the case but not what the reviewer judged. Systems built this way can create a surveillance flywheel that bypasses ordinary civil authorization.

A real-world analogue. In a federal lawsuit in Maine, the government acknowledged in a September 18, 2026 filing that an HSI agent's entry of an observer's name and vehicle information into Palantir's ICM system, made during a January 2026 enforcement operation, led to a March 2026 border stop; the plaintiffs say they faced enhanced screening while returning from Canada. The court has not ruled on the claims. DHS's 2016 privacy assessment shows ICM subject records are published to CBP's TECS screening platform.¹⁴ [CONFIRMED] No public documentation was located of how derived analytical links are corrected or deleted in ICE's new architecture when a source record is challenged. [OPAQUE]

The governance question is not only whether each data source was lawfully collected. It is whether the combined system creates a new decision-making architecture that was never separately authorized, explained, audited, or made contestable.

Section 08: China as a Cautionary Comparison: Data Fusion, Blacklists, and Weak Contestability

China's public-security and social-governance architecture is a cautionary comparison. Not every democratic deployment resembles it, but it shows where things lead when broad data fusion, weak contestability, administrative consequences, and state access are combined.

Human Rights Watch's 2019 reverse-engineering of the police app linked to Xinjiang's Integrated Joint Operations Platform (IJOP) documented a system that aggregates personal, travel, vehicle, phone, and location data, flags ordinary behavior as suspicious, and assigns investigative tasks to officers.⁵⁵ [CONFIRMED] China's social credit system, by contrast, is better understood not as a single universal citizen score but as a fragmented set of sectoral databases, administrative blacklists and redlists, and joint sanctions aimed mainly at businesses; local personal-scoring pilots were curtailed or made voluntary and reward-only.⁵⁶ [CONFIRMED]

The lesson is the institutional pattern: when security platforms, administrative blacklists, identity-linked databases, and weak avenues for challenge converge, the line between delivering civil services and exercising civil control weakens. The EU has drawn this line in law; since February 2025 the AI Act has prohibited social scoring and AI-based prediction of individual criminal offending based solely on profiling.^{19,57} [CONFIRMED]

Section 09: Confirmed / Reported / Inferred / Opaque: An Evidentiary Discipline

AI governance debates require evidentiary discipline. This framework separates established facts, claims that have been made on the record but not independently verified, reasonable

analytical conclusions, and unknowns hidden by procurement secrecy, classification, vendor confidentiality, or agency non-disclosure.

[CONFIRMED]	Publicly documented facts: official records and formal acts, court findings, procurement materials, or credible reporting based on documents or multiple sources. For an official act, the act itself is documented; a performance or outcome figure is Confirmed only when it comes from a completed review, audit, or adjudication. Establishes the factual baseline.
[REPORTED]	Claims made on the record by an official, agency, or vendor about its own figures, performance, or activities, or carried by a single press outlet, that no one has independently verified. The claim was made; its accuracy is not established.
[INFERRED]	Reasonable conclusions drawn from known capabilities, institutional incentives, or deployment patterns. Supports risk analysis without overclaiming.
[OPAQUE]	Material facts that could not be located in public sources, typically because of classification, proprietary terms, or agency non-disclosure. The text says “not located” where that is all a search can show; a failed search alone does not establish secrecy. Identifies where disclosure or audit is needed.

Where a figure or capability claim rests only on the claimant’s own statement or on a single press account, the text names who made it and the claim is tagged Reported. Opaque is reserved for material facts that could not be located in public sources.

What the tags verify. A tag records what the cited source establishes, not more. A Confirmed tag on an official act, such as a designation, contract award, memo, or court ruling, means the act is documented in the cited record. It does not mean that the performance claims made around that act are true: when an agency or vendor says how well its own system works, the act of saying so may be documented, but the claim is tagged Reported. Outcome figures, such as the MiDAS error rate, are Confirmed only when they come from a completed review, audit, or adjudication, and the tag then confirms the figure as that record states it, not an independent re-analysis. Inferred marks this brief’s own reasoning. Opaque marks facts this brief looked for and could not find in public sources; it is not, by itself, a finding that the facts are secret.

Claim	Tag	Basis
Brave1 Dataroom launched January 2026 on Palantir software	[CONFIRMED]	Ukrainian MoD release ³⁷
100+ firms using the Dataroom by mid-2026	[REPORTED]	Minister’s and MoD statements; not independently verified ^{13,39}
80+ models being trained in the Dataroom	[REPORTED]	Minister’s statement; not independently audited ³⁹
Ownership of derivative models and vendor reuse terms	[OPAQUE]	No public contract terms located
NATO finalized MSS NATO acquisition March 25, 2025	[CONFIRMED]	NCIA/SHAPE release ³

Claim	Tag	Basis
March 9, 2026 memo moves Maven from NGA to CDAO and toward program-of-record status	[CONFIRMED]	Memo as reported by Reuters and DefenseScoop ^{5,44}
Financial press reported Maven's program-of-record designation took effect (August 2026)	[REPORTED]	Single financial-press report ²²
Maven program-of-record designation completed by Sept. 30, 2026	[OPAQUE]	No DoD release located; officials still described the memo as a direction on Sept. 22 ^{6,22}
MSS NATO full operational capability and classified accreditation (June 2026)	[CONFIRMED]	NATO release ⁴
May 12, 2026 Karp meeting with Zelenskyy and Fedorov	[CONFIRMED]	Presidential and MoD readouts ^{41,42}
Palantir technology used in Ukrainian deep-strike planning	[REPORTED]	Then-Minister Fedorov's statement, as reported; not independently verified ^{39,41,42}
New collaboration agreed at the May 12 meeting	[OPAQUE]	No agreement announced; readout says areas were "explored" ⁴²
ImmigrationOS: ~\$30M (April 2025), prototype due Sept. 2025, runs to at least Sept. 2027	[CONFIRMED]	Contract justification as reported by WIRED ⁴⁶
Data sources and accuracy controls behind ELITE targeting	[OPAQUE]	No public documentation located
IRS-ICE address sharing was unlawful	[CONFIRMED]	District court and D.C. Circuit ^{8,29}
Shared vendor platforms lower the cost of military-to-civil migration	[INFERRED]	Analysis of procurement pattern
Robodebt, SyRI, MiDAS produced wrongful adverse outcomes	[CONFIRMED]	Royal Commission, court rulings, settlements ^{51,52,54}
LAPD ended LASER in 2019	[CONFIRMED]	LAPD OIG review; Los Angeles Times ^{49,50}
LASER-type ranking amplifies historical bias	[INFERRED]	This brief's analysis; not an OIG finding
IRS paid Palantir \$1.8M to improve a pilot audit-selection platform	[CONFIRMED]	Contract documents as reported by WIRED ²⁵
How the IRS pilot selects cases and whether criteria are disclosed	[OPAQUE]	No public disclosure of the criteria located

Claim	Tag	Basis
Maven helped strike 13,000 targets in 38 days; user base doubled to 100,000+	[REPORTED]	Officials' remarks in a single DefenseScoop report; no DoD data published ⁶
Ukraine's MoD and Brave1 ran standardized AI terminal-guidance trials (Sept. 2026)	[CONFIRMED]	MoD Ukraine release; The Defender ^{61,24}
Six of seven guidance modules passed	[REPORTED]	MoD Ukraine's own account; not independently verified ⁶¹
Tenfold rise in successful AI-guided strikes since January 2026	[REPORTED]	Ministry claim; underlying figures not published ⁶¹
A battlefield-validated model has been transferred into a civil system	No documented instance located	Pathways and civil fusion harms are documented; no documented instance was located, and the contract terms and model lineage were not found in public sources

This framework prevents both underreaction and overclaiming. It lets policymakers act on demonstrated risks while marking clearly where more transparency is required.

Section 10: Error Propagation and the Need for Oversight Cadence

AI systems used in rights-sensitive domains can be vulnerable to compounding error. Data-quality problems, model drift, biased feedback loops, incomplete human review, and institutional overreliance can amplify initial mistakes over time.

$$E_t = E_0 \cdot \prod_{i=1}^t \max(0, \alpha - \beta h_i)$$

Equation (1) · an illustration of assumptions, not evidence

In words: total error at time t equals the initial error multiplied, at each step, by the system's drift factor (α), reduced by the corrective effect of human oversight at that step (β times h_i). Each step's factor is floored at zero, because oversight can at most eliminate error, not make it negative. With no meaningful review ($h = 0$), error compounds geometrically as $E_0 \cdot \alpha^t$.

The equation illustrates assumptions; it is not evidence. It assumes that error compounds at a steady rate between reviews and that review removes a fixed share of it. The arithmetic cannot show that either assumption holds for any real system, or that review works at all. Its use is to make visible what an oversight design is betting on: if errors do compound between reviews, how often review happens matters as much as whether it happens. Whether they compound, and how much review reduces them, are empirical questions for the testing and audits described in Sections 12 and 13 and in the model bills.

Illustrative example. Hypothetical parameters, not an empirical model. The sliders start at $\alpha = 1.12$, $\beta = 0.35$, and $h = 0.55$, a per-step factor of about 0.93. At those values, error after 16 cycles is about 0.30 with oversight and about 6.1 with none, so unreviewed error ends roughly 20× higher. That ratio follows entirely from the chosen parameters; it does not show that review works, or how well, in any real system. Changing the parameters shows how much the result depends on them.

Illustrative output	Value	Note
Per-step factor	0.93	error shrinks
Error at t=16	0.30	with oversight
If unreviewed	6.1	$h = 0$
Ratio (illustrative)	20.4×	unreviewed ÷ reviewed

Parameter	Default	Meaning
α drift	1.12	1.00 = error carries over unchanged; above 1 it grows
β oversight effectiveness	0.35	How much of the per-step factor a unit of review removes
h human review intensity	0.55	0 = no meaningful review; 1 = full independent review each cycle

Section 11: The Domain-Translation Test

Before any battlefield-validated, security-grade, or enforcement-derived AI system is deployed in civil governance, the responsible agency should apply a five-part domain-translation test. The third column of each card notes where existing law or policy already supplies a partial hook. A system that fails this test should not be deployed in civil administration.

Test	Question	Existing hook (partial)
01 • Reversibility	Can the system be withdrawn, disabled, or separated from civil workflows without institutional collapse?	M-25-21 requires agencies to stop using non-compliant high-impact AI; M-25-22 requires anti-lock-in terms. ^{10,11}
02 • Contestability	Can affected people and independent overseers challenge the data, inference, ranking, or decision?	M-25-21 remedies and appeals; Colorado SB26-189 adverse-decision notices. ^{10,34}
03 • Feedback integrity	Are feedback signals auditable, tamper-resistant, and protected from biased enforcement loops?	M-25-21 ongoing monitoring; NIST AI RMF. ^{10,58}

Test	Question	Existing hook (partial)
04 • Lawful authorization	Is there explicit legal authority for this civil use, rather than authority inherited from a military or security deployment?	IRC 6103 and the Privacy Act, as the IRS–ICE litigation shows. ^{8,59}
05 • Domain separation	Are military, intelligence, enforcement, and civil-benefits systems separated by enforceable technical and legal firewalls?	Privacy Act computer-matching agreements, 5 U.S.C. 552a(o). ⁵⁹

Answer each question for a candidate system. A single failure is disqualifying.

Section 12: Minimum Civil Authorization Standard

Before deployment, the responsible agency should publish a civil authorization dossier, available before procurement lock-in, operational deployment, or irreversible integration into agency workflows. The standard puts into practice portability rights, data minimization, purpose limitation, sunset clauses, and vendor accountability.

Civil authorization dossier

- System purpose
- Legal authority
- Data inventory
- Model description
- Human-review standards
- Contestability procedures
- Audit-access rules
- Vendor obligations
- Data-retention limits
- Portability rights
- Prohibited uses
- Sunset date
- Independent review process

Much of this content overlaps with documents agencies already produce: M-25-21 AI impact assessments, privacy impact assessments, Privacy Act system-of-records notices, and computer-matching agreements.^{10,59} [CONFIRMED] The gap is the trigger. M-25-21 and M-25-22 do not apply to AI used as a component of a national security system,^{10,11} and the EU AI Act excludes systems used exclusively for military, defense, or national security purposes.¹² [CONFIRMED] The dossier requirement should therefore attach at the moment a system or its outputs cross from a national-security setting into civil use; under the AI Act's “exclusively” wording, that crossing should already bring the system back within scope. [INFERRED]

Section 13: Layered Audit Framework: Technical, Operational, and Democratic

Logs can reconstruct actions, but they do not prove judgment. A system may keep a detailed record of what occurred while still obscuring whether officials exercised meaningful independent review. Civil accountability requires three distinct audit layers.

Technical audit

Examines model performance, data quality, security, bias, drift, and robustness.

Does the system function as claimed? Are errors measurable and correctable?

Independent testers with code, model, and data access; agency CAIO

Operational audit

Examines how officials use the system inside real workflows.

Are humans independently reviewing outputs, or merely ratifying them?

Inspectors General; GAO; state auditors

Democratic audit

Examines legality, public authorization, rights impact, and institutional legitimacy.

Should this system be used in this domain at all?

Congress and state legislatures; courts where rights are litigated

A model may be technically functional and still democratically unauthorized.

Section 14: Policy Recommendations

Each recommendation names the actor best placed to act and, where one exists, the framework it can build on. “Agencies” means civil departments and their Chief AI Officers.

#	Recommendation	Owner	Build on
01	Require a civil authorization dossier before deployment.	OMB (require); agencies (produce)	M-25-21 impact assessments; PIAs; model federal text in Appendix E (E.2); model state bill in Appendix F (F.1)

#	Recommendation	Owner	Build on
02	Require separate legal authorization before a security-derived system is used in civil administration on the strength of its security record.	Congress; OMB as interim guidance; European Commission and EU Member States	Closes the national-security-system exclusion in M-25-21/M-25-22; clarifies that AI Act Art. 2(3) covers only “exclusively” military, defence, or national-security uses; model federal text in Appendix E (E.1, a lineage-plus-reliance gate, not a universal ban); model state bill in Appendix F (F.1), which covers any AI used in adverse decisions
03	Mandate independent technical, operational, and democratic audits.	Congress (fund and mandate); IGs; GAO	M-25-21 independent review; Section 13
04	Establish domain firewalls between military, intelligence, enforcement, and civil-benefits systems.	Congress; agencies	Privacy Act matching rules; IRC 6103
05	Require strict data minimization and purpose limitation.	Agencies; procurement officers	Privacy Act; M-25-22 government-data terms
06	Guarantee notice, explanation, and appeal rights for affected people.	Agencies; state legislatures; EU Member States	M-25-21 remedies; Colorado SB26-189; AI Act Art. 86 right to explanation
07	Require human command primacy for consequential civil decisions.	Agencies	M-25-21 human oversight
08	Prohibit opaque risk scores as the sole basis for benefit denial, enforcement action, audit targeting, or other adverse civil decisions.	Congress; state legislatures; EU co-legislators	SyRI and Robodebt lessons; AI Act Art. 5 profiling ban
09	Mandate portability rights to prevent vendor lock-in.	Procurement officers; OMB	M-25-22 lock-in protections
10	Require sunset clauses and periodic reauthorization.	Congress; state legislatures	No general requirement today
11	Require independent red-teaming before deployment in rights-sensitive domains.	Agencies; procurement officers	M-25-21 pre-deployment testing; NIST AI RMF
12	Preserve data-sovereignty protections for wartime or emergency datasets.	Ukraine’s MoD and partner governments; vendors by contract	Dataroom terms (not public)
13	Require post-action reviews for consequential deployments.	Agencies; IGs	M-25-21 ongoing monitoring

#	Recommendation	Owner	Build on
14	Publish prohibited-use rules before system integration.	Agencies; procurement officers	EU AI Act Article 5 as a model
15	Bring forward, or at least not further delay, safeguards for Annex III public-sector uses (migration, law enforcement, benefits).	European Parliament and Council; European Commission; national market-surveillance authorities	AI Omnibus timeline (December 2, 2027)

Three notes on actors

First, OMB's high-impact definition turns on whether AI output is a “principal basis” for a decision, and the memo itself notes AI can sit inside high-impact pipelines without meeting that test.¹⁰ [CONFIRMED] Procurement officers and Chief AI Officers should read the definition broadly for fused risk scores. [INFERRED]

Second, although EO 14365 seeks federal preemption of conflicting state AI laws, it expressly says the legislative proposal it calls for should not preempt state laws on state government procurement and use of AI.³⁰ [CONFIRMED] State legislatures therefore keep clear room to apply these safeguards to their own agencies. Colorado's replacement law, effective January 1, 2027, already reaches decisions about essential government services and public benefits.³⁴ [CONFIRMED]

Third, in the EU the AI Act already classifies AI used by public authorities to assess eligibility for public assistance benefits, many law-enforcement uses, and migration, asylum, and border-control uses as high-risk,²⁰ [CONFIRMED] but the AI Omnibus moved those obligations from August 2, 2026 to December 2, 2027.¹⁸ [CONFIRMED] The European Commission's implementing guidance and Member State market-surveillance authorities are therefore the main levers for treating a system that crosses from defense into civil use as in scope; that reading of Article 2(3) is this brief's, not settled law. [INFERRED]

Section 15: Conclusion

AI systems validated in war may be useful, powerful, and operationally impressive. That does not give them legal authority for civil use. The months since this brief first appeared have made the point less abstract: military AI infrastructure has become more permanent, the same platforms have spread across civil enforcement, and courts have had to stop tax and health data from crossing lines Congress drew.

The lesson is not to reject military AI categorically. It is to preserve the boundary between emergency operational systems and ordinary civil governance. Civil administration requires a higher standard of notice, proportionality, appeal, transparency, and public authorization.

No system built for war, intelligence, security fusion, or enforcement should become part of civil government unless it is separately and publicly authorized for that use.

Corrections

Corrections to the June 2026 Edition

The following factual statements in the original June 2026 brief were corrected in this October 7, 2026 revision. Other changes in this revision are listed in the changelog at the end.

Brave1 Dataroom

The “100+ firms / 80+ models” figures date from May–June 2026, not the January launch. Sources describe visual and thermal aerial-target datasets, not “intercepts”; access follows a security-compliance procedure. Palantir’s role as platform provider was added.

May 12, 2026 meeting

The meeting was with CEO Alex Karp. Deep-strike planning and intelligence processing were described by then-Minister Fedorov as existing areas of cooperation, not a new expansion agreed at the meeting.

Maven memo

The March 9, 2026 memo directed that Maven become a program of record by September 30, 2026, moving oversight from NGA to a CDAO program office. It did not move an existing program of record.

Reuters 2023

Reuters reported that Palantir would help prosecutors analyze evidence connected to the more than 78,000 reported war crimes, not that processing of all 78,000 incidents had been completed.

ImmigrationOS

September 2025 was the prototype deadline, not an extension. The ~\$30 million was a modification to an existing ICM contract, followed by a ~\$29.9 million continuation award in September 2025.

Fedorov’s status

Mykhailo Fedorov was Defence Minister at the time of the January and May 2026 events but left office in July 2026; he is described accordingly.

Maven program-of-record completion

Stated as fact in some secondary coverage, but not confirmed by DoD as of this revision; tagged Opaque.

LAPD LASER

The Inspector General's findings concerned inconsistent criteria, oversight, and lack of data on effectiveness. The point about bias amplification is kept as this brief's own inference and tagged as such.

Sources

References

Sources 1–60 are numbered in order of first citation; later additions follow. URLs were accessed between October 6 and 7, 2026 (UTC). Where a primary source was paywalled or blocked, a secondary report of the same document is listed alongside it. Superscript numbers in the text open the corresponding source.

1. Vera Bergengruen, "How Tech Giants Turned Ukraine Into an AI War Lab," TIME, February 2024.
<https://time.com/6691662/ai-ukraine-war-palantir/>
2. Reuters, "Data company Palantir to help Ukraine prosecute alleged Russian war crimes," Reuters / CNBC, April 22, 2023.
<https://www.reuters.com/world/europe/data-company-palantir-help-ukraine-prosecute-alleged-russian-war-crimes-2023-04-22/>
<https://www.cnbc.com/2023/04/22/data-company-palantir-to-help-ukraine-prosecute-alleged-russian-war-crimes.html>
3. NATO Communications and Information Agency, "NATO acquires AI-enabled warfighting system," NCIA, April 14, 2025 (acquisition finalized March 25, 2025).
<https://www.ncia.nato.int/newsroom/news/nato-acquires-ai-enabled-warfighting-system>
4. NATO Joint Warfare Centre (from SHAPE), "NATO Maven Smart System Achieves Full Technical Operational Capability," JWC NATO, Published August 12, 2026 (milestone June 22, 2026).
<https://www.jwc.nato.int/article/maven-achieves-ftoc/>
5. Brandi Vincent, "DOD components face 'aggressive' timeline for Maven Smart System transition," DefenseScoop, April 15, 2026.
<https://defensescoop.com/2026/04/15/palantir-maven-smart-system-pentagon-program-transition-feinberg/>
6. Jon Harper, "More than 100K personnel use Maven Smart System: Pentagon official," DefenseScoop, September 22, 2026.
<https://defensescoop.com/2026/09/22/maven-smart-system-ai-james-mazol-cameron-stanley-defensetalks/>
7. Makena Kelly, "DHS Opens a Billion-Dollar Tab With Palantir," WIRED, February 19, 2026.
<https://www.wired.com/story/department-homeland-security-ice-billion-dollar-agreement-palantir/>
8. Matt Bracken, "Appeals court keeps block on IRS from sharing taxpayer data with ICE," FedScoop; Center for Taxpayer Rights v. IRS, No. 26-5006 (D.C. Cir. Sept. 8, 2026), September 8, 2026.
<https://fedscoop.com/irs-ice-data-sharing-appeals-court-block/>
<https://law.justia.com/cases/federal/appellate-courts/cadc/26-5006/26-5006-2026-09-08.html>
9. Jude Joffe-Block, "ICE shared Medicaid data it wasn't supposed to have with Palantir," NPR, July 17, 2026.
<https://www.npr.org/2026/07/17/nx-s1-5898504/ice-medicaid-palantir-data>
10. Office of Management and Budget, "M-25-21, Accelerating Federal Use of AI through Innovation, Governance, and Public Trust," White House / OMB, April 3, 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>

11. Office of Management and Budget, "M-25-22, Driving Efficient Acquisition of Artificial Intelligence in Government," White House / OMB, April 3, 2025.

<https://www.whitehouse.gov/wp-content/uploads/2025/02/M-25-22-Driving-Efficient-Acquisition-of-Artificial-Intelligence-in-Government.pdf>

12. European Union, "Regulation (EU) 2024/1689 (AI Act), Article 2 (Scope)," EUR-Lex, 2024.

<https://eur-lex.europa.eu/eli/reg/2024/1689/oj>
<https://artificialintelligenceact.eu/article/2/>

13. Ministry of Defence of Ukraine, "Over 100 Ukrainian companies are already leveraging Brave1 Dataroom to train AI models," MoD Ukraine; Ukrainska Pravda, June 2026.

<https://mod.gov.ua/en/news/over-100-ukrainian-companies-are-already-leveraging-brave1-dataroom-to-train-ai-models>
<https://www.pravda.com.ua/eng/news/2026/06/11/8038856/>

14. Anthony Kimery, "ICE observer lawsuit shows how Palantir records can feed border screening, analytics," Biometric Update, October 4, 2026.

<https://www.biometricupdate.com/202610/ice-observer-lawsuit-shows-how-palantir-records-can-feed-border-screening-analytics>

15. SBS News / AAP, "Robodebt victims share in millions as Australia's largest-ever class action settlement approved," SBS, June 23, 2026.

<https://www.sbs.com.au/news/article/robotdebt-victims-class-action-settlement-approved/sd5arll0g>

16. The Guardian, "Volodymyr Zelenskyy dismisses Ukraine's defence minister on eve of Starmer visit," The Guardian, July 15, 2026.

<https://www.theguardian.com/world/2026/jul/15/volodymyr-zelenskyy-dismisses-ukraines-popular-defence-minister>

17. BusinessDay, "Zelensky faces fresh turmoil as Ukraine appoints new defence minister," BusinessDay, August 19, 2026.

<https://www.businessday.co.za/world/europe/2026-08-19-zelensky-faces-fresh-turmoil-as-ukraine-appoints-new-defence-minister/>

18. European Commission, "AI Omnibus enters into force; Regulation (EU) 2026/1744," European Commission / Official Journal, July 27, 2026.

<https://digital-strategy.ec.europa.eu/en/news/ai-omnibus-enters-force>
<https://eur-lex.europa.eu/eli/reg/2026/1744/oj/eng>

19. European Commission, "AI Act: application timeline and AI Omnibus," Shaping Europe's digital future, Accessed October 6, 2026.

<https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>

20. European Union, "AI Act Annex III, high-risk systems (public services, law enforcement, migration)," AI Act, 2024.

<https://artificialintelligenceact.eu/annex/3/>

21. Brandi Vincent, "Pentagon appoints new Maven Smart System program director in fresh push for C2 integration," DefenseScoop, August 5, 2026.

<https://defensescoop.com/2026/08/05/pentagon-appoints-new-maven-smart-system-program-director/>

22. Patrick Sanders, "Palantir's Maven Is Now an Official Pentagon Program of Record," The Motley Fool, August 25, 2026.

Secondary financial press; no official DoD announcement of completion located as of October 7, 2026.

<https://www.fool.com/investing/2026/08/25/palantirs-maven-is-now-an-official-pentagon-progra/>

23. UK Information Commissioner's Office, "Facial recognition in policing: earning public trust through strong data protection governance," ICO, August 2026.
<https://ico.org.uk/about-the-ico/media-centre/news-and-blogs/2026/08/facial-recognition-in-policing/>
<https://ico.org.uk/media2/uo1cdoxm/police-forces-frt-outcomes-report-202608.pdf>
24. Olha Zakrevska, "The number of successful strikes using AI guidance has increased tenfold," The Defender, September 8, 2026.
<https://thedefender.media/en/2026/09/ai-guided-strikes/>
25. WIRED, "The IRS Wants Smarter Audits. Palantir Could Help Decide Who Gets Flagged," WIRED, March 30, 2026.
<https://www.wired.com/story/documents-reveal-palantir-irs-contract-fraud-clean-energy-credits/>
26. The White House, "Executive Order 14179, Removing Barriers to American Leadership in Artificial Intelligence," White House, January 23, 2025.
<https://www.whitehouse.gov/presidential-actions/2025/01/removing-barriers-to-american-leadership-in-artificial-intelligence/>
27. The White House, "Executive Order 14243, Stopping Waste, Fraud, and Abuse by Eliminating Information Silos," Federal Register, Signed March 20, 2025 (published March 25, 2025).
<https://www.federalregister.gov/documents/2025/03/25/2025-05214/stopping-waste-fraud-and-abuse-by-eliminating-information-silos>
28. The White House, "Winning the Race: America's AI Action Plan," White House, July 23, 2025.
<https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf>
29. U.S. District Court for the District of Columbia, "Center for Taxpayer Rights v. IRS, No. 1:25-cv-00457 (D.D.C.)," Civil Rights Litigation Clearinghouse, November 21, 2025.
<https://clearinghouse.net/case/46138/>
30. The White House, "Executive Order 14365, Ensuring a National Policy Framework for Artificial Intelligence," 90 Fed. Reg. 58499, December 11, 2025.
<https://www.federalregister.gov/documents/2025/12/16/2025-23092/ensuring-a-national-policy-framework-for-artificial-intelligence>
31. Office of Management and Budget, "M-26-04, Increasing Public Trust in Artificial Intelligence Through Unbiased AI Principles," OMB, December 11, 2025.
<https://www.whitehouse.gov/wp-content/uploads/2025/12/M-26-04-Increasing-Public-Trust-in-Artificial-Intelligence-Through-Unbiased-AI-Principles-1.pdf>
32. California Privacy Protection Agency, "California Finalizes Regulations to Strengthen Consumers' Privacy," CPPA, September 23, 2025.
<https://cppa.ca.gov/announcements/2025/20250923.html>
33. Norton Rose Fulbright, "X.AI sues, DOJ intervenes, enforcement of Colorado's AI Act suspended," X.AI LLC v. Weiser, No. 1:26-cv-01515, D. Colo., May 2026.
<https://www.nortonrosefulbright.com/en-us/knowledge/publications/de3ad9de/xai-sues-doj-intervenes-enforcement-of-colorado-ai-act-suspended>
34. Goodwin Procter, "Colorado Enacts Law Repealing and Replacing Landmark Colorado AI Act (SB26-189)," Goodwin, June 10, 2026.
<https://www.goodwinlaw.com/en/insights/publications/2026/06/alerts-technology-fs-colorado-enacts-law-repealing-replacing-landmark-ai-act>
35. Federal Constitutional Court of Germany, "Legislation in Hesse and Hamburg regarding automated data analysis is unconstitutional," BVerfG press release No. 18/2023, February 16, 2023.
<https://www.bundesverfassungsgericht.de/SharedDocs/Pressemitteilungen/EN/2023/bvg23-018.html>

36. Steptoe LLP, "Federal AI Legislative Tracker," Steptoe, September 2026.
https://www.steptoe.com/a/web/dkVL6BDZnGdLmjjGzB3zCm/steptoe-federal-ai-legislative-tracker_september-2026.pdf
37. Ministry of Defence of Ukraine, "Ministry of Defence launches Brave1 Dataroom, a secure environment for training military AI solutions," MoD Ukraine, January 21, 2026.
<https://mod.gov.ua/en/news/ministry-of-defence-launches-brave1-dataroom-a-secure-environment-for-training-military-ai-solutions>
38. Digital State UA, "Ukraine Launches Brave1 Dataroom with Palantir to Train AI Models Using Battlefield Data," Ministry of Digital Transformation of Ukraine, January 2026.
<https://digitalstate.gov.ua/news/tech/ukraine-launches-brave1-dataroom-with-palantir-to-train-ai-models-using-battlefield-data>
39. Meduza, "Palantir CEO visits Kyiv; Ukraine's defense minister says company's technology helps plan deep strikes inside Russia," Meduza (summarizing Defence Minister Fedorov's Telegram post), May 12, 2026.
<https://meduza.io/en/news/2026/05/12/palantir-ceo-visits-kyiv-ukraine-s-defense-minister-says-company-s-technology-helps-plan-deep-strikes-inside-russia>
40. Palantir Technologies, "Palantir and Ministry of Economy of Ukraine Sign Demining Partnership," Palantir investor news, March 4, 2024.
<https://investors.palantir.com/news-details/2024/Palantir-and-Ministry-of-Economy-of-Ukraine-Sign-Demining-Partnership/>
41. Reuters, "Zelenskiy meets Palantir CEO as Ukraine expands use of AI in war," Reuters, May 12, 2026.
<https://www.reuters.com/world/europe/zelenskiy-meets-palantir-ceo-ukraine-expands-use-ai-war-2026-05-12/>
42. Ministry of Defence of Ukraine, "AI and Ukraine's defence strategy: Mykhailo Fedorov and Palantir Technologies discuss cooperation," MoD Ukraine, May 12, 2026.
<https://mod.gov.ua/en/news/ai-and-ukraines-defence-strategy-mykhailo-fedorov-and-palantir-technologies-discuss-cooperation-to-strengthen-security>
43. Tamara Rozouvan, "NATO agrees sole-source procurement of Palantir's Maven," Janes, April 23, 2025.
<https://www.janes.com/defence-intelligence-insights/defence-news/c4isr/nato-agrees-sole-source-procurement-of-palantirs-maven>
44. Reuters, "Exclusive: Pentagon to adopt Palantir AI as core US military system, memo says," Reuters, March 20, 2026.
<https://www.reuters.com/technology/pentagon-adopt-palantir-ai-as-core-us-military-system-memo-says-2026-03-20/>
45. DefenseScoop, "Feinberg's new Maven directive sets AI-enabled decision-making as 'the cornerstone' for CJADC2," DefenseScoop, April 3, 2026.
<https://defensescoop.com/2026/04/03/palantir-maven-feinberg-directive/>
46. Caroline Haskins, "ICE Is Paying Palantir \$30 Million to Build 'ImmigrationOS' Surveillance Platform," WIRED, April 18, 2025.
<https://www.wired.com/story/ice-palantir-immigrationos/>
47. U.S. Department of Homeland Security, "Privacy Impact Assessment DHS/ICE/PIA-032, FALCON Search & Analysis System," DHS, DHS/ICE/PIA-032.
<https://www.dhs.gov/publication/dhsicepia-032a-falcon-search-analysis-system-falcon>
48. U.S. ICE, "Investigative Case Management (ICM) Immigration OS Continued Support," Award 70CTD022FR000170-P000012 (SAM.gov / GovContractFinder), September 25, 2025.

<https://govcontractfinder.com/contracts/investigative-case-management-icm-immigration-os-continued-s-70ctd022fr000170-p000012>

49. Office of the Inspector General, Los Angeles Police Commission, "Review of Selected Los Angeles Police Department Data-Driven Policing Strategies," LAPD OIG, March 2019.
https://www.oig.lacity.org/_files/ugd/b2dd23_21f6fe20f1b84c179abf440d4c049219.pdf

50. Los Angeles Times, "LAPD ends another data-driven crime program touted to target violent offenders," Los Angeles Times, April 12, 2019.
<https://www.latimes.com/local/lanow/la-me-laser-lapd-crime-data-program-20190412-story.html>

51. Prime Minister of Australia, "Final report of the Royal Commission into the Robodebt Scheme," Australian Government; The Guardian, July 7, 2023.
<https://www.pm.gov.au/media/final-report-royal-commission-robodebt-scheme>
<https://www.theguardian.com/australia-news/2023/jul/07/robodebt-royal-commission-final-report-recommends-civil-criminal-prosecutions>

52. District Court of The Hague, "NJCM et al. v. The Netherlands (SyRI), ECLI:NL:RBDHA:2020:1878," Rechtspraak, February 5, 2020.
<https://uitspraken.rechtspraak.nl/details?id=ECLI%3ANL%3ARBDHA%3A2020%3A1878>

53. Benefits Tech Advocacy Hub, "Michigan Unemployment Insurance False Fraud Determinations," BTAH, Case study.
<https://btah.org/case-study/michigan-unemployment-insurance-false-fraud-determinations.html>

54. Michigan Department of Attorney General, "Class Action Settlement Approved in Bauserman v. State of Michigan Unemployment Insurance Agency," Michigan AG, January 30, 2024.
<https://www.michigan.gov/ag/news/press-releases/2024/01/30/class-action-settlement-approved-by-court-of-claims>

55. Human Rights Watch, "China's Algorithms of Repression: Reverse Engineering a Xinjiang Police Mass Surveillance App," HRW, May 1, 2019.
<https://www.hrw.org/report/2019/05/01/chinas-algorithms-repression/reverse-engineering-xinjiang-police-mass>

56. Vincent Brussee, "China's social credit score: untangling myth from reality," MERICS, February 11, 2022.
<https://merics.org/en/comment/chinas-social-credit-score-untangling-myth-reality>
<https://merics.org/en/report/chinas-social-credit-system-2021-fragmentation-towards-integration>

57. European Union, "AI Act Article 5, Prohibited AI practices," AI Act, 2024.
<https://artificialintelligenceact.eu/article/5/>

58. National Institute of Standards and Technology, "AI Risk Management Framework (AI RMF 1.0)," NIST, 2023.
<https://www.nist.gov/itl/ai-risk-management-framework>

59. United States Congress, "Privacy Act of 1974, 5 U.S.C. 552a (including matching agreements)," Cornell LII, As amended.
<https://www.law.cornell.edu/uscode/text/5/552a>

60. European Union, "AI Act Article 86, Right to explanation of individual decision-making," AI Act, 2024.
<https://artificialintelligenceact.eu/article/86/>

61. Ministry of Defence of Ukraine, "Ministry of Defence and Brave1 conduct large-scale testing of drones with AI guidance," MoD Ukraine, September 8, 2026.
<https://mod.gov.ua/en/news/ministry-of-defence-and-brave1-conduct-large-scale-testing-of-drones-with-ai-guidance>

Model Legislative Language

Note

Drafting aids, not legal advice. The constitutional and statutory cites need verification by counsel in the relevant jurisdiction before introduction. Notes on the cited authorities follow the model text. E.1 is a lineage-plus-reliance gate; E.2 and the state version in Appendix F (F.1), which is built on E.2's broader trigger, apply to AI used in rights-sensitive decisions whatever its origin.

E.1: No Civil Use of Security-Validated AI Without Separate Legal Authorization (a Lineage-Plus-Reliance Gate)

Note

Scope: E.1 is not a ban on security-derived AI in civil use. It covers a system only if the system has security-domain lineage and its security validation, accreditation, or use is cited or relied on for civil use, with a rebuttable presumption based on the vendor's product line (§2(d)). That limit is intentional: E.1 targets the case where a security track record is offered as a reason to trust a system in civil administration. A security-derived system that does not meet the reliance element is still covered by the general rules for rights-impacting AI, including OMB M-25-21's requirements for high-impact AI, and by E.2 where it is adopted; at the state level, F.1 covers it whatever its origin.

Section 1. Findings. Congress finds that:

- (1) AI systems developed, trained, or validated in military, intelligence, or national-security settings may be operationally effective and democratically unsuitable for civil administration;
- (2) the national-security-system carve-outs in OMB M-25-21 and M-25-22, and the "exclusively" military exclusion in EU AI Act Article 2(3), leave the security-to-civil crossing point inadequately governed;
- (3) existing law, including the Privacy Act, 5 U.S.C. 552a, and the tax confidentiality protections of 26 U.S.C. 6103, provides partial protections but does not require separate authorization for a system whose security-domain validation is cited as a basis for civil deployment; and
- (4) separate legal authorization is required before any such system enters civil use.

Section 2. Definitions. In this Act:

(a) AI system. Any machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments. The term shall be construed consistently with existing federal usage and, where applicable, with definitions adopted by the National Institute of Standards and Technology.

(b) Agency. An executive agency, as defined in 5 U.S.C. 105, for purposes of this Act, and any component thereof.

(c) National-security setting. An element of the intelligence community (50 U.S.C. 3003(4)), the Department of Defense, or a national security system (44 U.S.C. 3552(b)(6)).

(d) Covered system. An AI system, or a model, dataset, or output derived from it, that:

- (1) was developed, trained, fine-tuned, or accredited using data, infrastructure, or test results from a national-security setting; and
- (2) whose security-domain validation, accreditation, or use was cited or relied upon in:
 - (A) the agency's justification for acquiring or deploying it for civil use, including any sole-source justification, authorization to operate, or equivalent record; or
 - (B) the vendor's proposal, marketing materials, or other public statements about the system, its model family, or its product line.

Presumption of coverage. An AI system is presumed to be a covered system if the vendor supplies a system of the same product line or model family to a national-security setting. The agency head may rebut the presumption only by a written determination, published before the system is acquired or deployed for civil use, that states the factual basis for finding that the system does not meet paragraph (1) or does not meet paragraph (2), supported by the vendor's documentation of the system's training data, fine-tuning, and accreditation history. If the vendor does not provide that documentation, the presumption may not be rebutted. That the system is a commercial product (41 U.S.C. 103) may support such a determination but does not by itself rebut the presumption.

(e) Civil use. Deployment, procurement, or integration into workflows of an agency in a rights-sensitive domain.

(f) Rights-sensitive domain. Any of the following, and case management for any of them:

- (1) immigration enforcement;
- (2) immigration adjudication, including asylum, refugee, naturalization, visa, and immigration-benefits determinations;
- (3) border and port-of-entry screening;
- (4) policing;
- (5) criminal justice beyond policing, including pretrial release and detention, sentencing, probation and parole supervision, and corrections;
- (6) welfare and benefits administration; and
- (7) tax administration.

The term does not include (i) internal administrative functions such as payroll, hiring, or facilities management, or (ii) systems used solely for statistical or research purposes that do not produce decisions about identifiable individuals.

(g) Separate legal authorization. A statute or regulation that expressly authorizes civil use of the covered system, or an express written authorization for that use issued by the agency head and published in the Federal Register, together with any Privacy Act system-of-records notice (5 U.S.C. 552a(e)(4)) and computer-matching agreement (5 U.S.C. 552a(o)) required by law. An authorization by the agency head does not by itself constitute separate legal authorization where a statute restricts the collection, use, or disclosure of the data the system relies on, or restricts the use to which the system is put.

(h) Appropriate congressional committees. The committees of the Senate and the House of Representatives with legislative or oversight jurisdiction over the agency concerned.

Section 3. Prohibition. No agency may obligate funds for, deploy, or integrate into civil workflows any covered system unless the agency has obtained separate legal authorization for civil use and published a civil authorization dossier under Section 4. This section does not apply to a system for which the presumption in Section 2(d) has been rebutted by a published determination and that does not otherwise meet Section 2(d).

Section 4. Civil Authorization Dossier. The dossier shall include, at minimum:

- (1) system purpose, stated in plain language;
- (2) the specific legal authority for civil use;
- (3) a data inventory identifying every data source, its collection purpose, and its legal basis;
- (4) a model description, including architecture, training-data provenance, known limitations, and validation results;
- (5) human-review standards, including when a human must review, what the reviewer must document, and what the reviewer may override;
- (6) contestability procedures, including notice to affected persons, appeal timelines, and the mechanism for challenging data, inference, and decision;
- (7) audit-access rules, identifying who may inspect the system, at what layer, and under what authority;
- (8) vendor obligations, including data-return, deletion, portability, non-reuse, and cooperation with auditors;
- (9) data-retention limits, per data class, with a deletion-verification method;
- (10) portability rights, describing how data and models transition if the vendor relationship ends;
- (11) prohibited uses, enumerated by policy; and
- (12) a sunset date, not to exceed three years from publication, after which the authorization expires unless reauthorized under this section.

The dossier shall be reviewed by the agency's Chief AI Officer and legal counsel, and submitted to the agency's Inspector General. A summary shall be transmitted to OMB and to the appropriate congressional committees within 30 days of publication. Classified material may be filed in a classified annex, provided that the unclassified portion is sufficient for meaningful public and congressional oversight of the system's purpose, authority, data sources, and review procedures.

Reauthorization requires an updated dossier and independent review. Before the sunset date, the agency shall publish an updated dossier under this section and obtain an independent review of the system's performance in civil use, including its error rates, complaints and appeals, and any changes since the prior authorization, by an evaluator with no financial interest in the vendor. The review shall be summarized in the updated dossier and submitted

with it to the Inspector General. A covered system that has not been reauthorized by the sunset date is treated as lacking separate legal authorization.

Section 5. Domain Firewalls. The agency shall maintain enforceable technical and legal separation between military, intelligence, enforcement, and civil-benefits systems, including data-minimization, purpose-limitation, and audit-access controls. Nothing in this section prohibits cross-agency coordination that is (1) authorized by statute, (2) documented in a written agreement, and (3) subject to audit.

Section 6. Enforcement.

(a) Limitation on funds. An obligation, option exercise, or modification for a covered system without separate legal authorization and a published dossier is an obligation in excess of available appropriations, prohibited by 31 U.S.C. 1341(a)(1)(A), and reportable under 31 U.S.C. 1351.

(b) Oversight. Within 180 days of enactment, each agency shall report all covered systems, including withdrawn systems and systems operating under a waiver, to Congress and to the Comptroller General. The Comptroller General shall review the reports within one year and report findings to Congress. Classified details shall be filed in an annex to the relevant committees, with an unclassified summary.

(c) Complaints and judicial review. Any person, including an individual affected by a covered system or an organization representing such individuals, may file a written complaint with the agency's Chief AI Officer alleging that the agency has deployed, used, or continued to use a covered system in violation of this Act, including without separate legal authorization, without a published dossier or with a dossier that omits a required element or is materially inaccurate, in violation of Section 5, or after its authorization has expired under Section 4. The complaint shall identify the system, or describe it as specifically as the complainant can, state the facts supporting the allegation, and give a means of contacting the complainant. Within 90 days, the agency shall issue a written decision that states its findings on each allegation, the corrective action it will take or its reasons for denying the complaint, and the records on which it relied, which it shall make available to the complainant subject to Section 4's classified-annex provision and other law. An acknowledgment of receipt, a statement that the complaint is under review, or any other response that does not meet this subsection is not a decision. If the agency does not issue a decision meeting this subsection within 90 days, the complaint is deemed denied, and the denial is final agency action. A decision on a complaint, a deemed denial, and final agency action to deploy or continue a covered system in violation of this Act are reviewable under 5 U.S.C. 702–706. Relief is declaratory and injunctive only. Fees are available under the Equal Access to Justice Act, 28 U.S.C. 2412. Nothing in this subsection waives sovereign immunity beyond the waiver already provided in 5 U.S.C. 702.

(d) Waiver. The following rules govern waivers under this Act.

(1) Waiver. The agency head may waive Sections 3, 4, 6(a), and 7 for a specific covered system for not more than 180 days upon a written determination that the waiver is necessary to avert an imminent threat to human life.

(2) One renewal. A waiver may be renewed once, for not more than 180 days, upon a new written determination. No further waiver may be issued for the same system and purpose.

(3) Notice. Within 7 days of issuing or renewing a waiver, the agency head shall transmit the full determination, with any classified material in a classified annex, to the appropriate congressional committees and to the agency's Inspector General, and shall publish an unclassified summary.

(4) Limits. A waiver does not suspend Section 5, subsection (b) or (c) of this section, paragraph (3) of this subsection, or Section 8, and does not authorize any use otherwise prohibited by law.

(e) Transition. Notwithstanding Sections 3, 6(a), and 7, a covered system in use on the date of enactment may remain in use for 12 months after that date, during which the agency shall obtain separate legal authorization and publish a dossier under Section 4. Renewals and option exercises under contracts in effect on the date of enactment that occur within that 12-month period fall within this transition. New awards, and material modifications of existing contracts, made after the date of enactment are covered immediately. A covered system for which authorization has not been obtained and a dossier published by the end of the 12-month period shall be withdrawn from civil use.

Section 7. Appropriations. No funds appropriated or otherwise made available to any agency may be obligated or expended for the procurement, deployment, or integration of a covered system unless separate legal authorization has been obtained and the civil authorization dossier has been published.

Section 8. Rule of construction and savings. Nothing in this Act shall be construed to authorize any use of an AI system that is otherwise prohibited by law. Nothing in this Act shall be construed to supersede, limit, or diminish the Privacy Act, 5 U.S.C. 552a, the tax confidentiality protections of 26 U.S.C. 6103, or any other statutory protection for personal data.

Section 9. Condition on grants to states.

(a) Condition. As a condition of receiving funds under a federal grant program that funds state or local administration of a rights-sensitive domain, a recipient shall not deploy, or integrate into the administration of the program, a covered system as used by the recipient unless the recipient has:

- (1) identified the state or local legal authority that expressly authorizes that use; and
- (2) published a civil authorization dossier for that use containing the elements listed in Section 4.

The recipient shall certify compliance with this subsection to the federal agency administering the program before receiving funds and [annually thereafter] [Alternative: before deploying any new covered system and before any material change in a covered system's purpose, data sources, or model].

Drafting note

Annual certification is simple to administer but can lag a deployment by up to a year. The bracketed alternative ties certification to the events that change risk. The two can also be combined: an annual certification plus a certification before any new deployment or material change.

(b) Definitions. In this section:

(1) the term “recipient” means a State, a political subdivision of a State, or an agency or instrumentality of either, that receives funds under a grant program described in subsection (a), directly or as a subrecipient; and

(2) the term “covered system as used by a recipient” means an AI system, or a model, dataset, or output derived from it, that meets Section 2(d)(1) and whose security-domain validation, accreditation, or use was cited or relied upon in the recipient’s justification for acquiring or deploying it, or in the vendor’s proposal, marketing materials, or other public statements about the system, its model family, or its product line. The presumption of coverage in Section 2(d) applies, and may be rebutted only by a written determination of the recipient’s chief executive officer, or that officer’s designee, published before deployment and supported by the vendor documentation that Section 2(d) requires.

(c) Administration. The head of each federal agency administering a program described in subsection (a) shall state the condition in each notice of funding opportunity and include it in each grant agreement entered into or renewed after the date of enactment. After written notice and an opportunity to cure of not less than 90 days, the agency head may withhold, or recover, funds attributable to the use of a covered system in violation of this section. The condition applies only to awards made after the date of enactment.

Section 10. Severability. If any provision of this Act, or the application of any provision, is held invalid, the remainder of the Act and its application to other persons or circumstances shall not be affected.

E.2: Civil Authorization Dossier Requirement (Executive Branch)

Note

E.2 is written to stand on its own as an executive order or OMB guidance; it does not depend on E.1 being enacted. Section 1 sets out its own definitions and dossier contents, which match E.1’s.

Section 1. Definitions. In this order:

(a) AI system. Any machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments.

(b) Agency. An executive agency, as defined in 5 U.S.C. 105, and any component thereof.

(c) Rights-sensitive civil domain. Any of the following, and case management for any of them: immigration enforcement; immigration adjudication, including asylum, refugee, naturalization, visa, and immigration-benefits determinations; border and port-of-entry screening; policing; criminal justice beyond policing, including pretrial release and detention, sentencing, probation and parole supervision, and corrections; welfare and benefits administration; and tax administration. The term does not include internal administrative functions such as payroll, hiring, or facilities management, or systems used solely for statistical or research purposes that do not produce decisions about identifiable individuals.

(d) Deploy in a rights-sensitive civil domain. To use an AI system to make, or to materially support, a decision about an identifiable individual in a rights-sensitive civil domain. A system materially supports a decision if it produces an output, such as a score,

ranking, flag, classification, or recommendation, that a decision-maker is required or expected to consider, or that determines which individuals are selected for review, investigation, or action.

(e) Civil authorization dossier. A published document that includes, at minimum: (1) system purpose, stated in plain language; (2) the specific legal authority for the use; (3) a data inventory identifying every data source, its collection purpose, and its legal basis; (4) a model description, including architecture, training-data provenance, known limitations, and validation results; (5) human-review standards; (6) contestability procedures, including notice to affected persons, appeal timelines, and the mechanism for challenging data, inference, and decision; (7) audit-access rules; (8) vendor obligations, including data-return, deletion, portability, non-reuse, and cooperation with auditors; (9) data-retention limits, per data class, with a deletion-verification method; (10) portability rights; (11) prohibited uses; and (12) a sunset date under Section 5.

Section 2. Before deploying any AI system in a rights-sensitive civil domain, the agency shall publish a civil authorization dossier. The dossier shall be available before procurement lock-in, operational deployment, or irreversible integration into agency workflows. The dossier documents authorization; it does not substitute for the underlying legal authority.

Section 3. The dossier shall be reviewed by the agency's Chief AI Officer and legal counsel, and submitted to the agency's Inspector General. A summary shall be transmitted to OMB within 30 days of publication.

Section 4. The agency shall publish a prohibited-use list for each system covered by this requirement.

Section 5. Each authorization shall include a sunset date not to exceed three years from publication. Reauthorization requires an updated dossier and an independent review of the system's performance in civil use by an evaluator with no financial interest in the vendor.

Section 6. The agency's Inspector General is requested to review compliance with this order and to audit covered systems as the Inspector General considers appropriate, and is requested to report findings to the agency head and to Congress in the Inspector General's next semiannual report.

Notes on cited authorities

These notes flag where a citation is correct but narrower than it might look, and where a drafting choice needs counsel's attention.

Notes by authority

- **5 U.S.C. 105** defines "Executive agency" as an Executive department, a Government corporation, and an independent establishment, "for the purpose of this title." It is the right definition for covered federal agencies, and §2(b) incorporates it expressly "for purposes of this Act." APA review uses the broader "agency" definition in 5 U.S.C. 551(1) and 701(b)(1). [Text](#)
- **50 U.S.C. 3003** contains the National Security Act definitions. Paragraph (4) lists the elements of the "intelligence community"; §2(c) cites 3003(4). [Text](#)

- **44 U.S.C. 3552(b)(6)** defines “national security system”: systems involving intelligence or cryptologic activities, command and control of military forces, equipment integral to a weapons system, or systems critical to military or intelligence missions, plus systems protected as classified. It excludes routine administrative and business applications. The citation is correct. [Text](#)
- **41 U.S.C. 103** defines “commercial product.” The citation is correct, but the definition also covers products with “modifications of a type customarily available in the commercial marketplace” and “minor modifications” made to meet Federal Government requirements (§103(3)). A carve-out keyed to commercial products would therefore reach some modified systems. For that reason §2(d) does not exclude commercial products: commercial status may support a determination rebutting the presumption of coverage but does not by itself rebut it, and the lineage and reliance tests do the limiting work. Commercial services are defined separately in 41 U.S.C. 104. [Text](#)
- **5 U.S.C. 552a(e)(4)** requires a system-of-records notice to be published in the Federal Register. The citation is correct. A SORN gives notice; it does not grant legal authority. §2(g) treats it correctly, as an additional requirement alongside the authorizing statute, regulation, or published agency-head authorization. [Text](#)
- **5 U.S.C. 552a(o)** requires written computer-matching agreements. The citation is correct, but the requirement applies only to “matching programs” as defined in 552a(a)(8). That definition excludes, among other things, matches made in a criminal or civil law-enforcement investigation of named persons and certain tax-administration matches. It is a partial hook, not a general firewall.
- **31 U.S.C. 1351** is only the Antideficiency Act’s reporting provision. When an officer or employee violates 31 U.S.C. 1341(a) or 1342, the agency head must report to the President and Congress and send a copy to the Comptroller General. It does not prohibit any spending by itself. The operative bar on obligating funds against a funding limitation is **31 U.S.C. 1341(a)(1)(A)**, read with the purpose statute, **31 U.S.C. 1301(a)**. The administrative and criminal penalties are in 31 U.S.C. 1349–1350. §6(a) cites 1341(a)(1)(A) as the prohibition and 1351 as the reporting duty. §7’s appropriations limitation is what makes the amount available for an unauthorized covered system zero, so the two sections should stay together. [1351](#) · [1341](#)
- **5 U.S.C. 702–706** provide APA judicial review. The citation is correct. Section 702 waives sovereign immunity for relief other than money damages, which is what lets APA review do the work the earlier private right of action could not. Section 706 sets the scope of review. [702](#)
- **28 U.S.C. 2412** is the Equal Access to Justice Act provision on costs and fees. The citation is correct. Fee awards under 2412(d) are limited by party net-worth and size thresholds and are not available if the government’s position was “substantially justified.” [Text](#)
- **26 U.S.C. 6103** protects the confidentiality of returns and return information. The citation is correct; the D.C. Circuit applied it to the IRS–ICE address exchange (Section 06). [Text](#)
- **EU AI Act Article 3(1)** reads: “‘AI system’ means a machine-based system that is designed to operate with varying levels of autonomy and that may exhibit adaptiveness

after deployment, and that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments.” E.1 §2(a) keeps the inference element word for word and leaves out the autonomy and adaptiveness language, so it is somewhat broader than Article 3(1). That is a defensible drafting choice for a gate that turns on lineage and reliance, but it is not an exact match. The existing U.S. statutory definition in 15 U.S.C. 9401(3) is worded differently, and counsel should decide which one §2(a) is meant to be “consistent with.” [Text](#)

- **EU AI Act Article 2(3)** excludes AI systems “where and in so far” as they are placed on the market, put into service, or “used with or without modification exclusively for military, defence or national security purposes.” The exclusion covers defence and national security as well as military use. Recital 24 states that a system built for those purposes that is then used “temporarily or permanently” for civilian, law-enforcement, or public-security purposes “would fall within the scope” of the Regulation. Finding (2)’s shorthand (“military exclusion”) should be read with that wider scope. [Art. 2](#) · [Recital 24](#)
- **Executive Order 14365** (December 11, 2025). Section 3 directs the Attorney General to set up an AI Litigation Task Force within 30 days to challenge state AI laws. Section 8(b) (iii) provides that the legislative recommendation for a preemptive federal framework “shall not propose preempting otherwise lawful State AI laws relating to ... State government procurement and use of AI.” That carve-out limits the legislative proposal only. It does not on its face limit the Task Force’s litigation or the grant conditions in Section 5, and any text relying on it should say so. [Text](#)

Points for counsel

- **§2(d) presumption.** The presumption turns on whether the vendor supplies the same product line or model family to a national-security setting. Neither term is defined. One option is to define them functionally, for example as systems that share a base model, training pipeline, or core software platform, whatever the product name, so that rebranding does not defeat the presumption; another is to leave the terms to OMB guidance. Counsel should also decide whether the rebuttal standard (a published determination supported by vendor documentation) should require independent review.
- **Administrability and vendor due process.** The presumption burdens the agency’s acquisition decision: it determines what the agency must do before buying or deploying a system. It imposes no sanction, penalty, or debarment on the vendor, which remains free to sell to other customers and to supply the documentation that rebuts the presumption. That structure should not raise a due-process problem for vendors, but counsel should confirm it and should confirm that agencies can administer the product-line test.
- **Vendor non-cooperation.** If a vendor will not document a system’s training data, fine-tuning, and accreditation history, the presumption cannot be rebutted and the system stays covered. That is deliberate: lineage is usually known only to the vendor, so the burden of producing it sits with the vendor. Counsel may wish to back this with a contract clause requiring the documentation in any new award.

- **§2(d) reliance gate is intentional.** E.1 covers a system only if it has security-domain lineage and its security validation was cited or relied on, or is presumed to have been under the product-line test. A security-derived system acquired without reliance on its security record, whose presumption is properly rebutted, falls outside E.1 but remains subject to OMB's rules for high-impact AI and, where adopted, E.2. Counsel should confirm that scope, and consider whether reliance should instead trigger added review rather than limit coverage.
- **§6(c) complaints.** "Any person" is intentionally broad for filing a complaint, which costs the agency a reasoned written decision. Judicial review is narrower: a complainant still needs standing to sue under Article III and the APA, which the subsection does not and cannot expand.
- **§6(c) complaint reviewability.** The subsection requires a reasoned decision, not a mere response, and treats silence or a non-answer as a denial. Counsel should confirm which final agency action a court reviews after a deemed denial and on what record, and how the requirement fits with 5 U.S.C. 555(e) (notice of the grounds for denying a request) and 706(1) (agency action unreasonably delayed).
- **§6(d) waiver.** The waiver now reaches only Sections 3, 4, 6(a), and 7, the provisions that would otherwise stop the system, and preserves the firewalls, reporting, complaints, Inspector General notice, and the savings clause. Counsel should confirm the list, and that waiving §6(a) and §7 together is enough for an emergency deployment to proceed lawfully.
- **§4 reauthorization.** Counsel should decide whether the independent review at reauthorization should be done by the Inspector General or by an outside evaluator, as drafted.
- **§2(f) border screening.** Some border and port-of-entry screening systems may themselves be national security systems under 44 U.S.C. 3552(b)(6). The Act reaches them when a covered system is used in that domain for civil decisions about individuals; counsel should confirm that the interaction with §2(c) is intended.
- **§2(g) agency-head authorization.** An agency may not authorize itself where a statute restricts the data or the use. Counsel should confirm that Federal Register publication is the right vehicle and whether notice and comment should be required.
- **§6(e) transition.** The transition overrides §6(a) as well as §§3 and 7, because §6(a) would otherwise treat an option exercise during the window as a prohibited obligation. "Material modification" is not defined; counsel may wish to define it by contrast with the "minor modifications" in 41 U.S.C. 103(3).
- **§9 grant condition.** Conditions on federal grants to states must be stated unambiguously ([Pennhurst State School and Hospital v. Halderman, 451 U.S. 1, 17 \(1981\)](#)). §9 therefore states the condition and the recipient's obligations in the section itself, defines "recipient" and "covered system as used by a recipient" on their own terms rather than through §2(b), which reaches only federal agencies, and applies only to awards made after enactment.

- **E.2 §6.** An executive order cannot direct an Inspector General's work without raising independence concerns under the Inspector General Act, so §6 requests rather than directs review.
- **E.2 §1.** E.2 restates its definitions and dossier contents so that it works without E.1. If E.1 is enacted, the two sets of definitions should be kept identical.

Appendix F

State Version: Model Bill, Introduction, and Talking Points

Note

Drafting aids, not legal advice. Bracketed text marks choices that differ by state: the names of agencies and officers, the state administrative procedure act, whether political subdivisions are covered, and whether civil penalties are included. Counsel in the relevant state should verify every cite and choice before introduction.

F.1: Model State Bill, [State] Civil AI Authorization Act

Section 1. Short title. This Act may be cited as the [State] Civil AI Authorization Act.

Section 2. Findings. The legislature finds that:

- (1) state agencies make decisions about benefits, liberty, and livelihoods in which automated error can cause serious and widespread harm, as the record of automated unemployment-insurance fraud determinations shows;
- (2) that harm has occurred when automated decisions were deployed before legal authority, public disclosure, and a meaningful right to contest were in place, whatever the system's origin;
- (3) AI systems developed, trained, or validated in military, intelligence, and other national-security settings are offered for civil administration, and validation in those settings does not establish that a system is lawful, accurate, or contestable in civil use;
- (4) federal AI governance memoranda apply to federal agencies, not to state agencies; and
- (5) a state agency should identify the legal authority for, and publish the data sources, review standards, and appeal procedures of, any AI system before using it to make or support decisions about individuals.

Section 3. Definitions. In this Act:

(a) AI system. Any machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments. The term also includes any automated decision system, including a deterministic system that applies rules, thresholds, formulas, or matching criteria set by people, when it is used to make or materially support an adverse decision. [The term does not include a tool that only stores, retrieves, displays, or transmits data, or that performs a calculation a human decision-maker directs and checks, such as a spreadsheet or a database query.]

Drafting note

The second sentence of subsection (a) makes clear that a rules-based system like Michigan's MiDAS, which applied fixed criteria rather than a trained model, is covered when it makes or materially supports an adverse decision. The inference-based first sentence, drawn from the EU AI Act, might not reach such a system on its own. The bracketed exclusion keeps ordinary office

software out of scope. Counsel should test both sentences against the state's existing automated systems.

(b) State agency. Any department, board, commission, authority, or other agency of the executive branch of [State] [, and any political subdivision of [State] and any agency of a political subdivision]. [The term does not include the legislature or the judicial branch.]

(c) Rights-sensitive domain. Any of the following, and case management for any of them:

- (1) unemployment insurance, Medicaid, the Supplemental Nutrition Assistance Program, Temporary Assistance for Needy Families, housing assistance, state student aid, and other public benefits;
- (2) child welfare and child protective services;
- (3) policing, including facial recognition, predictive policing, risk scoring, and case prioritization;
- (4) criminal justice beyond policing, including pretrial release and detention, sentencing, probation and parole supervision, and corrections;
- (5) participation in immigration enforcement, including any agreement under 8 U.S.C. 1357(g), and the sharing of data for immigration-enforcement purposes;
- (6) state tax administration, including audit selection and collection enforcement; and
- (7) [other domains the state designates, such as professional or occupational licensing].

(d) Covered system. An AI system that a state agency uses to make, or to materially support, a decision about an identifiable individual in a rights-sensitive domain. The term does not include a system used solely for internal administrative functions such as payroll, hiring, facilities management, or general correspondence, or a system used solely for statistical or research purposes that does not produce decisions about identifiable individuals. That a system is a commercial product does not exclude it.

(e) Materially support. To produce an output, such as a score, ranking, flag, classification, or recommendation, that a decision-maker is required or expected to consider, or that determines which individuals are selected for review, investigation, or action.

(f) National-security setting. An element of the intelligence community (50 U.S.C. 3003(4)), the United States Department of Defense, a national security system (44 U.S.C. 3552(b)(6)), or an equivalent military, intelligence, or defense body of a foreign government.

(g) Security-domain lineage. A covered system has security-domain lineage if:

- (1) the system, or a model, dataset, or output from which it is derived, was developed, trained, fine-tuned, or accredited using data, infrastructure, or test results from a national-security setting; or
- (2) its validation, accreditation, or use in a national-security setting was cited or relied upon in the state agency's justification for acquiring or deploying it, or in the vendor's proposal, marketing materials, or other public statements about the system, its model family, or its product line.

Presumption. A covered system is presumed to have security-domain lineage if the vendor supplies a system of the same product line or model family to a national-security setting. The head of the state agency may rebut the presumption only by a written

determination, published in the dossier before deployment, that states the factual basis for finding that neither paragraph (1) nor paragraph (2) applies, supported by the vendor's documentation of the system's training data, fine-tuning, and accreditation history. If the vendor does not provide that documentation, the presumption may not be rebutted.

Basis stated in the dossier. The dossier shall state whether security-domain lineage is established under paragraph (1), established under paragraph (2), or presumed under this subsection and not rebutted, and the facts supporting that finding. A system with presumed lineage is a heightened-tier system, but the dossier shall not describe presumed lineage as established. The agency may rebut the presumption later, under this subsection, by publishing a revised dossier.

Drafting note

The presumption burdens the state agency's acquisition decision; it imposes no sanction on the vendor, which remains free to sell elsewhere and to supply rebuttal documentation. That should not raise a due-process problem for vendors, but counsel should confirm it, and confirm that agencies can administer the product-line test, under the state's constitution and procurement law.

(h) Separate legal authorization. A statute or regulation of [State] that expressly authorizes use of the covered system for the purpose stated in the dossier, or an express written authorization for that use issued by the head of the state agency and published in [the State Register], together with any data-sharing agreement required by law. An authorization by the agency head does not by itself constitute separate legal authorization where a state or federal statute restricts the collection, use, or disclosure of the data the system relies on, or restricts the use to which the system is put.

(i) Dossier. A civil authorization dossier published under Section 5: a short-form dossier for a Tier 1 system or a full dossier for a Tier 2 system.

(j) Heightened-tier system. A covered system that has, or is presumed to have, security-domain lineage.

(k) Department. The [Department of Information Technology], or another state agency the governor designates to maintain the inventory under Section 9.

(l) Adverse decision. A decision about an identifiable individual that:

- (1) denies, reduces, suspends, terminates, delays, or seeks recovery of a public benefit, service, payment, or license;
- (2) imposes or proposes a penalty, fine, tax assessment, overpayment determination, or collection action;
- (3) selects the individual for investigation, audit, fraud review, surveillance, or enforcement action;
- (4) initiates, or is considered in, a child-welfare investigation or a decision on removal or placement;
- (5) affects the individual's arrest, detention, pretrial release, sentence, or conditions of probation, parole, or other supervision;
- (6) discloses or makes available the individual's data for law-enforcement or immigration-enforcement purposes;

(7) grants a benefit, service, payment, license, or approval subject to conditions, monitoring, reporting, or restrictions that would not otherwise apply to the individual; or

(8) otherwise has a comparably significant adverse effect on the individual's rights, benefits, liberty, or livelihood.

(m) Tier 2 system. A covered system that makes or materially supports an adverse decision, and any heightened-tier system regardless of how it is used. A system whose outputs can lead only to a favorable decision, such as faster approval, is treated as not making or materially supporting an adverse decision if, and only for as long as, all of the following conditions are met, and the tier determination under Section 5(a) explains how each one is met:

(1) every individual the system does not select is processed as though the system had not been used, with no material delay, no missed deadline, no added scrutiny, and no reduction in the staff, resources, or review steps assigned to that individual's case;

(2) neither the system's output nor the fact that an individual was not selected is shown to or considered by a decision-maker in that processing;

(3) the system does not flag, score, rank, or queue any individual for investigation, audit, denial, or secondary review, meaning review beyond the processing the individual would receive if the system had not been used, and non-selection is not used as a flag, score, factor, or criterion in any other decision about the individual, including selection for review, investigation, or enforcement; and

(4) the system does not attach, and its output is not used to attach, conditions, monitoring, reporting, or restrictions to any approval, which would be an adverse decision under Section 3(l)(7).

Requiring a human to authorize the final outcome does not by itself satisfy these conditions. The agency shall monitor processing times for individuals the system does not select and report them annually under Section 9(c). If any of these conditions stops being met, the system is a Tier 2 system and Section 5(f) applies.

(n) Tier 1 system. A covered system that is not a Tier 2 system.

Section 4. Authorization required.

(a) Authorization and dossier. No state agency may deploy a covered system, or integrate it into the administration of a program, unless the agency has obtained separate legal authorization for that use, has published the dossier required by Section 5 for the system's tier at least [60] days before deployment, and, for a Tier 2 system, has met Section 6.

(b) Contracts. No state agency may enter into a new contract, or materially modify an existing contract, for a covered system unless the contract requires the vendor to provide the information the agency needs to complete the dossier, to disclose any security-domain lineage, and to cooperate with audits under Section 9. The contract shall also provide that, if the vendor knowingly supplies dossier information, a lineage disclosure, or an attestation under Section 5(f)(5) that is false or materially incomplete, the agency may terminate the contract, recover the fees paid for any period during which the information was false or materially incomplete, and refer the vendor for suspension or debarment under [state procurement law].

(c) Material changes. A material change in a covered system's purpose, use, data sources, or model requires a revised dossier before the change takes effect. A change described in Section 5(d) is a material change, and Section 5(f) governs its timing.

Section 5. Dossier tiers and contents.

(a) Tier determination. Before publishing a dossier, the state agency shall determine whether the covered system is a Tier 1 system or a Tier 2 system, and shall publish that determination in the dossier with a short statement of reasons that describes how the system's outputs are used in decisions about individuals. An agency that cannot make that determination shall treat the system as a Tier 2 system.

(b) Tier 1: short-form dossier. The dossier for a Tier 1 system shall include, at minimum:

- (1) system purpose, stated in plain language;
- (2) the specific legal authority for the use;
- (3) the categories of data the system uses, and their sources;
- (4) the vendor, if any, and the name and version of the product;
- (5) a statement of whether the system has, or is presumed to have, security-domain lineage, stating whether that lineage is established or presumed, the basis for the finding, and any rebuttal determination under Section 3(g);
- (6) a contact within the agency for questions about the system, and the procedure by which an individual may contest a decision in which the system was used;
- (7) prohibited uses, including any use that would make the system a Tier 2 system; and
- (8) a sunset date, not to exceed three years from publication, after which the authorization expires unless renewed with a new dossier.

(c) Tier 2: full dossier. The dossier for a Tier 2 system, including every heightened-tier system, shall include, at minimum:

- (1) system purpose, stated in plain language;
- (2) the specific legal authority for the use;
- (3) a data inventory identifying every data source, its collection purpose, and its legal basis;
- (4) a model description, including architecture, training-data provenance, known limitations, and validation results;
- (5) human-review standards, including when a human must review, what the reviewer must document, and what the reviewer may override;
- (6) contestability procedures, including notice to affected persons, appeal timelines, and the mechanism for challenging data, inference, and decision;
- (7) audit-access rules, identifying who may inspect the system, at what layer, and under what authority;
- (8) vendor obligations, including data-return, deletion, portability, non-reuse, and cooperation with auditors;
- (9) data-retention limits, per data class, with a deletion-verification method;

(10) portability rights, describing how data and models transition if the vendor relationship ends;

(11) prohibited uses;

(12) a statement of whether the system has, or is presumed to have, security-domain lineage, stating whether that lineage is established or presumed, the basis for the finding, and any rebuttal determination under Section 3(g);

(13) the acceptance criteria and evaluation summary under Section 6(a) and (b), and the [State Auditor]'s findings under Section 6(c) when issued; and

(14) a sunset date, not to exceed three years from publication, after which the authorization expires unless renewed with a new dossier.

(d) Escalation. If a change in the use of a Tier 1 system, including a change in how decision-makers use its outputs, would cause it to make or materially support an adverse decision, the change is a material change under Section 4(c). If the agency finds that a Tier 1 system has, or is presumed to have, security-domain lineage, the system is a Tier 2 system. Subsection (f) governs the timing in both cases.

(e) Reclassification. The [State Auditor], on its own initiative or at the request of the Department, a member of the legislature, or any person, may determine that a system the state agency classified as Tier 1 is a Tier 2 system. The [State Auditor] shall give the agency written notice of the determination and its reasons, and shall publish it. Subsection (f) governs what the agency must do after receiving the notice.

(f) Changes in tier: timing. This subsection governs every case in which a system becomes a Tier 2 system after its tier determination: a planned change in use under subsection (d); the agency's discovery that a Tier 1 system is already being used to make or materially support an adverse decision; a reclassification under subsection (e); and a finding that the system has, or is presumed to have, security-domain lineage.

(1) Bar on adverse-decision use. From the trigger date, the agency may not use the system to make or materially support an adverse decision until it has published a full dossier under subsection (c) and met Section 6. A planned change in use may not begin until the full dossier has been published for at least [60] days, as Section 4(a) requires for a new deployment.

(2) Deadline for a system in use. Except for a planned change in use, and subject to paragraphs (5) through (7), if the agency has not published a full dossier and made a complete submission to the [State Auditor] under Section 6(c) within [60] days after the trigger date, it shall stop all use of the system until it has met Section 6.

(3) Trigger date. The trigger date is the date the agency learns of the unplanned use, receives the [State Auditor]'s notice, or makes or receives the lineage finding, whichever applies.

(4) Competing triggers. If more than one trigger applies to the same system, the earliest trigger date governs the deadline in paragraph (2), a pending planned change does not extend it, and the bar in paragraph (1) continues until every requirement has been met.

(5) Provisional continuation for non-adverse functions. This paragraph applies only when a Tier 1 system becomes a Tier 2 system because of a finding that it has, or is presumed to have, security-domain lineage, or because of a reclassification under subsection (e) that rests only on such a finding. During the [60]-day period in

paragraph (2), and after it until the agency has met Section 6, the agency may continue to use the system for non-adverse functions only if, within [5] business days after the trigger date, it publishes in the dossier, and files with the Department and the [State Auditor], a written continuation certification signed by the head of the agency and accompanied by the vendor's written attestation, stating that:

- (A) the system is used only for functions that are not adverse decisions under Section 3(l) and that meet every condition in Section 3(m), or for internal administrative functions that do not produce decisions about identifiable individuals;
- (B) every module, model, or capability of the product that could make or materially support an adverse decision is disabled at the license or configuration level, and not merely hidden in the user interface;
- (C) any activation of, or configuration change to, such a capability is automatically logged, and the logs are available to the [State Auditor]; and
- (D) re-enabling such a capability is a material change under Section 4(c).

(6) Limits on continuation. If the agency does not publish a timely certification under paragraph (5), it shall stop all use of the system until it has published a full dossier and met Section 6. The [State Auditor] may revoke a certification at any time by written notice, effective immediately or on the date the notice states, and the agency shall then stop all use of the system until it has met Section 6. Continuation is also unavailable, and the agency shall stop all use of the system at once, if:

- (i) the system flags, scores, ranks, or queues individuals for investigation, audit, denial, or secondary review, as described in Section 3(m)(3); a requirement that a human authorize the final adverse outcome does not make that function non-adverse;
- (ii) selecting or accelerating some individuals causes material delay, missed deadlines, or added scrutiny for others;
- (iii) the system attaches conditions, monitoring, or restrictions to an approval; or
- (iv) the certification or the attestation is incomplete or false, or a capability described in paragraph (5)(B) is activated.

(7) Protection for individuals. Stopping a system, or any of its functions, under this subsection does not by itself revoke or reduce a benefit, service, payment, license, or approval already granted. The agency shall return affected cases to the processing it would have used without the system, and may not treat the stop as a ground to deny, delay, or reopen any individual's case.

(g) Publication. The agency shall publish the dossier on its website and file it with the Department. Information whose disclosure is prohibited by law, or that is a trade secret, may be placed in a confidential annex available to the Attorney General and the [State Auditor], provided that the public portion is sufficient for meaningful public and legislative oversight of the system's purpose, authority, data sources, and review procedures.

Drafting note

The two tiers answer the concern that a full dossier for every covered system could overwhelm state agencies. The short-form dossier still states the legal authority, data, vendor, lineage finding, contest procedure, and sunset, so every covered system stays on the public record.

States may add to the short-form elements, but a system that makes or materially supports an adverse decision, and every heightened-tier system, should stay in Tier 2. Subsection (f) sets one timing rule for every route into Tier 2: use in adverse decisions stops at once until the Tier 2 requirements are met, and a system already in use must have a full dossier and a complete auditor submission within [60] days or come out of service. Paragraphs (5) through (7) add one narrow exception. When a system moves into Tier 2 only because of its security lineage, the agency may keep using it for functions that harm no one, such as speeding approvals that would be granted anyway or routing internal paperwork, while it prepares the full dossier. To do so, the agency head must certify within [5] business days, with the vendor's written backing, that every capability that could feed an adverse decision is switched off at the license or configuration level, that any attempt to switch it back on is logged, and that switching it on would require a new dossier. The exception is lost if the system flags or queues people for investigation, slows others down, attaches conditions to approvals, or if the certification proves false; the Auditor can revoke it at any time. A fast-track is not harmless just because it only says yes. In a hypothetical case, a tool that fast-tracks 20 percent of tax refunds, if staff are moved to those cases and the other 80 percent miss their deadlines, fails Section 3(m)(1) and paragraph (6)(ii). If a system is stopped, people keep what they were already granted, and their cases go back to ordinary processing. Counsel should confirm whether state law requires a procedure for an agency to contest the [State Auditor]'s reclassification under subsection (e).

Section 6. Tier 2 safeguards; heightened tier. Before deploying a Tier 2 system, and in addition to the requirements of Sections 4 and 5, the state agency shall meet subsections (a) through (e), and shall meet subsection (f) while the system is in use. A heightened-tier system is always a Tier 2 system and requires a full dossier under Section 5(c), whatever its use; before deploying one, the agency shall also meet subsection (g).

(a) Independent testing. Obtain an evaluation of the system's accuracy, error rates, and differences in error rates across groups in the population on which it will be used, conducted by an evaluator with no financial interest in the vendor, and publish a summary in the dossier.

(b) Acceptance standard. Before testing begins, the evaluator shall set, and the agency shall publish in the dossier, written performance and disparity criteria specific to the system's intended use, including the maximum acceptable rate of erroneous adverse outcomes and the maximum acceptable differences in error rates across groups. The criteria may not be relaxed after testing begins except in a revised dossier that explains the change and is submitted to the [State Auditor]. If the evaluation finds that the system does not meet the criteria, or finds a material error or a disparate impact on a group protected under [state or federal civil-rights law], the agency may not deploy the system until the problems have been corrected and a new evaluation finds the criteria met. The agency may not deploy a system over a failed evaluation.

(c) Auditor review. Submit the dossier, the criteria, and the evaluation to the [State Auditor], who shall review them and, within [90] days after receiving a complete submission, report its findings to the agency and the legislature and publish them. The [State Auditor] shall tell the agency within [15] days if a submission is incomplete; otherwise the submission is complete when received. The agency may not deploy the system until the later of the end of the [60]-day period under Section 4(a) and the date the [State Auditor] reports its findings. If the [State Auditor] has not reported within [90] days,

the agency may proceed after publishing in the dossier a statement that the review was not completed in time; the [State Auditor] may complete its review afterward and may then act under subsection (d).

(d) Auditor action. If the [State Auditor] finds that the dossier is incomplete or materially inaccurate, that the evaluation was not independent or was not adequate to apply the criteria, or that the system has not met the criteria, the [State Auditor] may, by a published written determination stating its reasons, bar deployment until the deficiency is cured or, for a system in use, require the agency to suspend the system's use in adverse decisions or to remediate the deficiency within a period the [State Auditor] sets. [The agency may seek review of the determination under [the State Administrative Procedure Act]; the determination remains in effect pending review unless a court orders otherwise.]

(e) Human decision. Ensure that no adverse decision about an individual rests solely on the system's output, and that a human decision-maker with authority to override the output reviews each adverse decision.

(f) Continuing compliance. If, after deployment, an evaluation, an audit under Section 9(b), or a review under subsection (c) finds that the system no longer meets its criteria, or finds a material error or a disparate impact described in subsection (b), the agency shall within [30] days either suspend the system's use in adverse decisions or adopt and publish a remediation plan. Use in adverse decisions may not continue more than [90] days after the finding unless a new evaluation finds the criteria met.

(g) Heightened tier. Before deploying a heightened-tier system, the agency shall also:

(1) publish in the dossier a translation statement identifying the national-security setting in which the system was developed, trained, validated, or accredited, the population and data on which that validation was done, and how they, and the costs of each kind of error, differ from the civil use;

(2) transmit the dossier to [the appropriate standing committees of the legislature] at least [60] days before deployment; and

(3) set a sunset date not more than [two] years from publication.

Drafting note

Section 6 applies to every Tier 2 system, because a system that feeds adverse decisions needs testing, review, and a human decision whatever its origin; security-derived systems carry the added requirements in subsection (g). A state with limited audit capacity could [limit mandatory pre-deployment review under subsection (c) to heightened-tier systems and make review of other Tier 2 systems discretionary]; the acceptance standard in subsection (b) should apply either way. Whether the [State Auditor] may bar or suspend an executive agency's deployment under subsection (d) depends on the auditor's constitutional and statutory position. Where the auditor is a legislative officer or lacks that authority, counsel may vest the power to bar deployment in [the Attorney General or the Department] and leave the auditor with findings and reclassification.

Section 7. Notice, explanation, and appeal.

(a) Notice. When a covered system makes or materially supports an adverse decision about an individual, the state agency shall give the individual written notice that an AI

system was used, a plain-language explanation of the principal factors in the decision, and instructions for appeal.

(b) Appeal. The individual may challenge the data, the inference, and the decision before a human decision-maker under [the state's existing administrative hearing procedures].

(c) Expedited appeal. If the individual shows that the decision threatens urgent harm, such as the loss of income, food, housing, or medical care needed to meet basic needs, the agency shall decide the appeal within [10] business days.

(d) Interim protection. If an individual who shows urgent harm appeals within the time the notice allows, which may not be less than [10] days, a benefit, service, or payment that the decision reduces, suspends, or terminates shall continue at its prior level until the appeal is decided, and the agency shall take no collection or recovery action based on the decision until then. [Benefits continued under this subsection may be recovered afterward only as existing law allows.]

(e) Other rights. Nothing in this section limits any right under [the State Administrative Procedure Act] or federal law, including any right to continued benefits pending appeal.

Section 8. Domain firewalls. Each state agency shall maintain enforceable technical and legal separation between law-enforcement and immigration-enforcement systems and systems used for benefits, child welfare, or tax administration, including data-minimization, purpose-limitation, and audit-access controls. Data collected for benefits, child welfare, or tax administration shall not be used in a covered system for law-enforcement or immigration-enforcement purposes except through coordination that is (1) authorized by statute, (2) documented in a written agreement, and (3) subject to audit.

Section 9. Inventory, oversight, and reporting.

(a) Inventory. Within [180] days after the effective date, the Department shall publish an inventory of covered systems in use by state agencies, with each system's tier and a link to its dossier, and shall keep it current. For a system in use on the effective date whose dossier is not yet published, the inventory shall carry a transitional entry, clearly marked as transitional, that gives the system's name, vendor, and purpose, the agency's preliminary tier determination, and the date by which the dossier will be published, which may not be later than the end of the transition period under Section 12(a).

(b) Audit. The [State Auditor] may audit any covered system, including the agency's tier determination, and shall report to the legislature on state agencies' compliance with this Act at least once every [two] years.

(c) Annual report. Each state agency shall report annually to the Department and the legislature on the covered systems it uses, any waivers under Section 11, the processing times monitored under Section 3(m), and the number and outcome of complaints and appeals involving covered systems, including complaints deemed denied under Section 10(c).

Section 10. Enforcement.

(a) Attorney General. The Attorney General may bring an action in [court] for declaratory or injunctive relief to prevent the deployment, or stop the continued use, of a covered system in violation of this Act.

(b) Civil penalties [optional]. [Option A: A court may impose on a state agency that knowingly violates this Act a civil penalty of not more than [\$10,000] per violation, to be deposited in [a dedicated fund for independent evaluation under Section 6(a)].] [Option B: A court may impose on a vendor that knowingly provides false or materially incomplete information for a dossier, a lineage disclosure, or an attestation under Section 5(f)(5) a civil penalty of not more than [\$10,000] per violation.]

Drafting note

The \$10,000 figures are placeholders; set amounts to match the state's existing civil-penalty schedules. A civil penalty against a state agency is paid from state funds and, unless directed elsewhere, returns to the state treasury, so it can be circular and may simply reduce the budget of the program the Act protects. States may omit subsection (b), direct agency penalties to a dedicated fund (Option A), or apply penalties only to vendors that knowingly supply false or materially incomplete information (Option B). Option B reaches the vendor's continuation attestation under Section 5(f)(5) because the vendor controls whether a capability is actually disabled; the "knowingly" standard keeps honest mistakes outside the penalty. The contract terms required by Section 4(b) give the agency separate remedies: termination, recovery of fees, and referral for suspension or debarment. Whether the Attorney General may sue another state agency, rather than represent it, varies by state constitution and statute; where that authority is lacking, consider vesting enforcement in [the State Auditor] or relying on subsection (c).

(c) Complaints and judicial review. Any person, including an individual affected by a covered system or an organization representing such individuals, may file a written complaint with the state agency and the Department alleging use of a covered system in violation of this Act, including use without separate legal authorization, without a dossier or with a dossier that omits a required element or is materially inaccurate, in violation of Section 6, 7, or 8, or after the authorization has expired. The complaint shall identify the system, or describe it as specifically as the complainant can, state the facts supporting the allegation, and give a means of contacting the complainant. Within [90] days, the agency shall issue a written decision that states its findings on each allegation, the corrective action it will take or its reasons for denying the complaint, and the records on which it relied, which it shall make available to the complainant except as Section 5(g) allows. An acknowledgment of receipt, a statement that the complaint is under review, or any other response that does not meet this subsection is not a decision. If the agency does not issue a decision meeting this subsection within [90] days, the complaint is deemed denied, and the denial is final agency action. A decision on a complaint, a deemed denial, and final agency action to deploy or continue a covered system in violation of this Act are subject to judicial review under [the State Administrative Procedure Act]. [Relief under this subsection is limited to declaratory and injunctive relief.]

Section 11. Governor's waiver.

(a) Waiver. The governor may waive Sections 4 through 6 for a specific covered system for not more than [90] days upon a written determination that the waiver is necessary to avert an imminent threat to human life.

(b) One renewal. A waiver may be renewed once, for not more than [90] days, upon a new written determination. No further waiver may be issued for the same system and purpose.

(c) Notice to the legislature. Within [7] days of issuing or renewing a waiver, the governor shall transmit the determination to [the presiding officer of each house of the legislature and the appropriate standing committees] and publish it, withholding only information whose disclosure is prohibited by law.

(d) Limits. A waiver does not suspend Sections 7 through 10 or Section 13, and does not authorize any use otherwise prohibited by law.

Section 12. Transition.

(a) Systems in use. Notwithstanding Section 4, a covered system in use on the effective date may remain in use for [12] months after that date, during which the state agency shall obtain separate legal authorization, publish the dossier required by Section 5, list the system in the inventory under Section 9(a), and, for a Tier 2 system, meet Section 6.

(b) Existing contracts. Renewals and option exercises under contracts in effect on the effective date that occur within that period fall within this transition. New contracts, and material modifications of existing contracts, made after the effective date are covered immediately.

(c) Withdrawal. A covered system that does not comply with this Act at the end of the transition period shall be withdrawn from use.

Section 13. Savings; relation to federal law.

(a) No new authority. Nothing in this Act authorizes any use of an AI system that is otherwise prohibited by law.

(b) Privacy protections. Nothing in this Act supersedes, limits, or diminishes [the state's privacy and data-protection laws], the Privacy Act, 5 U.S.C. 552a, the tax confidentiality protections of 26 U.S.C. 6103, or any other protection for personal data.

(c) Federal requirements. Where federal law or a condition of federal funding imposes requirements on the same system, the state agency shall comply with both. A federal dossier may be incorporated by reference but does not satisfy this Act unless the state dossier identifies the state legal authority for the state's use. A state dossier does not satisfy a federal requirement unless the federal agency administering that requirement accepts it.

(d) Scope. This Act governs state agencies and their contractors. It does not regulate the development or use of AI by private parties acting on their own behalf.

Section 14. Severability. If any provision of this Act, or its application to any person or circumstance, is held invalid, the invalidity does not affect other provisions or applications of the Act that can be given effect without the invalid provision or application.

Section 15. Effective date. This Act takes effect [date]. Section 4(b) applies to contracts entered into or materially modified on or after the effective date.

Points for counsel on F.1

- **§3(a) rules engines.** The definition now reaches deterministic automated decision systems used in adverse decisions. Confirm that it covers MiDAS-type systems without sweeping in ordinary office software, and consider listing examples.

- **§3(m) favorable-only test.** Confirm that the three conditions are administrable, and decide how processing times for people not selected should be measured and against what baseline.
- **§3(l)(7) conditional approvals.** Confirm that “conditions... that would not otherwise apply to the individual” leaves out standard program conditions that apply to everyone, while reaching conditions a system adds.
- **§5(f)(5)–(7) continuation certification.** Decide whether a certification by the agency head, backed by the vendor’s attestation and subject to revocation by the [State Auditor], is enough, or whether continuation should also need independent confirmation or Auditor approval before it starts.
- **§5(f)(5)(B) verifying that capabilities are disabled.** Decide how “disabled at the license or configuration level” should be shown and checked, for example by license records, configuration exports, technical documentation the Auditor can inspect, testing that the capability cannot run, or review of the activation logs required by §5(f)(5)(C).
- **§6(b) acceptance standard.** Decide who sets the criteria (the evaluator, as drafted; the agency; or the Department by rule) and how “material error” and “disparate impact” map onto the state’s civil-rights law.
- **§6(c)–(d) auditor block authority.** Confirm whether the [State Auditor] may bar or suspend an executive agency’s deployment under the state constitution, whether the [90]-day review and the publish-and-proceed default fit the auditor’s capacity, and how an agency may contest an auditor determination or a reclassification under §5(e).
- **§7(c)–(d) interim protection.** Confirm how the expedited appeal and continued benefits interact with federal rules on benefits pending appeal in federally funded programs and with the state’s hearing system.
- **§10(c) complaint reviewability.** Confirm which action a court reviews after a deemed denial, what record it reviews, and whether the state’s standing rules let organizations sue.
- **Carried over.** The product-line definition and vendor due process (§3(g)); the Attorney General’s authority to sue another state agency (§10(a)); the meaning of “material modification” (§4(b), §15); and penalty amounts and where penalties are paid (§10(b)).

Introduction to the state version

Why a state version is needed

The federal statute in Appendix E governs federal agencies and attaches conditions to federal grants; it cannot directly govern state administration. Yet the documented harm from automated administration has largely occurred in programs that states run. MiDAS was a Michigan system. Robodebt was Australian, but its American analogues (automated unemployment insurance fraud determinations, Medicaid eligibility terminations, SNAP recertification failures) are administered by state agencies under federal grant conditions. State and local police departments deploy facial recognition, predictive policing, and risk scoring without federal oversight.

MiDAS, Robodebt, and SyRI were not built or validated in security settings, and a rule triggered only by security lineage would not have caught them. They show something else: what happens when automated decisions are deployed before anyone establishes the legal authority, publishes how the system works, or provides an effective way to contest it. That is why F.1 requires authorization and a published dossier for any AI system used in decisions about individuals in a rights-sensitive domain, and adds a heightened tier when security-domain lineage is present or presumed. The dossier is scaled to the system's role: a short-form dossier for a system that does not make or materially support an adverse decision, and the full dossier for one that does and for every security-derived system.

How the two versions interact

They are complementary, not duplicative. Four points of interaction matter.

1. Federal systems deployed in state programs. When a federal agency provides a system to a state agency (for example, an eligibility determination tool, a fraud detection platform, or a case management system), the federal statute governs the federal agency's authorization and dossier. The state statute governs the state agency's use. A state agency cannot satisfy its own obligation by pointing to the federal dossier; it must produce its own, because the legal authority for state use is different from the legal authority for federal use (F.1 §13(c)).

2. State agencies receiving federal funds. Federal §9 conditions certain grants on the recipient identifying its own legal authority and publishing a dossier for covered systems. The state statute gives the state its own mechanism for meeting that condition. Without a state statute, the state has no mechanism of its own for doing so, and enforcement of the condition is left to the federal grantor.

3. Preemption. Executive Order 14365 (December 2025) sought a federal framework for AI and directed the Attorney General to establish an AI Litigation Task Force to challenge state AI laws. The order also provides that the legislative recommendation it calls for "shall not propose preempting otherwise lawful State AI laws relating to ... State government procurement and use of AI." That carve-out limits only the legislative proposal; by its terms it does not limit the Task Force's litigation or the grant conditions in Section 5 of the order. A state civil authorization statute is therefore on comparatively strong ground: it regulates the state's own agencies, not private parties (F.1 §13(d)), and it is the kind of law the order's legislative carve-out describes.

4. Federal-state information sharing. Where a state agency shares data with a federal agency, or receives federal data, both statutes may apply. The state statute's domain-firewall provision (F.1 §8) and the federal statute's matching-agreement provision should be read together. Neither displaces the Privacy Act or 26 U.S.C. 6103.

What the state version covers that the federal version cannot

- State-administered benefits: unemployment insurance, Medicaid, SNAP, TANF, child welfare, state student aid
- State and local policing: facial recognition, predictive policing, risk scoring, case prioritization

- State criminal justice beyond policing: pretrial release, sentencing, probation and parole, corrections
- State tax administration: audit selection and collection enforcement
- State and local case management systems that produce decisions about individuals

The core standard, with a broader trigger

Both versions require separate legal authorization and a published civil authorization dossier before deployment, and both require a sunset. They differ in what triggers the requirement. E.1 applies to systems with security-domain lineage whose security validation was relied upon; it is a lineage-plus-reliance gate, not a ban on security-derived AI, and systems outside it remain subject to general rules for rights-impacting AI and to E.2. F.1, like the executive-branch text in E.2, applies to any AI system used to make or materially support decisions about individuals in a rights-sensitive domain. Every F.1 system that feeds adverse decisions must pass independent testing against published acceptance criteria, clear State Auditor review before deployment, and leave each adverse decision to a human who can override the output (F.1 §6(a)–(f)). F.1 then adds a heightened tier (a translation statement comparing the security and civil settings, legislative notice, and a shorter sunset) when security-domain lineage is established or presumed (F.1 §6(g)). F.1 also scales the dossier itself: a covered system that does not make or materially support an adverse decision needs only a short-form dossier, while a system that does, and every heightened-tier system, needs the full dossier (F.1 §5). E.2 has no short form; every system it covers needs the full dossier. The domain-translation test in Section 11 of this brief is the analysis behind the heightened tier; it is not written into either bill as a legal standard.

What the state version does differently

- A broader trigger with a heightened tier, as described above.
- A tiered dossier. Systems that do not make or materially support an adverse decision file a short-form dossier; the state auditor can move a system into the full-dossier tier, and a change in use that makes the system part of adverse decisions requires the full dossier first. One timing rule governs every route into Tier 2, with a narrow, certified exception that lets a system moved into Tier 2 only by its security lineage keep doing non-adverse work while the agency prepares the full dossier (F.1 §5(d)–(f)).
- Testing with consequences. An evaluator sets use-specific accuracy and disparity criteria before testing; a system that fails them cannot be deployed, and later findings of material error or disparate impact require suspension or remediation. The State Auditor must review every Tier 2 system before deployment and can bar or suspend one that fails (F.1 §6).
- Reasoned complaint decisions. An agency must answer a complaint with a written decision giving findings, corrective action or reasons for denial, and the record relied on; silence or a non-answer is a deemed denial that can be taken to court (F.1 §10(c)). E.1 §6(c) now does the same.
- AG enforcement instead of DOJ enforcement. Whether a state attorney general may sue another state agency, rather than represent it, varies by state constitution and statute, so

confirm that authority before introduction. Where it exists, the sovereign-immunity and separation-of-powers problems that block federal enforcement do not arise in the same way at the state level.

- Optional civil penalties. Penalties are bracketed in F.1 §10(b). A penalty paid by a state agency comes from state funds and can be circular, so the bill offers a dedicated-fund option and a vendor-only option. The federal version relies on the funds limitation and APA review.
- Governor's waiver instead of agency-head waiver. At the state level, the governor is the appropriate waiver authority. The waiver is limited to imminent threats to life, may be renewed once, must be reported to the legislature, and cannot suspend notice and appeal, firewalls, reporting, enforcement, or the savings clause (F.1 §11).
- No private right of action. Standing and exhaustion rules vary widely across state APAs. Complaints, judicial review under the state APA, and AG enforcement are more reliable mechanisms. States with robust APAs may consider adding a private right of action.

How to introduce it

The state version is designed to be introduced as a standalone bill. It does not depend on the federal statute passing. If the federal statute passes later, the state statute will operate alongside it; if the federal statute does not pass, the state statute still provides the authorization gate for state agencies.

A worked example (hypothetical)

Note

This example is invented to show how F.1 would operate. It describes no real agency, vendor, or system.

A state unemployment agency wants to license a commercial tool that scores weekly claims for possible fraud and sends the highest-scoring claims to investigators. The vendor sells nothing to national-security customers, so no lineage presumption arises.

Tier. Selecting a claimant for fraud review is an adverse decision (§3(l)(3)), so the tool is a Tier 2 system even though a person makes the final call (§3(e), §3(m)). The agency says so in its tier determination (§5(a)).

Who prepares the dossier. The program office prepares the full dossier (§5(c)) with agency counsel, using documentation the vendor must supply under its contract (§4(b)). The agency publishes it and files it with the Department at least [60] days before use (§4(a)).

Who checks it. Before testing begins, an evaluator with no financial interest in the vendor sets written limits on error rates and on differences across groups, then tests the tool on the state's own claims (§6(a)–(b)). The [State Auditor] reviews the dossier and the evaluation within [90] days. The agency cannot go live until that review is done, or until the [90] days have run and the agency has published that the review was late (§6(c)). If the tool fails its criteria, it cannot be deployed until it is fixed and retested, and the Auditor can bar deployment (§6(b), (d)).

If the vendor refuses documentation. Without the training-data provenance and validation results, the agency cannot complete the model description in the full dossier (§5(c)(4)), so it cannot deploy the tool (§4(a)). A new contract signed without the documentation clause would itself violate §4(b).

What the claimant gets. A claimant flagged and later told she owes an overpayment receives written notice that an AI system was used, the principal factors, and how to appeal (§7(a)). A person with authority to override the score must have reviewed the decision (§6(e)). She can challenge the data, the score, and the decision (§7(b)). If repayment would cost her the rent, she can ask for an expedited appeal, and collection pauses until it is decided (§7(c)–(d)).

If a complaint goes unanswered. A legal-aid group complains that the agency switched to a new model version without a revised dossier (§4(c)). Within [90] days the agency must issue a written decision with findings and either a fix or its reasons for refusing. A letter saying the complaint is “under review” does not count. If no decision arrives, the complaint is deemed denied, and the group can seek judicial review under the state APA, subject to the state's standing rules (§10(c)). The Auditor and the Attorney General can act separately (§6(d), §10(a)).

Legislator talking points

The 30-second version

“State agencies already use automated systems to decide who gets unemployment benefits, who keeps Medicaid, and who gets flagged for fraud. When Michigan ran one with little human review and notice and appeal procedures that failed in practice, about 40,000 people were automatically accused of fraud. Meanwhile, AI built and tested for war, intelligence, and border enforcement is being offered to civil government, and the federal pathways that could carry it there already exist. Battlefield validation is not civil authorization. This bill requires state agencies to publish the legal authority, data, and appeal process for any AI system before it is used in decisions about people’s benefits, liberty, or livelihoods, with a harder look at systems built for security. It doesn’t ban AI. It requires authorization.”

Why now

- Automated administrative harm has a documented record: Michigan’s MiDAS system automatically accused about 40,000 people of unemployment fraud; 93% of the roughly 22,000 determinations later reviewed did not involve fraud. The state paid \$20 million. MiDAS was not a security system. The failure was deployment with little human review, little transparency, and notice and appeal procedures that were defective in practice. This bill requires those safeguards first.
- The pathways from security work into domestic administration already exist at the federal level: the same vendor platforms serve defence, immigration enforcement, and tax and health agencies (Section 06). This brief located no documented case of a battlefield-validated model entering a state agency; the bill puts the gate in place before one does.

- Federal AI governance frameworks (OMB M-25-21 and M-25-22) apply only to federal agencies, not state agencies, and they carve out national-security systems.
- Executive Order 14365 says the federal preemption legislation it calls for should not propose preempting otherwise lawful state laws on state government procurement and use of AI. That carve-out does not limit the DOJ AI Litigation Task Force or the order's grant conditions, but it indicates that the federal legislative proposal will not target states' rules for their own agencies. This is the moment for states to act.

What the bill does

Four things.

- 1. Creates an authorization gate.** No state agency may use an AI system to make or materially support decisions about people in a rights-sensitive domain without separate legal authorization.
- 2. Requires a public dossier, scaled to the stakes.** Before deployment, the agency must publish the system's purpose, legal authority, data, vendor, appeal path, and sunset date. A system that makes or materially supports an adverse decision (a denial, cut-off, penalty, or selection for investigation) and every security-built system must also publish the full dossier: model description, human-review standards, audit access, retention limits, and the rest.
- 3. Requires testing that counts.** Every system that makes or materially supports an adverse decision must pass independent testing against published accuracy and fairness criteria (a failed test blocks deployment), go through State Auditor review before it goes live, and leave every adverse decision to a human who can override it. Systems with military, intelligence, or national-security lineage (presumed when the vendor sells the same product line to those customers) also need a published comparison of the security and civil settings, legislative notice, and a shorter sunset.
- 4. Provides rights and enforcement.** People get notice, a human appeal, and an expedited appeal with benefits continued where urgent harm is shown. The Attorney General can seek injunctive relief; the state auditor reviews agency compliance; any person can file a complaint, must get a reasoned written decision within 90 days, and can seek judicial review under the state APA. Whether to add civil penalties is left to each state.

What the bill does not do

- It does not ban AI. Agencies can use AI in rights-sensitive domains. They just need authorization first.
- It does not apply to internal administrative functions. Payroll, hiring, facilities management, and general correspondence are excluded.
- It does not apply to statistical or research systems that don't produce decisions about individuals.
- It does not single out commercial products. A commercial product with no security lineage faces only the base requirements of authorization and a dossier for its tier;

commercial status can help an agency rebut the presumption of security lineage but does not do so on its own.

- It does not displace any existing privacy protection. The Privacy Act, state privacy laws, and federal tax confidentiality protections remain in force.
- It does not prevent legitimate interagency coordination. Coordination authorized by statute, documented in writing, and subject to audit is expressly permitted.

The cost question

- Most of the dossier content is already produced: privacy impact assessments, system-of-records notices, procurement justifications, authorizations to operate. The bill adds the trigger and the publication, not the underlying analysis.
- The dossier scales with the stakes. A system that does not make or materially support an adverse decision needs only a short-form dossier; the full dossier is reserved for systems that can contribute to a denial, penalty, or investigation, and for security-built systems.
- Agencies that cannot produce these documents for a rights-sensitive deployment are not ready to deploy.
- The alternative (after-the-fact litigation, wrongful determinations, and settlement costs) is more expensive than the gate. Neither MiDAS nor Robodebt was a security system; both show the cost of automated decisions deployed before legal authority, review, and appeal were secured. Michigan's MiDAS system automatically accused about 40,000 people of unemployment fraud and ended in a \$20 million class settlement; Australia's Robodebt scheme has cost more than A\$2.4 billion.

Objections and answers

"This will slow down modernization."

The dossier is assembled largely from documents agencies already produce. The added work is publication, plus independent testing and auditor review for systems that feed adverse decisions. Systems that cannot contribute to an adverse decision file only a short-form dossier, so the full requirement falls on the systems where errors have the highest human cost.

"The legislature can't review every AI system."

The bill doesn't require legislative review. It requires agency authorization and public disclosure, State Auditor review for systems that feed adverse decisions, and notice to the legislature for security-derived systems. The legislature's role is oversight, not approval.

"This is anti-innovation."

The bill is pro-authorization. It creates a clear path for agencies to adopt AI where the legal basis exists. What it prevents is adoption without authorization.

"Federal law already covers this."

The federal AI governance memoranda (OMB M-25-21 and M-25-22) apply only to federal agencies, not state agencies, and they carve out national-security systems. The crossing point is inadequately governed at both levels.

“Vendors will leave the state.”

The bill regulates state agencies, not vendors. Vendors that want to sell to state agencies will comply. Those that can't produce documentation of their systems' provenance are the ones the bill is aimed at.

“The AG enforcement provision is unusual.”

The bill pairs attorney-general enforcement with state-auditor review and judicial review under the state APA. Whether the attorney general may sue another state agency, rather than represent it, varies by state constitution and statute, so confirm that authority before introduction.

The one-sentence ask

“No AI system should decide people's benefits, liberty, or livelihoods in state government until the state has authorized it in public, and no system built for war, intelligence, or enforcement should get there without a harder look.”

What to say if asked about the federal statute

“The federal bill and this bill do different things. The federal bill governs federal agencies and attaches conditions to federal grants. This bill governs our state agencies. They work together. This bill stands on its own.”

A note on sequencing

If you're introducing both, the state version should go first if the state has a functioning APA and an AG willing to enforce. The federal version is a longer legislative lift. The state version can pass in a single session and produce concrete results before the federal debate concludes.

If you're introducing only one, introduce the state version. It reaches the state programs where automated-administration harm has actually occurred, and its authorization gate applies whether or not a system came from a security setting.

Note

Both of these are drafting aids, not legal advice. The constitutional and statutory cites need verification by counsel in the relevant jurisdiction before introduction.

Revision history

Changelog

This changelog records what changed between the original June 2026 brief and the October 7, 2026 revision, other than the factual corrections listed in Section 16. Version 1.2 and 1.1 changes are summarized in plain language at the front of the brief and listed by provision below.

Version 1.2 (October 7, 2026)

- **F.1.** §3(l) adds new paragraph (7), conditional approvals, as an adverse decision; the catch-all becomes (8). §3(m) now has four conditions: no material delay, missed deadline, or added scrutiny for people not selected; no display or use of the output; no flagging, scoring, ranking, or queuing for investigation, audit, denial, or secondary review; and no conditions attached to approvals. It also states that a human sign-off does not by itself satisfy the conditions. §4(b) requires contract terms for termination, fee recovery, and referral for suspension or debarment when a vendor knowingly supplies false or materially incomplete dossier information, lineage disclosures, or continuation attestations. §5(f)(2) is made subject to new §5(f)(5)–(7): (5) a provisional continuation for non-adverse functions when a system enters Tier 2 only through lineage, on a certification by the agency head and the vendor's attestation published within [5] business days; (6) the consequences of no certification, revocation by the [State Auditor], and four disqualifiers; (7) protection for individuals when a system stops. The §5 drafting note explains continuation and uses a labeled hypothetical tax-refund example. §10(b) Option B covers lineage disclosures and continuation attestations, and its drafting note is updated. New counsel points cover conditional approvals, the sufficiency of self-certification, and how disabling is verified. The list of differences from the federal version mentions the continuation.
- **Presentation.** Appendix text now supports a third indentation level, so sub-items such as (A) and (i) inside a lettered subsection sit below their paragraph; a lettered subsection (i) that follows (h) is no longer mistaken for a roman numeral.

Version 1.1 (October 7, 2026)

- **F.1.** §3(a) adds deterministic automated decision systems used in adverse decisions, with a bracketed exclusion for ordinary office software and a drafting note. §3(g) requires the dossier to say whether lineage is established or presumed. §3(m) adds a three-part test for favorable-only systems and processing-time monitoring. §4(a) requires Tier 2 systems to meet §6 before deployment. §5(b)(5) and (c)(12) follow §3(g); new §5(c)(13) adds the acceptance criteria, evaluation, and auditor findings to the full dossier. §5(d)–(e) now point to new §5(f), a single timing rule for escalation, reclassification, unplanned adverse use, and lineage findings (the former 30-day grace period after reclassification is replaced by an immediate bar on adverse-decision use and a [60]-day deadline); publication moves to §5(g). §6 is restructured: independent testing, a new acceptance standard, mandatory pre-deployment auditor review, auditor action, the human-decision

rule, and continuing compliance apply to every Tier 2 system, and §6(g) keeps the heightened-tier translation statement, legislative notice, and shorter sunset. §7 adds expedited appeals and interim protection. §9(a) adds transitional inventory entries; §9(c) adds processing times and deemed-denied complaints. §10(c) requires a reasoned written decision. §11(d) preserves Sections 7 through 10 and 13 during a waiver. §12(a) refers to the inventory and to Tier 2 systems. New points for counsel follow the bill, and a hypothetical worked example follows the state-version introduction.

- **E.1.** Retitled to describe it as a lineage-plus-reliance gate, with a scope note. §4 requires an updated dossier and independent review at reauthorization. §6(b) reports include waived systems. §6(c) widens complaint scope and requires a reasoned written decision. §6(d) limits waivers to Sections 3, 4, 6(a), and 7, bars further waivers for the same system and purpose, and preserves the firewalls, reporting, complaints, Inspector General notice, and savings clause. New counsel notes cover the reliance gate, complaint reviewability, the waiver list, and reauthorization review.
- **E.2.** New §1 sets out definitions and dossier contents so that E.2 stands alone; the former §§1–5 are now §§2–6, and reauthorization requires an updated dossier and independent review by an evaluator with no financial interest in the vendor.
- **Brief.** New passages in Sections 01 and 03 on why security lineage warrants added scrutiny and what separate authorization means in the bills. Section 09 adds “What the tags verify” and revises the Confirmed and Opaque definitions; the cover and closing note match. Opaque passages now say what was not located rather than asserting that facts are not public. Section 10 presents the equation as an illustration of assumptions and drops the claim that a single audit cannot offset unreviewed cycles. Section 15 and Recommendation 2 match the bills’ mechanism and E.1’s scope. MiDAS is described as having defective notice and appeal procedures, not none. The AI-use note adds ChatGPT, which was used in review.

Evidence tags

- **Fourth tag added.** Claims are now tagged Confirmed, Reported, Inferred, or Opaque. Reported covers on-record claims by an official, agency, or vendor, and claims carried by a single press account, that no one has independently verified. Opaque now means only that material facts are not publicly available (refined in version 1.1 to distinguish facts not located in a search from facts shown to be withheld).
- **Retagged as Reported:** the Brave1 Dataroom figures of 100+ companies and 80+ models; officials’ statements that the Dataroom may later share algorithms with allies; the May 12 statements about Palantir’s role in deep-strike planning (the meeting itself remains Confirmed); the Ukrainian AI-guidance trial results and the tenfold-increase figure; the Maven 13,000-targets and 100,000-user figures; Palantir’s statement that it purged Medicaid data; WIRED’s report of the ELITE tool (the DHS purchase agreement remains Confirmed); TIME’s 2024 report that Palantir’s software was supplied free and used by more than half a dozen Ukrainian agencies; and the financial-press report that Maven’s program-of-record designation took effect (whether it did remains Opaque). The matching timeline entries were retagged.

- **Mixed claims split.** Sentences and table rows that combined a documented fact with an inference, an unverified claim, or an unknown now carry a separate tag for each part, including the national-security carve-outs, the IRS audit pilot, the Brave1 launch and its participation figures, and the May 12 meeting.
- **Transfer claim.** Whether a battlefield-validated model has been transferred into a civil system is no longer tagged. The text and claims table now say plainly that no documented instance was located; the contract terms and lineage that would show a transfer remain Opaque.

Model legislation (Appendices E and F)

- **Origin.** Appendices E and F were adapted from an outside legal review of the June brief and then revised.
- **Fact-check corrections to the review's text.** The talking points said MiDAS cost Michigan \$2.4 billion; MiDAS ended in a \$20 million settlement, and the A\$2.4 billion figure is Australia's Robodebt. The claim that federal AI guidance "expressly excludes" state agencies now says it applies only to federal agencies. Statements that state attorneys general "routinely" enforce state law against, or sue, state agencies were replaced with hedged language. The Executive Order 14365 passages now say that the Section 8(b)(iii) carve-out limits only the legislative recommendation. The claim that without a state statute the federal grant condition is "unenforceable" now says the state would lack its own certification mechanism.
- **E.1 citation corrections.** §2(b) incorporates 5 U.S.C. 105 "for purposes of this Act," because section 105 applies by its terms only "for the purpose of this title"; §2(c) cites 50 U.S.C. 3003(4); §6(a) cites the Antideficiency Act prohibition, 31 U.S.C. 1341(a)(1)(A), with 31 U.S.C. 1351 as the reporting duty.
- **E.1 §2(d) reliance test.** Reliance may now be shown by the vendor's proposal, marketing, or other public statements, not only the agency's acquisition record. A rebuttable presumption of coverage applies when the vendor supplies the same product line or model family to a national-security setting. The separate commercial-product exclusion was removed; commercial status may support rebuttal but does not by itself rebut the presumption. §3's cross-reference to a Section 2(d) "exclusion" now refers to the rebuttal.
- **E.1 §2(f).** Rights-sensitive domains now include immigration adjudication, border and port-of-entry screening, and criminal justice beyond policing.
- **E.1 §2(g).** "Executive authority" is narrowed to an express written authorization by the agency head, published in the Federal Register, which cannot serve alone where a statute restricts the data or use. A definition of "appropriate congressional committees" was added as §2(h).
- **E.1 §6(d) and §6(e).** The waiver is limited to imminent threats to life (national security removed), may be renewed once, and requires the full determination, with any classified annex, to go to the appropriate committees. The transition now begins "Notwithstanding Sections 3, 6(a), and 7"; §6(a) is included so that option exercises within the window are

not treated as prohibited obligations. Renewals and option exercises under existing contracts within the 12-month window fall within the transition; new awards and material modifications are covered immediately.

- **E.1 §9.** Rewritten to state the grant condition in the section itself and to define “recipient” and “covered system as used by a recipient” on their own terms, because §2(b) reaches only federal agencies.
- **E.2 §5.** The Inspector General “is requested to review,” rather than directed.
- **F.1 model state bill added.** Built on E.2’s broader trigger, with a heightened tier for security-domain lineage, optional civil penalties with a drafting note on circularity, and a governor’s waiver limited to imminent threats to life.
- **Appendix F framing.** MiDAS, Robodebt, and SyRI are now presented as showing why authorization comes first, not as cases a security-lineage rule would have caught. The 30-second pitch no longer says that war- or border-built systems are already appearing in state agencies. The talking point that state agencies are “acquiring AI systems at scale” through procurement without public justification was removed because no source was located for it.
- **Recommendations.** Recommendations 1 and 2 now point to the model federal text in Appendix E and the model state bill in Appendix F (F.1).

Second-round review fixes

- **Bottom line.** A one-page summary of five findings and five actions now opens Section 01. It is drawn only from the existing summary and recommendations and carries the same tags.
- **Maven program of record.** The Executive Summary now says the March 9 memo directed that Maven become a program of record by September 30, 2026, matching Section 05. The financial-press report that the designation took effect is tagged Reported, and whether it did is tagged Opaque, in the text, timeline, and claims table.
- **Ukraine trials.** The Ministry’s announcement that the trials took place is tagged Confirmed; the six-of-seven result and the tenfold figure are tagged Reported, in the text, timeline, and claims table.
- **Maven figures.** The Section 02 timeline entry now names James Mazol and Cameron Stanley.
- **Medicaid and Maine cases.** The wording now follows the sources: officials admitted the over-sharing and the judge paused CMS–ICE sharing; the transfer to Palantir rests on a state attorneys general motion citing discovery documents; and in the Maine case the government acknowledged in a filing that the agent’s entry led to a border stop. No court has ruled on either point.
- **Model text.** E.1 §6(c) and F.1 §10(c) now say who may complain and what a complaint must contain. E.1 §9 offers a bracketed event-based alternative to annual certification. E.1 §2(d), E.1 §9, and F.1 §3(g) state that the presumption cannot be rebutted without vendor documentation. F.1 §10(b) uses bracketed \$10,000 placeholders, and F.1 §13(c) adds that a state dossier does not satisfy a federal requirement unless the federal

agency accepts it. New counsel notes cover product-line definitions, vendor non-cooperation, and complaint standing.

- **Presumption due process.** Counsel notes in Appendix E and F.1 now address whether the product-line presumption is administrable and fair to vendors.
- **Text export.** The illustrative error-model figures are now a single labeled table in the PDF and Word versions.

Tiered dossier (October 7, 2026)

- **F.1 tiered dossier.** Responding to the concern that full dossiers for every covered system could overwhelm state agencies, F.1 now scales the dossier to the system's role. New §3(l) defines "adverse decision," and new §3(m) and §3(n) define Tier 2 and Tier 1 systems, using the existing §3(e) definition of "materially support." §5 is restructured: (a) the agency publishes its tier determination with a short statement of reasons and must treat the system as Tier 2 if it cannot decide; (b) a Tier 1 system files a short-form dossier (purpose, legal authority, data categories and sources, vendor, lineage finding, agency contact and contest procedure, prohibited uses, sunset); (c) a Tier 2 system, which includes every heightened-tier system whatever its use, files the existing full dossier; (d) a change in use that makes a Tier 1 system part of adverse decisions is a material change requiring a full dossier [60] days before the changed use; (e) the State Auditor may reclassify a system to Tier 2; (f) the publication and confidential-annex rule is unchanged. A drafting note follows §5.
- **Conforming changes.** F.1 §3(i), §4(a), §4(c), §6, §9(a), §9(b), and §12(a) now refer to the tiered dossier. §3(g), §7, §10, §11, and §13 are unchanged. E.1 and E.2 are unchanged; the state-version introduction now notes that E.2 has no short form.
- **Appendix F prose.** The state-version introduction, the list of differences from the federal version, and the talking points (what the bill does, what it does not do, the cost question, and the modernization objection) now describe the two tiers.
- **MiDAS wording.** The talking points now say MiDAS "automatically accused" about 40,000 people of unemployment fraud, rather than "falsely accused," matching Section 06 and the sources: the system made the accusations automatically, and a review of about 22,000 determinations found 93% did not involve fraud.
- **Tier 1 terminology.** The talking points now call the Tier 1 filing a "short-form dossier," matching F.1 §3(i) and §5(b), instead of a "short-form public notice."

Presentation and structure

- **Byline.** The brief now names its author, Alec F., under the title and in the closing note of the web version, on the title page of the PDF and Word versions, and in the Word file's author property.
- **Byline and punctuation.** The byline now reads "Alec F., Independent Researcher" under the web title, on the title page of the PDF and Word versions, and in the closing note. Em dashes throughout the brief, including the appendices and bill text, were

replaced with commas, colons, parentheses, or separate sentences; en dashes in number ranges are unchanged.

- **Equation (Section 10).** The original described a formula but did not display it, and raw markup leaked into the text. The equation is now displayed, every symbol is defined, and each step's factor is floored at zero. The interactive figure is labeled an illustrative example with hypothetical parameters, and the 20.4× ratio appears only there.
- **Risk-log example (Section 7).** Labeled a hypothetical illustration constructed for this brief. It does not depict any real agency system, vendor product, record, or person.
- **Harm cases (Section 06).** A paragraph now explains that Robodebt, SyRI, and MiDAS were not security-derived and are included to show why authorization comes first.
- **Corrections and changelog.** Section 16 now lists only factual corrections to the June edition; all other revision notes moved to this changelog.
- **Cross-references.** Recommendation 3 points to the audit framework in Section 13, not Section 12.

Sources

- Source 61, the Ukrainian Ministry of Defence's September 8, 2026 release on AI-guidance trials, was added; the trials are now sourced to it as well as to The Defender.
- The Maven strike and user figures are attributed to the officials who gave them, as reported by DefenseScoop.
- Every reference URL was checked against the live page. The SBS, BTAH, DHS privacy impact assessment, LAPD Inspector General, and Brussee entries a reviewer flagged were already correct; one broken link (Goodwin) was fixed. Reuters and Justia pages refuse automated requests, so those URLs could not be loaded in the check.

Verification status

- U.S. Code citations in Appendix E were checked against current law; the notes after the model text record where a citation is narrower than it looks. Legislative counsel should still confirm them, and should verify every cite and bracketed choice in F.1 for the relevant state.
- The EU AI Act Article 3(1) definition was confirmed against the current text.

By Alec F., Independent Researcher. Revised edition, October 7, 2026. Interactive presentation of the policy brief "Battlefield Validation Is Not Civil Authorization." Version 1.2. Evidence tags record what the cited sources establish, as defined in Section 09: Confirmed means an act or record is documented, not that every performance claim around it is true; Reported claims have not been independently verified; Inferred marks this brief's own analysis; and Opaque marks facts not located in public sources, which is not by itself proof that they are secret.