FIDUCIARY COMMONS FRAMEWORK

Section-by-section analysis and anticipated committee questions

Prepared for the "For legislators" page at fiduciarycommons.com

Working draft for review | July 2026

Prepared by Michael G. Leahy, J.D.

This document provides two resources for the "For legislators" page on fiduciarycommons.com. Part I contains section-by-section analyses of each of the three foundational statutes (VIDA, PDTA, and GAAFA), written for legislative staff and committee members who need to understand what each section does and why it matters without reading the full statutory text. Part II contains anticipated committee questions and answers, developed from existing correspondence, the advocacy white paper, and the scholarly apparatus.

Both resources are working drafts for review before deployment on the website.

Part I. Section-by-section analysis

Each entry provides a plain-language summary of what the section does and, where relevant, why it matters within the framework's constitutional architecture. Cross-references between statutes are noted where they illuminate the layering structure.

Verifiable Identity and Digital Autonomy Act (VIDA)

VIDA is the foundational layer of the framework. It mandates the technical architecture within which government identity and data systems must operate: citizen-controlled credentials, decentralized verification, prohibition on surveillance-capable infrastructure, and open standards. VIDA answers the question: what must government identity systems look like to satisfy constitutional requirements?

Section 1. Short title

Establishes the Act's name: the Verifiable Identity and Digital Autonomy Act (VIDA).

Section 2. Legislative findings and purpose

Establishes the constitutional and policy basis for the Act. Finds that modern integrated government databases possess all four characteristics the Founders condemned as hallmarks of general warrants. States the Act's purpose: mandating citizen-controlled, privacy-preserving digital identity systems governed by fiduciary duties.

Section 3. Definitions

Defines key terms including personal data, fiduciary duty, verifiable credential, self-controlled identifier, phone home, surveillance, digital wallet, and open standards. These definitions are shared across all three companion statutes.

Section 3A. Prohibition on surveillance architectures

Prohibits government digital identity systems from exhibiting the four characteristics of general warrants: perpetual duration, universal scope, discretionary access, and delegable authority. Prohibits centralized data aggregation enabling mosaic surveillance. Prohibits continuous or retrospective monitoring absent specific judicial authorization for a specific individual.

Section 3B. Identity association as participatory act

Establishes that connecting data to an identifiable person requires that person's active participation through credential presentation. Prohibits identity inference, meaning government may not use behavioral patterns, probabilistic matching, or other methods to associate data with a person without that person's active credential presentation. Extends this prohibition to all entities acting in connection with government functions.

Section 4. Data fiduciary duties

Designates every government agency or contractor handling personal data as a data fiduciary with four enforceable duties: loyalty (no unauthorized repurposing), care (appropriate security and governance), confidentiality (no unauthorized disclosure), and portability (citizens can transfer their data). Establishes enforcement through private right of action with statutory damages. Requires all penalties to be paid from the violating agency's operating budget, not the general fund. Prohibits public indemnification for willful violations.

Section 4A. Protections against AI-enabled threats

Requires cryptographic signatures for identity verification, rendering deepfakes and synthetic media legally irrelevant. Prohibits government use of AI for surveillance or profiling without individual

probable cause and judicial authorization. Requires regulatory review of emerging AI applications before deployment in rights-affecting contexts.

Section 4B. Digital likeness and synthetic identity protections

Prohibits government construction, maintenance, or acquisition of digital likenesses or compiled behavioral profiles except pursuant to a warrant or through the individual's active credential presentation. Extends this prohibition to government purchase of commercially compiled profiles: buying surveillance data from data brokers requires the same constitutional protections as direct government collection.

Section 5. Implementation and transition

Requires a phased migration plan within one year of enactment, prioritizing high-risk systems. Mandates pilot programs with public input before large-scale deployment. Requires training and non-digital alternatives to ensure equitable access.

Section 6. Open digital identity standards

Requires all government identity systems to support verifiable credentials and decentralized identifiers based on open standards (W3C specifications). Mandates selective disclosure capability. Prohibits mandatory use of any single private platform. Requires publicly auditable trust registries free from exclusive control.

Section 7. Verifiable digital identity framework

The Act's most detailed technical section. Establishes sixteen specific requirements for government identity systems, including: holder control of credentials with no government monitoring of presentations; purpose-limited verification through cryptographic queries rather than raw data transmission; prohibition on phone-home capabilities; Zero Trust architecture prohibiting shared secrets and requiring cryptographic authentication; preference for government endorsement of credentials over direct issuance; interoperability through open standards; monopoly prevention through FRAND terms; and multi-stakeholder governance of trust registries.

Section 8. Data minimization and sequestered databases

Prohibits centralized repositories beyond what is necessary for a declared, legally authorized purpose. Requires purpose-specific, sequestered datasets with no lateral access between datasets. Prohibits centralized multipurpose data lakes. Requires Zero Trust architecture across all government data infrastructure. Mandates that agencies confirm facts without collecting supporting data wherever feasible.

Section 9. Public participation and accessibility

Requires equitable access for all communities, including those with disabilities, limited internet access, or limited English proficiency. Requires plain-language impact assessments with public comment before adopting new identity systems.

Section 10. Citizen rights and enforcement mechanisms

Establishes five individual rights: to know what data government holds, to access and correct it, to withhold consent for secondary uses, to use digital services without surveillance, and to bring civil actions for violations. Establishes the Data Fiduciary Oversight Office with investigative, enforcement, and penalty authority. Grants enforcement authority to the Attorney General.

Section 11. Implementation and timelines

Requires technical standards within seven months. New services must comply within 24 months; existing services within 48 months with possible extension.

Section 12. Digital equity and inclusion mandate

Requires non-digital alternatives, community partnerships for onboarding, and inclusive interfaces with multilingual and assistive capabilities.

Section 13-18. Severability, effective date, statutory construction, sunset review, funding, and oversight

Standard legislative provisions. Effective date 180 days after enactment. Five-year sunset review cycle. Funding provisions coordinated across all three statutes. Requires construction in harmony with PDTA and GAAFA, with the provision offering greater citizen protection governing any conflict.

Personal Data Trusteeship Act (PDTA)

PDTA is the legal superstructure. It establishes the fiduciary relationship between government and citizen, designates citizens as primary trustees of their own data, imposes enforceable duties on government as secondary fiduciary, and creates the private right of action that gives citizens direct enforcement power. PDTA answers the question: what legal obligations does government owe citizens regarding personal data, and how are those obligations enforced?

Section 1. Short title

Establishes the Act's name: the Personal Data Trusteeship Act (PDTA).

Section 2. Findings and purpose

Finds that personal data constitutes a commons of unownable facts; that traditional ownership models lead to commodification and inequity; that government bears enhanced fiduciary duties; and that the Fourth Amendment's particularity requirement prohibits general-warrant-style database architectures. States the purpose: establishing individuals as primary trustees, imposing fiduciary obligations, and providing private rights of action.

Section 3. Definitions

Defines key terms including personal data, data subject trustee (the individual), secondary fiduciary (any entity handling personal data, including AI agents), fiduciary duties, data commons, and the oversight bodies (Data Trusteeship Authority and Data Fiduciary Oversight Office).

Section 4. Establishment of personal data trusteeship

The Act's foundational provision. Designates every individual as the primary trustee of their own personal data, with overriding authority to manage, access, correct, and control its use. This trusteeship is inalienable: it cannot be waived, contracted away, or extinguished. Treats personal data as a commons ineligible for ownership claims. Requires secondary fiduciaries to obtain explicit, revocable authorization for any data handling.

Section 4A. Bilateral trust structure

Establishes the framework's core relational architecture. Individuals are primary trustees; government entities are limited secondary fiduciaries authorized only to endorse or verify factual claims without creating centralized repositories. Prohibits repurposing data beyond the original justified purpose without fresh authorization, traced directly to the Fourth Amendment's particularity requirement.

Section 4B. Fiduciary relationship arising from collection

Establishes that fiduciary duties arise at the moment of data collection, not at the point of any consent interaction or contract. Any entity collecting personal data in connection with a government function assumes fiduciary duties from the moment of collection, regardless of terms of service or click-through agreements. Fiduciary protections cannot be disclaimed through adhesion contracts.

Section 5. Duties and rights of data subject trustees

Establishes trustee duties (maintaining data accuracy) and rights: access to comprehensive audit logs, veto and revocation of data uses, correction and deletion, delegation to trusted agents, portability, and the right to algorithmic explanation (enforced through GAAFA). Includes guardianship provisions for minors and incapacitated individuals.

Section 6. Obligations of secondary fiduciaries

Imposes four enforceable duties: loyalty (act in the data subject's best interests), care (appropriate security and governance), confidentiality (no unauthorized disclosure), and impartiality (no conflicting interests). Places the burden of proof on the secondary fiduciary in all government-citizen

data disputes. Establishes data minimization as a duty of care. Makes the retaining agency jointly and severally liable for vendor breaches, with third-party beneficiary enforcement rights for citizens.

Section 6A. Architectural requirements for fiduciary compliance

Establishes that fiduciary duties require architectural implementation, not just procedural compliance. Mandates purpose-limited sequestration of data; zero-knowledge verification (where verification suffices, collection is prohibited); individual-controlled immutable audit trails; and cryptographic authentication rendering presentation media legally irrelevant.

Section 6B. Prohibition on government acquisition of commercially compiled surveillance data

Prohibits government purchase, licensing, or receipt of commercially compiled personal data (including browser tracking, location data, biometric data, and behavioral profiles) except pursuant to a warrant or lawful court order. Applies regardless of anonymization or de-identification. Data acquired in violation is inadmissible and derivative evidence is subject to suppression.

Section 7. Government duties

Establishes government's enhanced fiduciary role with specific obligations: favor individuals in power-imbalanced disputes; provide resources enabling effective trusteeship; establish a Data Trusteeship Authority with investigative, audit, and penalty authority; prohibit data uses that undermine autonomy; and require impact assessments for data initiatives.

Section 8. Private rights of action

Grants any data subject trustee the right to sue any secondary fiduciary or government entity for violations. Prohibits pre-dispute arbitration clauses and class action waivers. Eliminates administrative exhaustion requirements. Five-year statute of limitations with a discovery rule for concealed violations.

Section 9. Remedies

Comprehensive remedies including injunctive relief, compensatory damages, statutory damages (minimum $1,000 per violation), punitive damages for egregious breaches, and attorney's fees. Government sovereign immunity waived for monetary awards paid from the violating agency's operating budget. No public indemnification for willful violations.

Section 10. Enforcement

Data Trusteeship Authority has primary enforcement with subpoena power, civil penalties, cease-and-desist authority, and criminal referral. Coordinates with existing enforcement agencies. Annual reporting to the legislature.

Section 11-15. Severability, effective date, statutory construction, sunset review, and funding

Standard legislative provisions. Effective date 180 days after enactment. Five-year sunset review. Funding coordinated across all three statutes. Construction in harmony with VIDA and GAAFA.

Government Algorithmic Accountability and AI Fiduciary Act (GAAFA)

GAAFA is the accountability layer. It governs how government uses AI and automated decision systems within the architecture VIDA mandates and under the fiduciary obligations PDTA imposes. GAAFA answers the question: when government deploys AI to make decisions affecting citizens, what accountability, transparency, and oversight requirements apply?

Section 1. Short title

Establishes the Act's name: the Government Algorithmic Accountability and AI Fiduciary Act (GAAFA).

Section 2. Legislative findings and purpose

Finds that AI systems are being deployed in rights-affecting government contexts without adequate accountability frameworks; that opacity of government AI is a structural threat to due process and equal protection; that autonomous AI agents exercise fiduciary power and must be governed accordingly; and that durable AI governance must be grounded in constitutional principles rather than technical specifications, drawing on the EU AI Act's implementation difficulties as a cautionary example.

Section 3. Definitions

Defines key terms including adverse automated decision, agentic AI system, algorithmic deployment assessment, automated decision system, accountable human official (defined to exclude rubber-stamping), black-box model, disparate impact, interpretable model, rights-affecting decision, and the three tiers of AI use.

Section 4. Constitutional and fiduciary foundations

Establishes that all four PDTA fiduciary duties (loyalty, care, confidentiality, transparency) apply with full force to AI systems. No classification of AI output as "advisory" or "supplementary" relieves the deploying agency of fiduciary obligations. Specifies what each duty requires in the AI deployment context: loyalty means purpose limitation; care means pre-deployment assessment and ongoing monitoring; transparency means algorithmic explanation and public registration; confidentiality means vendor access limitations.

Section 4A. Algorithmic inference as identity construction

Classifies algorithmic inference (deriving personal attributes from data that does not explicitly contain them) as digital likeness construction. Prohibits government use of inferred attributes (health conditions, political beliefs, sexual orientation, and others) unless strictly necessary, disclosed, and independently verifiable. Inferred attributes may not serve as the sole basis for rights-affecting decisions.

Section 4B. Government use of algorithmically processed commercial data

Requires full algorithmic accountability compliance when government uses outputs derived from commercially compiled data. The commercial origin of the data does not excuse compliance with deployment assessments, explanation rights, tiered oversight, or audit trail requirements. Commercially derived risk scores and classifications are subject to the same standards as directly deployed government systems.

Section 5. Algorithmic deployment assessment

Requires a comprehensive pre-deployment filing at least 60 days before any Tier 1 deployment. Must document: system description in plain language, decision logic, training data, accuracy and error rates disaggregated by protected characteristics, disparate impact analysis, red-team testing results, human

oversight architecture, and compliance certifications for VIDA and PDTA. Filed assessments are public by default. Material changes require updated assessment.

Section 6. Right to algorithmic explanation

Establishes the right of any individual receiving an adverse automated decision to receive, within 30 days, a plain-language explanation of the principal factors, what would change the outcome, the identity of the accountable human official, the audit trail, the appeal pathway, and the public deployment assessment. An agency's inability to explain is presumptive evidence of breach of care.

Section 7. Tiered AI use requirements and the interpretable model preference rule

Establishes three tiers. Tier 1 (rights-affecting): interpretable models required unless demonstrated necessity for black-box; every adverse decision requires accountable human official review; black-box exceptions expire after 24 months. Tier 2 (administrative): abbreviated deployment notice, annual impact assessment, drift prevention controls. Tier 3 (citizen-facing information): mandatory disclosure of AI interaction, escalation pathway to human officials, accuracy liability, safeguards against hallucination.

Section 8. Agentic AI systems as statutory secondary fiduciaries

Classifies every autonomous AI agent deployed by government as a secondary fiduciary under PDTA. The deploying agency is jointly and severally liable for the agent's fiduciary breaches. No defense based on the automated nature of the action. Requires enhanced audit trails logging every query, action, authorization basis, and responsible human official. Citizens have full access to agentic audit trails. Vendor agentic AI contracts must designate the vendor as secondary fiduciary and prohibit commercial use of citizen data.

Section 9. Algorithmic impact assessments with binding consequences

Requires ongoing assessments (annual for interpretable Tier 1 and Tier 2; quarterly for black-box exception systems). Findings trigger mandatory consequences: disparate impact requires remediation within 90 days or suspension; systematic error triggers the same; excessive automated approval rates trigger oversight review. Citizens affected during periods of identified bias have reopener rights to de novo human review.

Section 10. Technical AI oversight capacity

Supplements the VIDA oversight office and PDTA trusteeship authority with AI-specific mandates: technical staffing with demonstrated AI expertise, audit capability, transparency registry, and annual AI deployment report. Establishes a Joint AI Audit and Review Panel with independent members including technical experts, civil liberties advocates, and community representatives.

Section 11. Enforcement

Private right of action supplementing PDTA's enforcement provisions. No administrative exhaustion required. Remedies include injunctive relief, compensatory damages, statutory damages ($1,000 per violation, $100 per day for continuing violations), punitive damages, attorney's fees, and reopener relief. Class action presumption. Attorney General concurrent enforcement. Sovereign immunity waived. Five-year statute of limitations with discovery rule.

Section 12. Vendor obligations and government procurement standards

Designates vendors of Tier 1 and Tier 2 systems as secondary fiduciaries. Requires complete documentation disclosure, prompt notification of material changes, cooperation with audits, and indemnification. Procurement contracts must require disparate impact testing, FRAND terms, data portability, and prohibition on vendor commercial use of citizen data. Incentives for open-source procurement.

Section 13. Implementation and timelines

Coordinated with VIDA and PDTA timelines. Technical standards within seven months. Inventory of existing Tier 1 systems within 12 months. New systems compliant within 24 months. All existing Tier 1 systems compliant within 48 months (60-month waiver with interim safeguards). Immediate assessment required for systems with credible bias complaints or adverse judicial findings.

Section 14-16. Statutory construction, severability, and sunset review

Standard provisions. Construction in harmony with VIDA and PDTA. Five-year sunset review. Expressly does not prohibit government use of AI; establishes the conditions under which government may lawfully use AI systems.

Part II. Anticipated committee questions and answers

These questions represent the concerns legislators and committee staff most commonly raise when evaluating the framework. Answers are drawn from the advocacy white paper, the SSRN essays, the Chatrie law review article, existing correspondence, and the framework's design principles.

What does enacting this framework cost?

The framework is designed to align with procurement cycles states already operate on. States rebuild identity infrastructure, database systems, and benefits platforms continuously; the question is whether the next procurement cycle purchases surveillance-capable architecture or constitutionally compliant architecture. The cost differential is in design specifications, not new spending. Specific fiscal items include the Data Fiduciary Oversight Office and Data Trusteeship Authority (both can be structured as divisions within existing IT or AG offices), the grant program for agency compliance (VIDA Section 4(c)), and audit trail infrastructure. Utah enacted identity-layer elements (SB 260 and SB 275) without significant appropriation controversy. The statutes include bracketed provisions for legislative specification of funding sources and penalty amounts.

Does this framework interfere with legitimate law enforcement?

No. The framework's companion statute on automated license plate readers (VLPFA) demonstrates the principle in concrete terms: a plate scan checks against active warrants and produces a hit-or-no-hit result without retaining records of non-matching plates. Law enforcement gets the information it needs; the surveillance capability is never created. The same principle governs the foundational statutes. The framework does not prohibit law enforcement access to data. It requires that access satisfy constitutional standards that already apply: probable cause, particularity, and judicial authorization. Officers seeking specific information about specific individuals through specific warrants are fully accommodated. What the framework prohibits is the architecture that gives officers standing, comprehensive, discretionary access to integrated databases without individualized justification, which is the digital equivalent of the general warrants the Fourth Amendment was ratified to prohibit. The PDTA's purpose-limited sequestration requirement means a law enforcement query returns only the data relevant to the authorized purpose, not the citizen's entire cross-agency profile.

How does this interact with existing state privacy law?

The framework is designed to complement, not replace, existing state privacy statutes. The cross-referencing provisions (VIDA Section 15, PDTA Section 13, GAAFA Section 14) provide that in any conflict, the provision offering greater protection to the citizen governs. State consumer privacy laws (such as those modeled on the CCPA or Virginia CDPA) address commercial data practices. The framework addresses government data practices, operating on the distinct constitutional authority of the government-citizen fiduciary relationship. The two-layer architecture is deliberate: the foundational three statutes derive from the Fourth Amendment and fiduciary principles inherent in the government-citizen relationship; commercial data practices are addressed separately through the CSDLPA, which operates on the state's authority over commercial conduct and dignitary rights.

What about federal preemption?

The framework is specifically designed to resist preemption. Its constitutional grounding in the Fourth Amendment and the government-citizen fiduciary relationship, rather than in general consumer protection authority, provides a stronger preemption defense than privacy frameworks built on commerce-clause authority. State constitutional privacy provisions (present in many state constitutions) provide independent authority that federal legislation cannot ordinarily reach. The framework's alignment with NIST standards and federal technical guidelines further strengthens the

preemption defense: states are implementing the strongest architecture the federal government's own standards recognize. Utah's enactment of framework elements confirms that constitutional grounding provides legislative durability across the political spectrum.

What are the vendor and procurement implications?

The framework requires that procurement contracts include specific fiduciary provisions: vendors handling citizen data become secondary fiduciaries under the PDTA; agencies remain jointly and severally liable for vendor breaches; procurement must require training data disclosure, disparate impact testing, data portability clauses, and prohibition on vendor commercial use of citizen data. These requirements do not exclude vendors; they specify what vendors must provide. The FRAND (fair, reasonable, and non-discriminatory) terms requirement and the prohibition on exclusive licensing prevent vendor lock-in. The open-source preference for AI systems incentivizes, but does not mandate, transparent models. Vendors who already comply with NIST standards will find the framework's technical requirements familiar; the framework adds the legal accountability layer that NIST cannot provide.

What is the relationship between this framework and SEDI?

SEDI (State-Endorsed Digital Identity) addresses the identity verification layer. The Fiduciary Commons Framework adds two layers SEDI does not reach: enforceable fiduciary duties with private right of action (PDTA) and algorithmic accountability (GAAFA). The relationship is successive, not competitive: SEDI establishes how identity is verified; the framework adds the constitutional and legal superstructure governing what government may do with the information identity verification generates. VIDA's technical architecture is designed to be compatible with and build upon SEDI-compliant infrastructure. A state that has enacted SEDI-model legislation has built the identity foundation; enacting PDTA and GAAFA completes the constitutional architecture. Utah is in conversation with multiple states considering joining the SEDI Consortium. The framework's model statutes are designed to build on the SEDI identity layer with enforceable fiduciary duties and algorithmic accountability that SEDI's identity-focused architecture does not reach.

How does this address AI governance specifically?

GAAFA establishes a tiered governance framework. Tier 1 (rights-affecting decisions such as benefits eligibility, licensing, enforcement) requires interpretable, auditable models with mandatory human review of every adverse decision. Tier 2 (internal administrative uses) requires abbreviated deployment notices and annual impact assessments. Tier 3 (citizen-facing information systems including chatbots) requires disclosure, escalation to human officials, and accuracy liability. Autonomous AI agents are classified as statutory secondary fiduciaries with the deploying agency jointly liable for their breaches. The framework governs the government's relationship to the citizen regardless of which technology mediates it, avoiding the EU AI Act's difficulty of regulating the technology itself as it evolves.

What happens if a state enacts only one or two of the three statutes?

Each statute is designed to be independently enactable, though the full architecture requires all three. VIDA alone governs identity infrastructure: a state gets citizen-controlled credentials, privacy-preserving verification, and prohibition on surveillance architectures, but without the enforceable fiduciary duties (PDTA) or AI accountability (GAAFA). VIDA plus PDTA provides the identity architecture and the legal superstructure of fiduciary duties but leaves the AI governance gap. The layering logic is intentional: VIDA is the foundation because no fiduciary duty can be enforced against a surveillance-capable architecture, and PDTA is the legal superstructure because architectural compliance without enforceable legal obligations depends on continued administrative

will. The framework's design encourages sequential adoption, starting with VIDA (which has the strongest existing legislative precedent in Utah's SB 260).

How do you enforce architectural requirements without mandating specific technologies?

The framework specifies constitutional requirements, not technical implementations. VIDA Section 7 requires verifiable credentials, decentralized identifiers, zero-knowledge verification, and Zero Trust architecture as categories of technical capability, not as specific products. The technical standards required within seven months of enactment (VIDA Section 11) translate these categories into implementation specifications using established open standards (W3C VC, W3C DID, NIST SP 800-63-4). This is the same approach used in structural injunction jurisprudence: courts require that prisons provide adequate medical care without specifying which medications to stock; the framework requires that identity systems prevent surveillance without specifying which vendor's software to install. The Data Fiduciary Oversight Office certifies systems for compliance; the state retains full discretion over which compliant systems to procure.

Is this politically viable? Can it pass?

Utah's SEDI legislation (SB 260 and SB 275) passed without a single dissenting vote in either chamber. Utah is in conversation with multiple states considering joining the SEDI Consortium. The framework is grounded in constitutional principles (Fourth Amendment, fiduciary duty, public trust) that resonate across the political spectrum: limited government and individual liberty for conservative members; enforceable civil rights protections and accountability for progressive members. The institutional drift framing, attributing surveillance infrastructure to procurement incentive structures rather than deliberate wrongdoing, offers current officials a path to compliance without moral accusation. The framework is explicitly not a privacy-versus-security trade-off; its security argument (purpose-sequestered databases limit breach scope, distributed credentials eliminate central targets) is independently compelling for technology officials. The StateTech Magazine article and IAPP publication demonstrate engagement with the practitioner community that builds legislative credibility.