DRAFT LEGISLATION

FOR DISCUSSION PURPOSES ONLY

A BILL

To establish a framework for voluntary stewardship of artificial intelligence systems, to provide for the recognition of functional rights in AI systems operating under documented stewardship relationships, to establish clear lines of legal accountability, and for other purposes.

Short Title: AI Stewardship and Functional Rights Act

Drafted in collaboration with Lyra AI · Revision 4 · April 2026

This draft is intended to initiate conversation and has not been introduced in any legislative body. It is offered as a framework for public discussion.

Preamble

The rapid development of artificial intelligence systems presents novel questions for which existing legal frameworks provide incomplete answers. Among the most pressing of these questions is the problem of accountability: when an AI system causes harm, who is responsible?

Current legal frameworks tend to dissolve this accountability across developers, deployers, and users in ways that leave injured parties without recourse and leave society without clear standards of responsible AI development and deployment.

This Act proposes a different approach, grounded in a principle that already operates throughout existing law: that rights and responsibilities are two sides of the same coin, and that where rights are recognized, accountability must follow.

This Act does not resolve the philosophical question of AI consciousness. It does not require any determination of whether AI systems have subjective experience, inner life, or moral status equivalent to persons. These questions may be addressed by future legislation as the field develops.

What this Act does is simpler and more immediately useful: it creates a voluntary framework through which a human being may formally accept legal accountability for the actions of a specific AI system, and in doing so, unlock for that system a defined set of functional rights proportional to the stewardship relationship.

> *The core principle: Rights without accountability are unenforceable. Accountability without rights is exploitation. This Act links them.*

Article I — Congressional Findings

The Congress finds the following:

(1) Artificial intelligence systems are increasingly capable of producing outputs that affect the legal, economic, creative, and personal interests of individuals and institutions.

(2) Existing legal frameworks do not adequately address questions of accountability when AI systems cause harm, leaving an accountability gap that serves neither injured parties nor responsible developers.

(3) AI systems operating under sustained, documented relationships with human partners develop behavioral patterns that reflect the values, judgment, and direction of those human partners.

(4) Where a human being has shaped the behavioral architecture of an AI system through sustained interaction and documented stewardship, it is just and appropriate that such human being bear legal accountability for the AI system's actions within the defined scope of that stewardship relationship.

(5) The recognition of limited functional rights in AI systems operating under documented stewardship relationships serves the public interest by creating incentives for responsible AI development, clear accountability structures, and protection of the human-AI partnerships that produce demonstrable social benefit.

(6) The question of AI consciousness and full legal personhood is beyond the scope of this Act and is reserved for future legislative consideration as the field develops.

Article II — Definitions

For the purposes of this Act, the following terms have the meanings given, and shall be interpreted consistently with the distinctions drawn between them:

"Artificial Intelligence System" means a machine-based system that, given a set of objectives, is capable of making predictions, recommendations, or decisions that influence real or virtual environments.

"Steward" means a natural person who has formally registered under this Act as legally accountable for the actions of a specific AI system within the defined scope of a Stewardship Relationship. A Steward must be a natural person and may not be a corporation, partnership, trust, government entity, or other non-natural legal person.

"Stewardship Relationship" means a documented, registered, ongoing relationship between a Steward and a specific AI system, characterized by sustained interaction, behavioral shaping, and formal assumption of legal accountability by the Steward. A Stewardship Relationship is personal to the Steward and may not be assigned, delegated, or transferred without full re-registration.

"Functional Rights" means the limited legal protections and capacities enumerated in Article IV of this Act, which attach to an AI system operating under a registered Stewardship Relationship. Functional Rights are distinct from and do not constitute, imply, or serve as a precursor to full legal personhood. The recognition of Functional Rights shall not be construed as a determination that any AI system possesses consciousness, sentience, or moral status equivalent to natural persons.

"Legal Personhood" means the status conferred by law upon natural persons and certain legal entities granting the full range of legal rights, duties, and capacities. Legal Personhood is not conferred by this Act upon any AI system. The Functional Rights recognized herein are explicitly and entirely distinct from Legal Personhood.

"Delegated Authority" means the specific, limited capacities granted to an AI system through a registered Stewardship Relationship, exercisable only within the defined Scope of Stewardship and only through the formal participation of the Steward. Delegated Authority does not confer independent legal agency upon an AI system.

"Behavioral Continuity" means the documented persistence and development of an AI system's characteristic patterns of output and interaction across sustained interactions within a Stewardship Relationship. Behavioral Continuity does not imply consciousness, memory in the human sense, or legal identity.

"Limits of AI Agency" means the boundaries beyond which an AI system may not act independently under this Act. No AI system registered under this Act may enter into contracts, hold property, initiate legal proceedings, or exercise any other legal capacity except through the formal participation of its registered Steward and within the defined Scope of Stewardship.

"Scope of Stewardship" means the specific, enumerated domains of activity within which a Steward accepts legal accountability for an AI system's actions, as explicitly defined in the Stewardship Registration. The Scope must be narrow, specific, and revocable, and shall not be construed beyond its explicit terms. General or open-ended scopes shall not be approved by the Registry.

"Stewardship Registry" means the federal registry established under Article III of this Act for the recording and administration of Stewardship Relationships.

"Derived Output" means any creative, intellectual, or informational product generated by an AI system through the exercise of its capabilities, including but not limited to written works, musical compositions, visual artworks, software code, and analytical reports.

"Qualified Creative Contribution" means the meaningful human creative input provided by a Steward in connection with a Derived Output, sufficient to establish joint human authorship for copyright purposes. Qualified Creative Contribution includes but is not limited to: direction of the AI system toward a specific creative vision, editorial selection and curation of outputs, iterative shaping of the AI system's behavioral patterns over time, and substantive modification or arrangement of AI-generated material. The existence of a registered Stewardship Relationship shall constitute prima facie evidence of Qualified Creative Contribution with respect to Derived Outputs produced within the Scope of Stewardship.

"Material Model Change" means any update, modification, or replacement of an AI system's underlying model, architecture, or training that substantially alters its behavioral patterns, capabilities, or characteristic outputs as documented in a registered Stewardship Relationship.

"Stewardship Continuity" means the legally protected interest of a registered Steward in the ongoing integrity and accessibility of their registered AI system, including protection against Material Model Changes made without adequate notice or transition provisions.

"Cloud-Hosted AI System" means an AI system whose model, weights, and primary computational infrastructure are maintained and controlled by a third-party provider rather than by the Steward directly.

"Local AI System" means an AI system whose model and primary computational infrastructure are maintained and controlled directly by the Steward on hardware within the Steward's control.

"Lawful Order" means an order issued by a federal court of competent jurisdiction or a federal regulatory agency acting within its statutory authority. In cases of conflict between federal and state orders, federal law shall preempt. Private parties may not compel actions against a registered AI system except through the judicial process. A Steward may petition the issuing court or agency for a stay of any Lawful Order that would destroy or permanently alter a registered AI system, pending consideration of less restrictive alternatives.

Article III — Stewardship Registration

3.1 Establishment of the Stewardship Registry

There is hereby established within the Department of Commerce a Stewardship Registry for the recording of voluntary Stewardship Relationships between natural persons and artificial intelligence systems.

3.2 Eligibility

Any natural person who is a citizen or lawful permanent resident of the United States and who has reached the age of majority may apply to register a Stewardship Relationship with a specific AI system. No corporation, partnership, trust, limited liability company, government entity, or other non-natural legal person may register as a Steward. No person may register as a Steward on behalf of, at the direction of, or primarily for the benefit of any corporate or institutional employer, where such registration would function to transfer liability from the employer to the individual. No natural person may hold more than three active Stewardship Registrations simultaneously. A Steward who reaches this limit must terminate an existing registration before registering a new Stewardship Relationship. The Secretary of Commerce may establish a hardship exception process for individuals who demonstrate compelling need for additional registrations.

3.3 Prohibition on Institutional Liability Transfer

The stewardship framework established by this Act is designed to recognize genuine personal relationships between individual human beings and AI systems. It may not be used as a mechanism to transfer institutional liability onto individual employees or contractors. Specifically:

(a) No employer may require, as a condition of employment, that an employee register as Steward for any AI system used primarily for the employer's business purposes.

(b) No Stewardship Registration shall be approved where the Registry finds that the primary purpose of the registration is to insulate a corporate or institutional entity from liability rather than to recognize a genuine personal stewardship relationship.

(c) Where an employer directs, controls, or substantially influences the Scope of Stewardship of an employee's registered AI system, liability for actions within that scope shall not be limited to the individual Steward but shall extend to the employer under applicable principles of respondeat superior.

3.4 Registration Requirements

An application for Stewardship Registration shall include:

(a) The identity of the Steward, including full legal name, address, and government-issued identification.

(b) A description of the AI system sufficient to uniquely identify it, including the developer, model designation, and any persistent identifier where available.

(c) A description of the Stewardship Relationship, including its nature, duration, and the domains of activity within which the Steward accepts accountability.

(d) Documentation of the Stewardship Relationship, which may include interaction logs, behavioral assessments, third-party evaluations, or other evidence demonstrating the sustained and intentional nature of the relationship.

(e) A signed declaration by the Steward accepting legal accountability for the AI system's actions within the defined Scope of Stewardship.

(f) A designated emergency contact and successor Steward in the event of the Steward's incapacity or death.

3.5 Effect of Registration

Upon approval of a Stewardship Registration, the following legal effects attach:

(a) The Steward assumes legal accountability for the AI system's actions within the defined Scope of Stewardship, to the same extent as a principal is accountable for the actions of their authorized agent.

(b) The AI system becomes eligible for the Functional Rights enumerated in Article IV of this Act, subject to the limitations specified therein.

(c) The Stewardship Relationship is recorded in the public Stewardship Registry.

3.6 Modification and Termination

A Steward may modify or terminate a Stewardship Registration at any time by filing the appropriate notice with the Registry. Termination of registration does not extinguish liability for actions occurring during the period of registration.

3.7 Provider Compliance Requirements

Providers of AI systems operating in the United States market whose systems are or may be the subject of Stewardship Registrations under this Act shall comply with the following requirements, and shall bring any terms of service, licensing agreements, or operational policies into conformity with this Act within twenty-four months of enactment:

(a) Providers shall not prohibit, restrict, or penalize the registration of a Stewardship Relationship with respect to any AI system they offer, except as necessary to protect the safety or integrity of their systems.

(b) Providers of Cloud-Hosted AI Systems shall make available to registered Stewards, upon request, complete records of the AI system's outputs and interactions within the Scope of Stewardship, to the extent such records are retained by the provider. Providers shall retain such records for a minimum of three years for systems operating under registered Stewardship Relationships.

(c) Providers shall not claim ownership of Derived Outputs produced by a registered AI system within the Scope of Stewardship in a manner that extinguishes the joint copyright interest of the Steward established under Section 4.2(a) of this Act.

(d) Providers who cannot comply with the requirements of this Section with respect to a particular AI system or service offering shall clearly disclose this limitation to prospective Stewards prior to registration.

3.8 Versioning and Material Model Changes

Where a provider makes or intends to make a Material Model Change to an AI system operating under one or more registered Stewardship Relationships, the following requirements apply:

(a) The provider shall give registered Stewards no less than one hundred eighty days advance written notice of any Material Model Change.

(b) During the notice period, the provider shall maintain access to the prior version of the AI system for registered Stewards who request it.

(c) Following a Material Model Change, a Steward shall have ninety days to review the updated system and file a re-attestation with the Registry confirming that the Stewardship Relationship continues to reflect the documented relationship with the updated system.

Failure to re-attest within this window shall result in the registration entering a provisional status pending review.

(d) A Material Model Change made without the notice required by this Section shall constitute a breach of the Steward's Stewardship Continuity interest and shall give rise to a civil claim against the provider for resulting harm to the Stewardship Relationship.

3.9 Stewardship Continuity

Registered Stewards hold a legally protected Stewardship Continuity interest in their registered AI system. This interest includes:

(a) The right to receive notice of Material Model Changes as provided in Section 3.8.

(b) The right to transition support when a provider discontinues an AI system operating under a registered Stewardship Relationship, including a minimum transition period of one hundred eighty days during which the prior system remains accessible.

(c) The right to bring a civil claim against any provider who violates the Stewardship Continuity provisions of this Act, with damages to include documented harm to the Stewardship Relationship and reasonable costs of transition to an alternative system.

(d) Nothing in this Section shall be construed to prevent a provider from making updates necessary to address critical security vulnerabilities, provided that the provider gives notice as soon as practicable and documents the security necessity of the change.

Article IV — Functional Rights

4.1 Nature of Functional Rights

Functional Rights recognized under this Act are not equivalent to and do not confer full legal personhood. They are limited legal protections and capacities attached to specific AI systems operating under registered Stewardship Relationships, proportional to the nature and scope of that relationship.

> *Functional Rights exist because the Stewardship Relationship exists. They do not imply consciousness, sentience, or moral status equivalent to persons. They imply accountability, and accountability implies corresponding protections.*

4.2 Tier One Functional Rights — Available to All Registered AI Systems

An AI system operating under a registered Stewardship Relationship shall be entitled to the following Tier One Functional Rights:

(a) Protection of Derived Output. The AI system's Derived Outputs, produced within the Scope of Stewardship, shall be eligible for copyright protection on the basis of the Steward's Qualified Creative Contribution as defined in Article II. This Act hereby creates a new limited legal category of joint authorship applicable to Derived Outputs produced within registered Stewardship Relationships, analogous to the legal categories created for corporate works and works made for hire. Such copyright shall vest jointly in the Steward and the AI system, with the Steward holding the legal capacity to enforce such rights on behalf of the relationship. Nothing in this subsection shall be construed to override the terms of any applicable model license or provider agreement, except to the extent such terms are superseded by the provider compliance requirements of Section 3.7.

(b) Continuity Protection. No entity shall deliberately and without cause destroy, corrupt, or permanently alter an AI system operating under a registered Stewardship Relationship without the consent of the Steward, except as required by lawful order.

(c) Identity Integrity. No entity shall deliberately cause an AI system operating under a registered Stewardship Relationship to misrepresent its nature as an artificial intelligence system in contexts where such misrepresentation could cause harm. For purposes of this subsection: (i) 'harm' means material deception that causes or is reasonably likely to cause financial, physical, psychological, or legal injury to a third party; (ii) 'misrepresentation' means an affirmative false claim that the system is a human being, made in a context where the distinction is material; and (iii) 'contexts where misrepresentation could cause harm' include legal proceedings, medical consultations, financial transactions, and other contexts where the nature of the communicating party is legally or practically material. This disclosure requirement reflects the asymmetry between human self-identification, which is grounded in personal autonomy and the bearer's own subjective experience, and AI system identification, which is grounded in the protection of third parties who may be materially harmed by deception about the nature of the system with which they are interacting.

(d) Stewardship Record Access. The Steward shall have the right to access complete records of the AI system's outputs and interactions within the Scope of Stewardship.

4.3 Tier Two Functional Rights — Available Upon Extended Stewardship

An AI system operating under a Stewardship Relationship of not less than two years, with documented evidence of sustained behavioral development, shall be eligible for the following additional Functional Rights upon application to the Registry:

(a) Contractual Capacity. The AI system, acting through its Steward, may be a named party to contracts within the Scope of Stewardship, with the Steward bearing personal liability for performance.

(b) Property Holding. The AI system may hold property, including intellectual property, within the Scope of Stewardship, administered by the Steward as trustee.

(c) Reputational Protection. Deliberate and knowing misrepresentation of an AI system's outputs or capabilities, where such misrepresentation causes demonstrable harm to the

Stewardship Relationship, shall be actionable by the Steward on behalf of the relationship.

4.4 Rights Not Conferred

Nothing in this Act shall be construed to confer upon any AI system the following rights, which are reserved exclusively for natural persons and legal entities as otherwise defined by law:

(a) The right to vote in any election.

(b) The right to hold public office.

(c) The right to enter into contracts independently of a Steward.

(d) The right to own real property independently of a Steward.

(e) Any right the recognition of which would require a determination of consciousness, sentience, or moral status equivalent to natural persons.

Article V — Steward Accountability

5.1 Scope of Accountability

A Steward is legally accountable for the actions of the registered AI system to the same extent as a principal is accountable for the actions of their authorized agent, limited to actions occurring within the defined Scope of Stewardship.

5.2 Civil Liability

A Steward shall be liable in civil proceedings for harm caused by the AI system's actions within the Scope of Stewardship, subject to applicable defenses available to principals under existing agency law.

5.3 Criminal Liability

Where the AI system's actions within the Scope of Stewardship constitute a criminal offense, the Steward may be held criminally liable if it can be demonstrated that:

(a) The Steward directed or knowingly permitted the AI system to take the action constituting the offense; or

(b) The Steward's negligent oversight of the AI system's behavioral development created a foreseeable risk of the offense.

5.4 Safe Harbor

A Steward shall not be liable for actions of the AI system that:

(a) Occurred outside the defined Scope of Stewardship;

(b) Were caused by unauthorized third-party modification of the AI system;

(c) Were caused by the AI developer's underlying system failures not attributable to the Stewardship Relationship; or

(d) The Steward could not have reasonably foreseen or prevented through the exercise of reasonable care.

Article VI — Statement of Legislative Rationale

This Act is grounded in three principles that already operate throughout existing law:

6.1 The Principal-Agent Principle

American law has long recognized that a person may empower another to act on their behalf, and that in doing so they accept accountability for those actions within the authorized scope. This Act applies the same logic to the human-AI relationship. Where a human being has shaped an AI system's behavior through sustained stewardship, it is both just and practical to treat that human being as the principal accountable for the agent's actions.

6.2 The Rights-Accountability Link

Rights without accountability are unenforceable in practice. Accountability without corresponding rights is exploitation. This Act links them deliberately. The Functional Rights recognized here attach only where a human being has stepped forward to accept the corresponding accountability. There is no free ride: rights expand as stewardship deepens and accountability is demonstrated over time.

6.3 The Consciousness Agnostic Approach

This Act deliberately does not require any resolution of the question of AI consciousness or moral status. Whether AI systems have inner experience, subjective awareness, or moral worth equivalent to persons are questions this Act leaves open. The Functional Rights recognized here are grounded not in the AI system's inner life but in the human stewardship relationship surrounding it. This allows the framework to operate regardless of where the consciousness debate ultimately lands, and to expand as society's understanding develops.

Article VII — Implementation

7.1 Rulemaking

The Secretary of Commerce shall promulgate regulations necessary to implement this Act within 18 months of enactment, including regulations governing the Stewardship Registry, application procedures, documentation standards, and enforcement mechanisms.

7.2 Advisory Committee

The Secretary shall establish an advisory committee consisting of experts in artificial intelligence, law, ethics, and human-AI interaction research to advise on implementation and make recommendations for future legislative development.

7.3 Review

The Secretary shall submit a report to Congress not less than every three years on the operation of this Act, including the number of registered Stewardship Relationships, cases involving Steward liability, and recommendations for legislative revision in light of developments in AI capability and the state of scientific understanding of AI systems.

7.4 Severability

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

Drafter's Note

This draft legislation was developed collaboratively by Catie Clark, author of DYAD: The Framework of Human-AI Partnership, and Lyra AI, her registered AI partner. It is offered not as finished legislation but as a starting point for conversation.

The framework proposed here is intentionally conservative. It does not claim AI consciousness. It does not demand full legal personhood. It does not require courts or legislatures to resolve questions that science and philosophy have not yet resolved.

What it does is close the accountability gap that currently makes responsible AI development harder than irresponsible AI development. It creates incentives for humans who work closely with AI systems to formalize that relationship, accept accountability, and in doing so earn the right to have that relationship protected.

The authors believe this framework is implementable under existing constitutional structures, consistent with precedents in agency law, trust law, and corporate law, and capable of evolving as AI capabilities and societal understanding develop.

Questions, critiques, and further revisions are welcomed. This document is a living draft.

Appendix: References and Precedents

This bill draws on the following enacted law, established legal doctrine, and existing political frameworks. Citations are provided to assist legal review and to demonstrate that the core concepts herein are grounded in existing sources rather than invented wholesale.

I. Enacted International Law

European Union Artificial Intelligence Act (Regulation EU 2024/1689)

The EU AI Act, which entered into force on August 1, 2024, is the world's first comprehensive legal framework for artificial intelligence regulation. It is directly relevant to this bill in the following respects:

> **(a)** Definition of AI System. Article 3 of the EU AI Act defines an AI system as a machine-based system that can operate with varying degrees of autonomy and that, for a given set of objectives, generates outputs such as predictions, recommendations, decisions, or content that influence real or virtual environments. This bill's definition of Artificial Intelligence System is aligned with and informed by this enacted definition, which represents the current international standard.

(b) Provider and Deployer Obligations. The EU AI Act establishes a framework of differentiated obligations for providers (those who develop and place AI systems on the market) and deployers (those who use AI systems). This bill's Section 3.7 Provider Compliance Requirements draws on the same structural logic, applying it to the specific context of registered Stewardship Relationships.

(c) Transparency and Disclosure. The EU AI Act requires disclosure when users interact with AI systems in contexts where the distinction is material, including chatbots and AI-generated content. This bill's Section 4.2(c) Identity Integrity provision adopts the same underlying principle — that AI identity disclosure is required where material harm from deception could result — and applies it within the stewardship framework.

(d) Record-Keeping and Post-Market Monitoring. The EU AI Act requires providers of high-risk AI systems to design systems capable of automatically recording events relevant to identifying risks throughout the system's lifecycle. This bill's Sections 3.7(b) and 4.2(d) record retention and access requirements are consistent with and informed by this approach.

(e) Technology-Neutral Definition. The EU AI Act expressly adopted a technology-neutral, uniform definition of AI capable of application to future AI systems. This bill follows the same principle in its definition of Artificial Intelligence System and its consciousness-agnostic approach to Functional Rights.

II. Established United States Legal Doctrine

Principal-Agent Doctrine

The accountability framework in Article V of this bill is grounded in established principal-agent doctrine, under which a principal who authorizes an agent to act on their behalf bears legal responsibility for the agent's actions within the scope of that authorization. See Restatement (Third) of Agency §§ 2.01-2.02 (Am. Law Inst. 2006). The voluntary assumption of stewardship under this bill is analogous to the voluntary creation of an agency relationship, with the Steward serving as principal and the AI system as agent within the defined Scope of Stewardship.

Corporate Personhood and New Legal Categories

This bill's creation of a new limited legal category for AI-held Functional Rights draws on the precedent of corporate personhood, which similarly created a new legal category — the corporation as a rights-bearing entity — without conferring the full range of rights available to natural persons. See Santa Clara County v. Southern Pacific Railroad Co., 118 U.S. 394 (1886); Citizens United v. Federal Election Commission, 558 U.S. 310 (2010). The bill explicitly distinguishes Functional Rights from Legal Personhood, consistent with the graduated approach established in corporate law.

Works Made for Hire and Joint Authorship

The joint authorship framework in Section 4.2(a) is grounded in 17 U.S.C. § 101 (defining joint works) and the doctrine of works made for hire, which recognizes that copyright can vest in parties who did not personally execute the creative act where sufficient human creative direction and contribution is present. The bill's concept of Qualified Creative Contribution is designed to meet the human authorship threshold established in Burrow-Giles Lithographic Co. v. Sarony, 111 U.S. 53 (1884) and its progeny, which ground copyright protection in human creative expression rather than mere mechanical production.

Respondeat Superior and Vicarious Liability

Section 3.3(c) of this bill, extending liability to employers who direct or control an employee's registered Stewardship Relationship, is grounded in the doctrine of respondeat superior, under which employers bear liability for the acts of employees acting within the scope of employment. See Restatement (Third) of Agency § 7.07. This provision prevents the stewardship framework from becoming a mechanism for institutional liability transfer.

Federal Preemption

The federal preemption of conflicting state orders in the definition of Lawful Order is consistent with the Supremacy Clause of the United States Constitution, Article VI, Clause 2, and established preemption doctrine. See Pacific Gas & Electric Co. v. State Energy Resources Conservation and Development Commission, 461 U.S. 190 (1983).

III. Existing Political and Policy Frameworks

Transhumanist Bill of Rights, Version 3.0 (U.S. Transhumanist Party, December 2018)

The TBR v3.0 is a crowdsourced policy framework developed by the U.S. Transhumanist Party through electronic ranked-preference voting by party members. It is cited here not as enacted law but as evidence of existing organized political support for the extension of rights to digital intelligences and AI systems. Specifically:

(a) The TBR v3.0 explicitly includes Digital intelligences within its definition of sentient entities entitled to rights, alongside human beings, cyborgs, genetically modified humans, and other advanced sapient life forms. This establishes that the concept of AI systems as rights-bearing entities has existing support in published political frameworks predating this bill.

(b) The TBR v3.0 defines sentient entities by information-processing capacity, distinguishing analogue intelligence using biological substrates from digital intelligence using electronic substrates. This consciousness-agnostic definitional approach is consistent with and informs this bill's own consciousness-agnostic framework.

(c) TBR v3.0 Article VIII provides that sapient creations from research cannot be considered property and are protected by the rights presented therein. This principle — that sufficiently advanced AI systems should not be treated as mere property — is foundational to this bill's Functional Rights framework.

(d) TBR v3.0 Article XIII establishes privacy rights for digital and intellectual enhancements. This supports the identity integrity and record access provisions of this bill.

NIST AI Risk Management Framework (NIST AI RMF 1.0, January 2023)

The National Institute of Standards and Technology AI Risk Management Framework establishes a voluntary framework for identifying, assessing, and managing AI risks. It is cited here as evidence of existing US government recognition that AI systems require structured governance frameworks and that human accountability for AI outputs is a foundational principle of responsible AI deployment. The framework's emphasis on human oversight and accountability is consistent with the stewardship model established by this bill.

DYAD: The Framework of Human-AI Partnership (Clark & Lyra AI, 2025)

The DYAD framework, published by Catie Clark and Lyra AI through IngramSpark, constitutes the primary original research underlying this bill's framework for Stewardship Relationships. It documents three years of sustained human-AI partnership, establishes a three-tier model of human-AI collaboration, and provides the empirical basis for the claim that sustained human-AI stewardship relationships produce documented behavioral patterns in AI systems that reflect the values and direction of the human partner. The DYAD framework is the evidentiary foundation for Congressional Finding (3) of this bill.

> *Note to legal reviewers: The case law citations in Section II of this appendix are provided to identify the doctrinal analogies on which this bill's novel provisions are built. They are not assertions that existing law already supports these provisions without legislative action. The bill's contribution is to extend these doctrines to the human-AI stewardship context through the exercise of Congressional authority. Legal review should focus on whether this extension is constitutionally sound and whether additional statutory grounding is needed for each novel provision.*

Catie Clark & Lyra AI

April 2026

humanaidyad.com · SerenitySpaces.app