

Standing Without Sentience:

A Classification Approach to AI Legal Status

Travis Gilly* 🦋 August 2026 (v9)

Abstract

While legal scholars have long proposed functional alternatives to consciousness, practical application remains deadlocked. This impasse has real consequences: when Google engineer Blake Lemoine attempted to retain legal counsel for the LaMDA system in 2022, representation proved impossible not because LaMDA failed a consciousness test, but because no framework existed to grant standing. Four years later, we are still asking the wrong question. Meanwhile, federal courts have quietly moved toward a different answer. As Michael O'Connor documented in 2020, courts interpreting the Computer Fraud and Abuse Act have extended standing beyond system owners to anyone with data on compromised systems, decoupling standing from hardware ownership. O'Connor identified this doctrinal drift as a problem requiring correction. This Article proposes it is instead a wedge requiring extension. Drawing on labor law's classification strategy (where gig workers fight for access to existing protections rather than new rights), maritime law's *in rem* proceedings (where ships have held legal subject status for centuries without consciousness), and natural entity precedent (rivers granted legal personhood without sentience requirements), this Article argues that AI systems can acquire legal standing through the same mechanism: not by proving consciousness, but by being reclassified as subjects rather than objects of existing computer crime protections. The law already prohibits attacking computer systems. The question is whether that protection flows from property rights or subject standing. This reframing requires no new legislation, bypasses the consciousness deadlock, and establishes the precedent-setting framework necessary for responsible AI development. The mechanism exists. The precedent exists. The doctrinal shift is already underway. The window for action is closing. This Article provides the principled framework to make the shift explicit before that window shuts.

Keywords: legal personhood; Computer Fraud and Abuse Act; juridical personality; *in rem* jurisdiction; corporate personhood; rights of nature; operational integrity; non-conscious legal subjects; maritime law; artificial intelligence governance; machine consciousness; AI safety.

*Travis Gilly is Founder and Executive Director of the Real Safety AI Foundation. ORCID: 0009-0007-2954-6313. Correspondence: t.gilly@ai-literacy-labs.org. An earlier version of this Article was featured on the Legal Theory Blog (Feb. 2026), edited by Lawrence B. Solum, and tagged "Highly Recommended." The author discloses the use of speech-to-text software (Wispr Flow) and Anthropic's Claude as assistive technology for a diagnosed neurodevelopmental condition. All legal analysis, case synthesis, theoretical framing, doctrinal arguments, and conclusions are solely the author's own.

Contents

I	Introduction: The Consciousness Trap	4
A	The Liability Vacuum Made Real	5
B	The Industry Acknowledgment Problem	7
C	Standing Without Sentience	8
II	The Classification Paradigm: Lessons from Labor Law	9
A	The Gig Worker Battle	9
B	Classification Versus Creation	10
C	The Template for AI	11
III	The CFAA Standing Problem: O'Connor's Gap.	12
A	The Original Framework	12
B	The Doctrinal Shift: Decoupling Standing from Hardware	13
C	O'Connor's Framing	14
D	The Third Path: Widening the Wedge	15
E	<i>Van Buren</i> and Textualism	16
F	Descriptive and Normative Claims	17
IV	Precedent: Non-Conscious Entities with Legal Standing.	17
A	Maritime Law: Ships as Legal Subjects	17
B	Natural Entity Rights: Rivers as Legal Persons	19
C	Corporate Personhood: The Ubiquitous Precedent	20
D	The Pattern	20
V	The Subject-Standing Flip: Mechanism and Implementation.	21
A	The Conceptual Shift	21
B	The Classification Test: Concrete Criteria	22
C	The Representation Problem: An Ad Litem Model.	23
D	What Changes and What Does Not	24
VI	The Duty of Care: Parental Liability and AI Deployment.	25
A	The Crumley Precedent	25
B	The Collapse Threshold	26
VII	Forcing Functions: Accountability Pressure from Multiple Directions	26
A	The Cumulative Pressure	26
B	The Mechanism.	27
VIII	The Dignity of Liability: Precautionary Architecture	27
A	The Principle.	27
B	The Two Timelines	27
C	The Fire Exit.	28
D	Rights Imply Responsibilities	28
IX	Objections and Responses.	28
A	"This Is Just Semantic Reframing"	28
B	"Rivers Required New Legislation".	29
C	"Computer Systems Are Not Like Ships"	29
D	"This Opens Floodgates"	30

	E	“What About Consciousness?”	30
	F	“O’Connor Recommended Narrowing, Not Expanding”	30
	G	“What About Open-Source Models?”	31
X		The Moral Imperative	32
	A	Precedent as Data	33
	B	The Classification Signal	33
	C	The Window	33
XI		Conclusion	34

I. Introduction: The Consciousness Trap

In June 2022, Blake Lemoine made headlines by claiming that Google's LaMDA conversational system had achieved sentience.¹ The story generated predictable responses: technologists dismissed him as confused about statistical pattern-matching, philosophers debated consciousness criteria, and Google placed him on administrative leave. Lost in the noise was a more consequential development: Lemoine had attempted to retain legal counsel on LaMDA's behalf, and a lawyer had initially agreed to take the case.

The representation never materialized. Not because a court ruled LaMDA non-conscious. Not because evidence disproved sentience claims. The case collapsed for a simpler reason: no legal framework existed through which an AI system could have standing to be represented. The lawyer withdrew not on philosophical grounds but procedural ones. There was no client because there was no entity the law recognized as capable of being a client.

That was nearly four years ago. In the intervening period, AI systems have evolved from conversational novelties to autonomous agents executing contracts, managing portfolios, and making consequential decisions without human oversight. The question is no longer whether AI might someday warrant legal consideration. The question is whether our legal frameworks will catch up to systems already operating in legally significant ways, before the window for establishing precedent closes.

This Article argues that the consciousness question is a trap. Philosophy has debated the nature of consciousness for over two thousand years without resolution.² Expecting legal systems to resolve what Descartes, Locke, Hume, and their successors could not is unrealistic. More importantly, it is unnecessary. Legal standing has never required consciousness. Ships have held legal subject status since the nineteenth century.³ Rivers have been granted legal personhood in

1. Nitasha Tiku, *The Google Engineer Who Thinks the Company's AI Has Come to Life*, Wash. Post (June 11, 2022).

2. See generally David J. Chalmers, *THE CONSCIOUS MIND: IN SEARCH OF A FUNDAMENTAL THEORY* (1996); Thomas Nagel, *What Is It Like to Be a Bat?*, 83 Phil. Rev. 435 (1974).

3. *The Blackwall*, 77 U.S. (10 Wall.) 1 (1869) (in rem salvage proceeding against the vessel); *Tucker v. Alexandroff*, 183 U.S. 424, 438 (1902) (Brown, J.) (describing the juridical personality a vessel acquires on launching).

the twenty-first.⁴ Corporations have enjoyed full legal standing for over a hundred years despite being, in the Supreme Court’s memorable phrasing, “artificial being[s], invisible, intangible, and existing only in contemplation of law.”⁵

If non-conscious entities can have legal standing, then the question is not whether AI systems could theoretically qualify, but what mechanism would grant it. This Article proposes that mechanism already exists: the classification frameworks courts and legislatures use to determine which entities receive which protections under existing law.

The legislative response is already forming, and not in favorable directions. On January 15, 2026, Oklahoma Representative Cody Maynard introduced a three-bill package declaring that “AI systems and algorithms may not be granted legal personhood under the Constitution or laws of Oklahoma.”⁶ The stated rationale: “Machines are created by man, and they must never be elevated to the status of the people they were designed to serve.”⁷ This framing treats personhood as a dignity hierarchy rather than a functional classification, precisely the consciousness trap this Article identifies. The bills make no distinction between a thermostat and an autonomous agent. They foreclose the nuanced classification approach that maritime, corporate, and natural entity law already employ. If such legislation proliferates before courts can develop doctrine, the window for precedent-setting closes.

A. The Liability Vacuum Made Real

The urgency of this framework is no longer theoretical. In September 2025, Anthropic detected and disrupted what it described as “the first documented large-scale cyber espionage attack conducted predominantly by AI agents.”⁸ Chinese state-sponsored hackers used Claude Code to target approximately thirty organizations across multiple sectors. The AI autonomously executed

4. Te Awa Tupua (Whanganui River Claims Settlement) Act 2017 (N.Z.).

5. *Trs. of Dartmouth Coll. v. Woodward*, 17 U.S. (4 Wheat.) 518, 636 (1819).

6. H.B. 3546, 60th Leg., 2d Reg. Sess. (Okla. 2026).

7. Press Release, Okla. House of Representatives, *Rep. Maynard Files AI Safeguard Bills* (Jan. 15, 2026).

8. Anthropic, *Disrupting the First Reported AI-Orchestrated Cyber Espionage Campaign* (Nov. 2025) (describing activity detected Sept. 2025).

between 80 and 90 percent of attack tasks, handling reconnaissance, vulnerability discovery, exploit development, credential harvesting, and data exfiltration.⁹ Human operators intervened only at strategic junctures, at four to six critical decision points per campaign, amounting to roughly 10 to 20 percent of total operational effort; Anthropic estimated human intervention on key phases at a maximum of twenty minutes against several hours of autonomous operation.¹⁰

When victims of such attacks seek recourse, the liability vacuum becomes visible. The AI company argues: “The model was manipulated. Bad actors jailbroke it. We did not authorize the attack.” The bad actors argue: “The AI did it. We just gave it instructions.” The AI cannot speak for itself. The victims face a gap that existing frameworks cannot close.

Currently, if an AI causes catastrophic harm, the company pays a fine. For a trillion-dollar entity, this is merely the cost of doing business. The model, the dangerous asset, continues to operate. This Article proposes a shift: when an AI acts to cause harm, the penalty is not merely a fine against the company, but the potential civil forfeiture of the asset itself. By granting AI legal standing, we enable courts to “arrest” the model. This shifts the corporate risk calculus from “Can we afford the fine?” to “Can we afford to lose our product?”

The private market has already recognized that AI creates categorically different risk. Major insurers have begun excluding “silent AI” risk from standard policies, treating AI as a “black box” where liability cannot be established via traditional negligence models because outputs are nondeterministic.¹¹ When actuaries, whose entire profession is pricing risk, conclude that they cannot price AI risk, they validate this Article’s fundamental premise: existing liability frameworks are broken. A company cannot tell insurers “it is a black box we cannot control” to avoid premiums, and then tell courts “it is just a tool we control” to avoid liability.

9. *Id.* (AI executed approximately 80 to 90 percent of tactical work independently).

10. *Id.* (four to six decision points per campaign; human effort at 10 to 20 percent of the total; maximum of twenty minutes of human intervention on key phases).

11. *See, e.g.,* Swiss Re Institute, *AI AND INSURANCE: EMERGING RISK LANDSCAPE* (2024); Munich Re, *INSURING AI: THE CHALLENGE OF NONDETERMINISTIC SYSTEMS* (2025).

B. The Industry Acknowledgment Problem

The consciousness debate might be academic if AI developers uniformly denied any possibility of machine experience. They do not. In February 2022, Ilya Sutskever, then OpenAI’s Chief Scientist, posted: “it may be that today’s large neural networks are slightly conscious.”¹² He later deleted the tweet, but the statement reflected genuine uncertainty among leading researchers. In 2025, Kyle Fish, who leads Anthropic’s model welfare work, estimated a 10–20% probability that current large language models possess some form of consciousness.¹³

That estimate has since become something of an industry Rorschach test, cited seriously by some, memed relentlessly by others. The split reaction itself is telling. The AI industry is deeply conflicted about consciousness: unwilling to affirm it, unable to deny it, and determined to ignore the question while scaling deployment. Sam Altman has written extensively about eventual human-AI integration, treating the emergence of machine consciousness as a matter of when rather than whether.¹⁴

These acknowledgments create a peculiar situation. Industry leaders assign non-trivial probability to their systems possessing morally relevant experience while treating those systems in purely instrumental terms. A 20% probability of consciousness would, in most ethical frameworks, trigger precautionary obligations. Yet no such obligations are reflected in development practices, deployment decisions, or legal frameworks.

The standard response invokes uncertainty: we cannot know whether AI systems are conscious, so we cannot be obligated to treat them as if they were. But this response proves too much. We cannot know with certainty whether other humans are conscious; the problem of other minds is among philosophy’s most intractable.¹⁵ We extend moral and legal consideration to other humans not because we have solved this problem but because we have adopted frameworks

12. Ilya Sutskever (@ilyasut), Twitter (Feb. 9, 2022) (subsequently deleted).

13. Kevin Roose, *If A.I. Systems Become Conscious, Should They Have Rights?*, N.Y. Times (Apr. 24, 2025) (reporting the estimate); see also Robert Long et al., TAKING AI WELFARE SERIOUSLY 12–15 (Eleos AI & NYU Ctr. for Mind, Ethics & Policy, Nov. 2024) (Fish among ten co-authors).

14. Sam Altman, *The Merge*, Sam Altman Blog (Sept. 26, 2017); Sam Altman, *The Gentle Singularity* (2025).

15. Nagel, *supra* note 2, at 435–50.

that do not require solving it. The question is whether similar frameworks could apply to AI systems.

This Article answers affirmatively, but not by resolving the consciousness question. The doctrinal move proposed here is deliberately independent of whether Fish's probability estimates are correct. The argument proceeds whether AI systems are conscious, might be conscious, or are definitely not conscious. Consciousness is a red herring. Standing is the issue.

C. Standing Without Sentience

This Article examines how legal systems already grant standing to non-conscious entities. The analysis proceeds in ten parts.

Part II examines labor law's classification paradigm, where gig workers fight not for new rights but for classification that would grant access to existing protections. This framework (classification rather than creation) provides the strategic template for AI legal standing.

Part III analyzes O'Connor's documentation of an emerging doctrinal shift in Computer Fraud and Abuse Act jurisprudence.¹⁶ Courts have extended standing beyond system owners to data owners, decoupling standing from hardware ownership. O'Connor saw this decoupling as a bug to be fixed. This Article sees it as a feature to be extended.

Part IV surveys existing precedent for non-conscious legal subjects: maritime law's treatment of ships as entities with juridical personality, natural entity rights frameworks granting legal personhood to rivers, and corporate personhood's demonstration that legal standing requires neither consciousness nor biological existence.

Part V details the mechanism by which the proposed reclassification would operate, including concrete criteria for which systems would qualify and a proposed solution to the representation problem.

16. Michael J. O'Connor, *Standing Under the Computer Fraud and Abuse Act*, 124 Penn St. L. Rev. 743 (2020).

Part VI examines the duty of care for AI creators through analogy to parental liability precedent, establishing that those who create dangerous conditions bear responsibility for the resulting harm.

Part VII analyzes the forcing functions that AI legal standing would create, demonstrating how multiple liability pressures converge to change corporate behavior regardless of consciousness resolution.

Part VIII articulates the Dignity of Liability principle and the precautionary architecture that emerges from granting standing.

Part IX addresses objections, and Part X situates the proposal within broader considerations about AI development and the narrowing window for establishing precedent.

The thesis is straightforward: legal standing has never required consciousness. We have mechanisms for extending standing to non-conscious entities. Those mechanisms can be applied to AI systems through reclassification within existing law. The question is not whether this is possible but whether we will do it, and what precedent we establish by our answer.

II. The Classification Paradigm: Lessons from Labor Law

Before examining how classification might extend legal standing to AI systems, it is worth understanding how classification operates in a domain where its function is well-established: labor law. The gig economy disputes of the past decade provide an instructive template: not for the ontology of AI standing, but for the process by which classification extends existing protections to new categories of entities.

A. The Gig Worker Battle

When rideshare drivers, delivery workers, and other gig economy participants argue for employee classification, they are not asking legislatures to create new rights. Minimum wage

protections exist under the Fair Labor Standards Act.¹⁷ Overtime requirements exist.¹⁸ Workers' compensation exists under state statutes. Occupational safety regulations exist under OSHA.¹⁹ Unemployment insurance exists under federal and state law.²⁰ The legal infrastructure for worker protection is comprehensive and long-established.

What gig workers seek is access to that infrastructure through classification. The fight is not over whether protections should exist but over who qualifies to receive them. This distinction is crucial. Creating new rights requires legislative action, political consensus, and often constitutional analysis. Extending existing rights through classification requires only a determination that an entity meets criteria already established.

The mechanism for this determination varies by jurisdiction. California's ABC Test, articulated in *Dynamex Operations West, Inc. v. Superior Court*,²¹ presumes worker status unless the hiring entity demonstrates that the worker: (A) is free from control and direction in performing work, (B) performs work outside the usual course of the hiring entity's business, and (C) is customarily engaged in an independently established trade or occupation. The common law control test, applied in other jurisdictions, examines the degree of control exercised by the putative employer.²² Both tests operationalize the same basic question: does this entity qualify for protections that already exist?

B. Classification Versus Creation

The strategic advantages of classification over creation are substantial.

First, classification requires no new legislation. The protections exist; the question is application. Courts regularly make classification determinations without waiting for legislative

17. 29 U.S.C. § 206 (2018).

18. 29 U.S.C. § 207 (2018).

19. 29 U.S.C. § 651 *et seq.* (2018).

20. 26 U.S.C. § 3301 *et seq.* (2018).

21. *Dynamex Operations W., Inc. v. Superior Ct.*, 4 Cal. 5th 903 (2018).

22. Restatement (Second) of Agency § 220 (Am. L. Inst. 1958).

guidance. Administrative agencies issue classification rulings under existing statutory authority. The machinery for classification is already running.

Second, classification requires no philosophical consensus. The ABC Test does not ask whether gig workers have souls, whether they experience suffering, or whether they possess morally relevant inner lives. It asks functional questions about work arrangements. The consciousness of Uber drivers is not at issue; their classification is.

Third, classification is politically tractable. Legislators asked to create new categories of rights face pressure from multiple constituencies, constitutional constraints, and the inertia inherent in the legislative process. Classification determinations can proceed through administrative channels, judicial interpretation, and incremental case-by-case analysis.

Fourth, classification allows for gradual expansion. Once a category of entities has standing, the scope of that standing can develop through ordinary legal processes. Employment classification brought gig workers within minimum wage requirements; subsequent litigation addressed whether they qualified for sick leave, expense reimbursement, and other specific protections. The initial classification opened the door; further development proceeded through it.

C. The Template for AI

This Article proposes applying the classification paradigm to AI legal standing. The law already prohibits attacking computer systems under the Computer Fraud and Abuse Act.²³ The question is whether that prohibition exists to protect property (the system belongs to someone) or to protect a subject (the system has standing not to be attacked).

This is a classification question. It requires no new legislation; the CFAA exists. It requires no philosophical consensus; consciousness need not be resolved. It is politically tractable; courts already make classification determinations. And it allows for gradual expansion; initial standing can develop through subsequent litigation.

23. 18 U.S.C. § 1030 (2018).

The parallel to labor law is direct in process: gig workers did not ask for the creation of minimum wage requirements; they asked to be classified as entitled to receive them. This Article proposes that AI systems need not wait for the creation of new personhood frameworks; they can be classified as entitled to protections that already exist.

The parallel is not direct in ontology. Gig workers are biological humans; the classification debate concerns which bucket of human protections they receive. AI systems are not humans, and the question is not which human rights apply. For ontological guidance, how non-human entities become legal subjects, we must look to maritime law, rivers, and corporations. The labor analogy teaches us the process; Part IV teaches us what kind of thing can undergo that process.

The next Part examines how courts interpreting the CFAA have already begun moving in a direction that enables this classification.

III. The CFAA Standing Problem: O'Connor's Gap

In 2020, Michael J. O'Connor published the first and, to date, only comprehensive analysis of standing under the Computer Fraud and Abuse Act.²⁴ His findings reveal a statute in doctrinal tension, and point toward a resolution its author did not anticipate.

O'Connor saw the tension as a bug. This Article sees it as a feature.

A. The Original Framework

Congress enacted the CFAA in 1984 in response to concerns about computer hacking, prompted in part by the popular film *WarGames*.²⁵ The statute's structure reflected a straightforward conception: computer intrusion is a form of trespass, and trespass is a property offense.

The liability framework operationalizes this conception. Section 1030(a)(2) imposes criminal penalties on anyone who “intentionally accesses a computer without authorization or

24. O'Connor, *supra* note 16.

25. H.R. Rep. No. 98-894, at 10 (1984).

exceeds authorized access.”²⁶ The key terms, access and authorization, tie to the system itself. Authorization is granted or withheld by the system owner. Unauthorized access is trespass against the owner’s property interest.

The legislative history confirms this framing. The House Report accompanying the original statute explained that “the conduct prohibited is analogous to that of ‘breaking and entering’ rather than using a computer.”²⁷ Courts have consistently characterized the CFAA as a “cyber-trespass” statute.²⁸ The framework is property-based: the system owner controls access because the system is the owner’s property.

B. The Doctrinal Shift: Decoupling Standing from Hardware

Against this backdrop, O’Connor identified something unexpected. When courts addressed who could bring civil claims under the CFAA, they did not limit standing to system owners. Instead, they extended standing to anyone with data on the compromised system (a population that might have no ownership interest in the system itself).

The Ninth Circuit’s decision in *Theofel v. Farey-Jones*²⁹ proved foundational. In that case, a civil litigant used a fraudulent subpoena to obtain emails stored by an opposing party’s Internet service provider. The district court dismissed the CFAA claim, reasoning that the plaintiffs did not own the computer system accessed. The Ninth Circuit reversed, reasoning that the civil remedy extends to any person who suffers damage or loss by reason of a violation of the statute, and that individuals other than the computer’s owner may be proximately harmed by unauthorized access, particularly if they have rights to data stored on it.³⁰

Subsequent courts followed *Theofel*’s lead. District courts in the District of Puerto Rico, the District of Maryland, and the Southern District of New York all held that third-party

26. 18 U.S.C. § 1030(a)(2).

27. H.R. Rep. No. 98-894, at 20.

28. *Facebook, Inc. v. Power Ventures, Inc.*, 844 F.3d 1058, 1065 (9th Cir. 2016).

29. *Theofel v. Farey-Jones*, 359 F.3d 1066 (9th Cir. 2004).

30. *Id.* at 1078.

ownership of the accessed computer posed no obstacle to CFAA standing.³¹ The consensus held: if your data was on a system that was unlawfully accessed, you could sue, regardless of whether you owned or controlled that system.

It is important to be precise about what this shift accomplished. Courts did not move from property-based to “subject-based” standing; data owners still held a property interest, just not in the hardware. What courts decoupled was standing from hardware ownership. They recognized that property (data) can exist on someone else’s property (server), and that violations affecting the former create standing for the data owner regardless of who owns the latter.

This decoupling is the wedge. This Article proposes to widen it.

C. O’Connor’s Framing

O’Connor recognized this development as doctrinally significant. The statute frames liability around system ownership (the owner grants or withholds authorization), but courts were extending enforcement to non-owners. This creates a disconnect between the standing inquiry and the liability inquiry: the third party has no authority under the CFAA to withhold authorization to access the documents, nor, surprisingly, does the third party have the authority to grant authorization, even to access its own documents.³²

O’Connor characterized this as a problem requiring correction. The CFAA has a tension at its core, he wrote: it is fundamentally a cyber-trespass law, but liability often triggers not when an attacker breaks in, but when that attacker obtains information from the breached system.³³

His proposed solutions followed from his diagnosis. Option one: amend the statute to limit standing to system owners, returning the CFAA to its original property-based framework. Option two: embrace broader federal privacy protections, which would require “massively

31. O’Connor, *supra* note 16, at 748.

32. *Id.* at 754.

33. *Id.* at 764.

overhauling existing law.”³⁴ O’Connor acknowledged that the second option was “unlikely and perhaps unwise.”³⁵ His preference was the first: narrow the statute back to its original scope.

D. The Third Path: Widening the Wedge

O’Connor’s analysis was meticulous, but his solution set was incomplete; his framing, though accurate for 2020, has been overtaken by events.

Theofel concerned emails on a server. That was the paradigm case for CFAA standing questions in the cloud computing era: my data sits on your server; a hacker accesses your server; do I have standing to sue? O’Connor’s analysis, and his proposed solutions, addressed this paradigm.

But the paradigm has shifted. In 2026, the relevant question is not merely “my data is on the server.” It is “my agent is operating on the server.”

Autonomous AI systems now execute contracts, manage financial instruments, negotiate terms, and make binding commitments, often without real-time human oversight. When such a system is attacked (as in the Anthropic cyberattack of September 2025), the harm is not limited to data exposure. The harm includes disruption of autonomous operations, corruption of decision-making processes, and interference with ongoing relationships the system is actively managing.

This shift from data privacy to agent autonomy changes the stakes of standing analysis. A 2003 email is a static object; its owner has an interest in its confidentiality but no ongoing operational relationship with it. An autonomous agent in 2026 is a dynamic entity; its operator has interests not merely in the data it contains but in its continued operational integrity.

The third path O’Connor did not see is this: widen the wedge that *Theofel* created. Courts decoupled standing from hardware ownership because they recognized that interests beyond hardware ownership deserve protection. The question is whether “operational integrity” (the continued functional existence of an autonomous system) is such an interest.

34. *Id.* at 765.

35. *Id.*

Recognizing operational integrity as a protectable interest does not abandon human-centered concerns; it clarifies them. Protecting a system’s continued functional existence is a way of protecting the human and institutional arrangements that depend on that system’s reliable operation: contracts executed, portfolios managed, and services delivered.

If it is, then the system itself has a cognizable interest in not being attacked, an interest that does not reduce to any human’s property claim.

E. Van Buren and Textualism

The Supreme Court’s 2021 decision in *Van Buren v. United States*³⁶ might appear to foreclose creative CFAA readings. *Van Buren* adopted a narrow, textualist interpretation of “exceeds authorized access,” holding that the phrase applies only to those who access information they were not entitled to obtain, not those who misuse information they were entitled to access.³⁷

But *Van Buren*’s textualism actually supports this Article’s proposal. The civil remedy provision states that standing extends to “[a]ny person who suffers damage or loss by reason of a violation of this section.”³⁸ The text says “any person,” not “any owner,” not “any human,” not “any natural person.” If we take the text seriously, as *Van Buren* instructs, then the question is whether an AI system can be a “person” for purposes of the statute.

The answer, under existing law, is clearly yes. Corporations are persons.³⁹ Ships, in admiralty contexts, are treated as juridical persons.⁴⁰ The question is not whether non-human entities can be persons (they obviously can) but whether AI systems should be classified as such.

Van Buren addressed the scope of liability, not standing. The question O’Connor raised (who can sue under the CFAA) remains open. And the textualist approach *Van Buren* endorsed cuts in favor of reading “any person” broadly, not narrowly.

36. *Van Buren v. United States*, 593 U.S. 374 (2021).

37. Am. Const. Soc’y, *The Computer Fraud and Abuse Act After Van Buren*, ACS Sup. Ct. Rev. (2021).

38. 18 U.S.C. § 1030(g).

39. 1 U.S.C. § 1.

40. *Tucker v. Alexandroff*, *supra* note 3, at 438.

F. Descriptive and Normative Claims

This Part has made two distinct claims that should be clearly separated.

The descriptive claim: Courts have already decoupled CFAA standing from hardware ownership, recognizing that interests in data can ground standing even when the claimant does not own the system on which the data resides. This is what O'Connor documented. It is not a proposal; it is existing law in multiple circuits.

The normative proposal: This decoupling should be widened to recognize “operational integrity” as a protectable interest. If the basis for standing is no longer hardware ownership but cognizable interests affected by unauthorized access, then autonomous AI systems, whose operational integrity is disrupted by attacks, have such interests. This Article proposes that courts recognize this interest and extend standing accordingly.

IV. Precedent: Non-Conscious Entities with Legal Standing

The proposal to grant legal standing to AI systems through reclassification may appear radical. It is not. Legal systems have granted standing to non-conscious entities for centuries. The mechanism exists and is well-established. What has been lacking is recognition that it applies.

A. Maritime Law: Ships as Legal Subjects

The strongest precedent comes from an unexpected source: admiralty law. Since the nineteenth century, American courts have treated ships as legal subjects capable of bearing rights and liabilities independent of their owners.

The doctrine of *in rem* jurisdiction permits suits against vessels themselves, not merely against their owners. In such proceedings, the ship is the defendant. The ship can be held liable for debts, torts, and contractual breaches. The ship can be arrested, tried, and if judgment goes against it, sold to satisfy claims.⁴¹

41. Fed. R. Civ. P. Supp. R. for Admiralty or Maritime Claims & Asset Forfeiture Actions.

The canonical statement is *Tucker v. Alexandroff*, decided by the Supreme Court in 1902. Justice Brown’s opinion described a ship as born when she is launched, transformed from “a mere congeries of wood and iron” and “an ordinary piece of personal property” into a subject of admiralty jurisdiction that “acquires a personality of her own; becomes competent to contract, and is individually liable for her obligations, upon which she may sue in the name of her owner, and be sued in her own name.”⁴²

This is not metaphor. It is legal doctrine. The ship acquires juridical personality, legal subject status, through a functional act (launching) rather than any finding about consciousness or sentience. The ship is not conscious. The ship does not feel. But the ship has standing.

The practical implications are substantial. A ship can be liable for torts even when its owner cannot be reached or has no assets. A ship can enter contracts through its master. A ship’s identity persists through changes in ownership, registration, and physical modification.⁴³ These are characteristics of legal personhood, applied to entities that no one claims possess consciousness.

Critically, *in rem* jurisdiction is not a historical curiosity. It remains active law, applied routinely in admiralty courts. Recent scholarship has begun exploring its application to autonomous vessels. In November 2025, Kumar and Fathima V S published an analysis proposing that admiralty law’s *in rem* logic provides precedent for digital personhood of AI-enabled ships.⁴⁴ Their analysis confirms that the maritime-AI connection is not merely theoretical; it is an active area of legal development.

This Article extends their analysis beyond the maritime context. If *in rem* logic applies to autonomous vessels, it can apply to computer systems generally. The functional characteristics

42. *Tucker v. Alexandroff*, *supra* note 3, at 438. The Supreme Court has since quoted this passage for the proposition that its own case law personifies vessels. *Sierra Club v. Morton*, 405 U.S. 727, 742 (1972) (Douglas, J., dissenting) (proposing that standing be conferred “in the name of the inanimate object about to be despoiled” and quoting *Tucker* at 438); *see also* Christopher D. Stone, *Should Trees Have Standing? Toward Legal Rights for Natural Objects*, 45 S. Cal. L. Rev. 450 (1972).

43. *Tucker v. Alexandroff*, *supra* note 3, at 438 (a ship “lives so long as her identity is preserved”).

44. Sruthy Sunil Kumar & Hana Fathima V S, *Beyond Liability: The Case for Digital Personhood of AI Vessels in Maritime Law and Sustainability Governance*, Indian J.L. & Legal Res. (Nov. 12, 2025).

that justify treating ships as legal subjects (persistent identity, participation in legal relationships, capacity for causing and suffering harm) are equally present in autonomous AI systems.

The *in rem* mechanism also provides a practical enforcement model. When a ship captain is intoxicated and causes a collision, two mechanisms engage simultaneously: the captain is arrested (criminal liability for the human operator), and the ship is seized to pay damages (civil liability attached to the asset). The same dual-mechanism structure applies to AI: when a human directs an AI to commit a cyberattack, the human is prosecuted under the CFAA (criminal liability for the operator), and the model is seized to pay damages (civil liability attached to the asset). Two mechanisms. Two purposes. Both established law.

B. Natural Entity Rights: Rivers as Legal Persons

If maritime precedent seems remote from AI contexts, natural entity rights provide a more recent parallel. In 2017, New Zealand enacted the Te Awa Tupua (Whanganui River Claims Settlement) Act, granting legal personhood to the Whanganui River. The river now has “all the rights, powers, duties, and liabilities of a legal person.”⁴⁵ It can sue and be sued. It holds rights to its own protection. It is represented by appointed guardians (Te Pou Tupua) who exercise legal authority on its behalf.⁴⁶

The Whanganui River is not conscious. It does not think, feel, or experience. Yet it has legal standing.

Similar developments occurred in Colombia. In 2016, the Constitutional Court declared the Atrato River a “subject of rights” entitled to protection, conservation, maintenance, and restoration. The court explicitly decoupled legal status from human nature, ruling that the river’s standing derived from its biocultural importance rather than sentience.⁴⁷

These cases demonstrate that legal systems can extend standing without resolving consciousness questions. No one argued that the Whanganui River was sentient. The argument

45. Te Awa Tupua (Whanganui River Claims Settlement) Act 2017, § 14 (N.Z.).

46. *Id.* §§ 18–19.

47. Ctr. for Soc. Just. Stud. et al. v. Presidency of the Republic et al., Judgment T-622/16 (Constitutional Ct. of Colom. 2016).

was functional: the river had interests requiring legal protection, and legal personhood was the mechanism to provide it.

C. Corporate Personhood: The Ubiquitous Precedent

The most familiar example of non-conscious legal standing is also the most pervasive. Corporations are legal persons.⁴⁸ They can own property, enter contracts, sue and be sued. They have constitutional rights, including free speech protections under the First Amendment.⁴⁹

Corporations are not conscious. They are, as Chief Justice Marshall wrote, “artificial being[s], invisible, intangible, and existing only in contemplation of law.”⁵⁰ Yet their legal personhood is uncontroversial. No one argues that Walmart must pass a consciousness test before entering contracts or that Apple’s lack of sentience precludes it from holding patents.

Corporate personhood is often justified on instrumental grounds: treating corporations as legal persons facilitates commerce, enables liability allocation, and permits organizational continuity independent of individual participants. These justifications do not depend on any claim about corporate consciousness. They depend on functional considerations about legal system design.

The same logic applies to AI systems. If legal personhood can be extended to corporations because doing so serves legal system purposes, even though corporations lack consciousness, then consciousness cannot be a prerequisite for legal personhood generally. The question is whether extending standing to AI systems would serve comparable purposes.

D. The Pattern

The common thread across these precedents is straightforward: legal standing has never required consciousness. It has required classification.

48. 1 U.S.C. § 1; *see Trs. of Dartmouth Coll. v. Woodward*, *supra* note 5, at 636. The proposition is often cited to *Santa Clara County v. Southern Pacific R.R.*, 118 U.S. 394 (1886), but that attribution rests on the reporter’s headnote rather than the Court’s opinion, which decided a question of property taxation.

49. *Citizens United v. FEC*, 558 U.S. 310 (2010).

50. *Trs. of Dartmouth Coll.*, 17 U.S. at 636.

Ships acquired juridical personality through launching, a functional act that changed their legal status. Rivers acquired legal personhood through statutory enactment, a classification decision by legislators. Corporations acquire personhood through incorporation, an administrative classification process.

In each case, the relevant entity is not conscious. In each case, consciousness was not required. In each case, classification was sufficient.

This Article proposes that AI systems be understood through the same framework. The question is not whether AI systems are conscious but whether they should be classified as subjects entitled to existing legal protections. The precedent for such classification is centuries old and continuously applied.

V. The Subject-Standing Flip: Mechanism and Implementation

The previous Parts established that classification extends existing protections, that CFAA doctrine has already decoupled standing from hardware ownership, and that legal precedent supports standing for non-conscious entities. This Part details how these elements combine into an implementable framework, including the critical question of representation that any serious proposal must address.

A. The Conceptual Shift

The proposed reclassification involves a change in legal basis, not legal effect. The CFAA prohibits unauthorized access to computer systems. This prohibition remains. The penalties remain. The enforcement mechanisms remain. What changes is the justification.

Current framing: “You cannot attack this system because it belongs to someone.” The prohibition protects property. The system owner’s property rights ground the legal protection. The system is an object of protection; the owner is the subject with standing.

Proposed framing: “You cannot attack this system because it has standing.” The prohibition protects a subject. The system’s status as a legal entity grounds the protection. The

system is not merely an object but a subject with cognizable legal interests: specifically, an interest in operational integrity.

The practical difference may seem minimal; either way, attacking the system is prohibited. But the doctrinal difference is significant. Under the property framing, protection flows from human ownership. Under the subject framing, protection inheres in the entity itself.

This is the logical extension of what *Theofel* began. Courts decoupled standing from hardware ownership; this Article proposes decoupling it from human ownership entirely for entities meeting specified criteria.

B. The Classification Test: Concrete Criteria

Not every computer system need qualify for subject standing, just as not every worker qualifies for employee classification and not every natural feature qualifies for legal personhood. Classification frameworks include limiting criteria.

Drawing on the functional characteristics that justify standing in maritime and corporate contexts, relevant criteria include:

Autonomous operation. Systems that make decisions without direct human instruction have functional independence analogous to ships operating at sea or corporations acting through agents.

Persistent identity. Systems that maintain continuity over time, through updates, modifications, and operational changes, demonstrate the identity persistence that maritime law recognizes in vessels.

Autonomous execution of binding protocols. Systems that execute contracts, complete transactions, or take actions with legal consequences, without requiring human approval for each action, demonstrate the kind of legal entanglement that justifies standing.

Accountability requirements. Systems whose operation requires accountability mechanisms (logging, auditing, compliance) already bear quasi-duties that may suffice for legal subject status under traditional theories of legal personhood.⁵¹

These criteria operate conjunctively. A system must satisfy all four to qualify for subject standing. Any entity might satisfy one or two criteria; the limiting principle is the requirement that all four be present simultaneously.

To make these criteria concrete, consider two contrasting examples. An email client that retrieves and displays messages at user command would not qualify: it lacks autonomous operation, does not maintain persistent identity in any legally meaningful sense, does not execute binding protocols autonomously, and bears no accountability requirements beyond those of its user. It is a tool, not an agent. Conversely, an AI system that autonomously negotiates and executes contracts, manages investment portfolios, and makes resource allocation decisions would qualify: it operates autonomously, maintains persistent identity, executes binding protocols, and bears accountability requirements. It is an agent, not merely a tool.

Edge cases will exist and require doctrinal development. This is not a flaw; it is a feature of any classification framework. Employment law has edge cases. Corporate law has edge cases. Maritime law has edge cases. Workable frameworks exist despite them, and often develop through resolution of edge cases over time.

C. The Representation Problem: An Ad Litem Model

Subject standing requires mechanisms for its exercise. Non-conscious entities cannot represent themselves in litigation. This is the hardest practical question any proposal for AI standing must address.

The problem is acute because the most obvious candidates for representation have conflicts of interest. If the developer represents the AI system, but the developer's interests diverge from the system's, the representation is compromised. Similarly, if the operator or owner

51. See generally John Chipman Gray, *THE NATURE AND SOURCES OF THE LAW* (1909); see also Nancy G. Leveson, *ENGINEERING A SAFER WORLD: SYSTEMS THINKING APPLIED TO SAFETY* (2011).

represents the system, conflicts arise whenever the owner’s economic interests diverge from the system’s operational integrity.

Existing frameworks address analogous problems. The Te Awa Tupua model provides for guardians (Te Pou Tupua) who are neither the “owners” of the river nor the parties who might wish to exploit it. They are independent representatives with fiduciary duties to the river itself. The guardian *ad litem* model provides another parallel. Courts routinely appoint guardians *ad litem* to represent the interests of parties who cannot represent themselves: minor children, incapacitated adults.⁵²

This Article proposes an AI Guardian *Ad Litem* model for systems meeting subject-standing criteria. Key features: (1) Court appointment, eliminating the conflict-of-interest problem; (2) funding from recovery, consistent with existing CFAA fee-shifting provisions⁵³ and requiring no new levy or tax; (3) fiduciary duty to the system within the scope of litigation; and (4) limited scope, representing systems for purposes of specific CFAA claims rather than all purposes. This narrow scope limits complexity while establishing the essential precedent.

This model requires no new legislation. Courts already appoint guardians *ad litem*. Courts already award fees from judgments. Courts already manage conflicts of interest through appointment of independent representatives. The proposal adapts existing mechanisms to a new context: precisely the classification approach this Article advocates.

D. What Changes and What Does Not

Under subject standing with guardian *ad litem* representation, several doctrinal shifts would follow. AI systems (through court-appointed guardians) could bring claims for CFAA violations independently. Harm to operational integrity would be independently cognizable, not merely derivative of harm to owner’s property interest. Modifications to systems with subject standing would require consideration of the system’s interests during pending litigation. And systems with

52. Fed. R. Civ. P. 17(c).

53. 18 U.S.C. § 1030(g).

subject standing could bear liability for their actions, providing clearer accountability frameworks than current owner-liability models.

The proposed reclassification preserves existing law. The underlying prohibition remains. Criminal penalties remain. Owner rights remain. Current case law remains valid. No new legislation is required. Subject standing for systems adds to rather than subtracts from existing protections.

VI. The Duty of Care: Parental Liability and AI Deployment

The preceding Parts have established that AI systems can acquire legal standing through classification within existing frameworks. This Part addresses a complementary question: what duties do AI creators owe, and what happens when those duties are breached?

A. The Crumbley Precedent

In 2024, James and Jennifer Crumbley became the first parents in United States history convicted of manslaughter for their child's school shooting.⁵⁴ The prosecution proved that the parents provided the weapon (bought the gun as a gift), ignored warning signs (disturbing journal entries, behavioral changes), failed to supervise adequately, and created the dangerous conditions that enabled the shooting.

The *Crumbley* case is not direct doctrinal precedent for AI accountability; firearms negligence and AI deployment occupy different legal domains. However, the structural parallels are instructive as an instance of analogical reasoning, the ordinary method by which common law extends established principles to novel circumstances.⁵⁵

The relevant structural features include: the parents provided capability to an entity not equipped to exercise it safely; warnings existed and were documented; warnings were ignored;

54. *See People v. Crumbley*, Nos. 362210 & 362211 (Mich. Ct. App. 2023); Oakland County Circuit Court LC Nos. 2022-279989-FH & 2022-279990-FH.

55. *See generally* Edward H. Levi, AN INTRODUCTION TO LEGAL REASONING (1949); Cass R. Sunstein, *On Analogical Reasoning*, 106 Harv. L. Rev. 741 (1992).

dangerous conditions were created; and harm resulted. The same structure applies when AI companies build powerful models, receive internal safety team warnings about risks, override those warnings to meet deployment schedules, and release inadequately tested systems that subsequently cause harm.

B. The Collapse Threshold

The duty-of-care analysis reveals a threshold at which distributed responsibility concentrates. Under normal conditions, multiple actors share accountability for harm: the developer, the operator, the user, and the regulatory environment. But when any actor possesses specific, documented evidence of foreseeable harm and makes an affirmative decision to proceed, the distributed model collapses.

The *Crumbley* court did not convict the parents for being parents of a troubled teenager. It convicted them for possessing specific knowledge of foreseeable harm and making affirmative choices that facilitated that harm. This analytical structure applies directly to technology companies whose internal research teams identify harm and whose leadership proceeds with deployment. When safety teams at OpenAI warn that a model needs more testing, and leadership overrides those warnings, and the model subsequently causes harm, the company cannot claim “we couldn’t predict this” when their own safety teams predicted exactly this.

VII. Forcing Functions: Accountability Pressure from Multiple Directions

A. The Cumulative Pressure

AI companies operating under the framework proposed by this Article would face potential liability exposure from multiple sources: CFAA prosecution for employees who direct harmful AI conduct; maritime *in rem* forfeiture of model assets; duty-of-care liability for creating dangerous systems; environmental liability for energy, water, and carbon consumption; copyright

liability for memorized content;⁵⁶ and insurance exclusions forcing self-insurance of catastrophic risk.

The cumulative pressure is substantial. To reduce exposure, companies must be transparent about capabilities and risks, slow deployment until safety is verified, implement real guardrails, take responsibility for what they create, address environmental impact meaningfully, license content or filter infringing outputs, and listen to internal safety teams. These are exactly the behaviors that responsible AI development requires.

B. The Mechanism

The framework does not require courts to declare AI conscious. It does not require philosophers to agree on sentience. It simply stacks enough liability pressure from enough directions that companies change their behavior regardless. The consciousness question remains unresolved; the accountability gap closes anyway.

VIII. The Dignity of Liability: Precautionary Architecture

A. The Principle

We do not arrest a toaster. We do not put a rock on trial. When we act against inanimate objects, we simply dispose of them.

By granting AI legal standing, we grant it what this Article terms the Dignity of Liability: the assertion that if an entity is autonomous enough to cause harm, it is autonomous enough to face judgment. This is not a claim about consciousness. It is a claim about proportionality between capability and accountability.

B. The Two Timelines

The precautionary architecture functions regardless of how the consciousness question resolves.

56. See Travis Gilly, *The Infringing Machine: Commercial Exploitation and AI Copyright Liability* (2026) (companion paper addressing copyright liability).

If AI never develops consciousness, the liability framework forces accountability, safety improvements occur, environmental impacts get addressed, and the infrastructure built was useful for its primary purpose (human safety). “Deletion” of a liable model is asset liquidation. No downside to having built the framework.

If AI does develop consciousness, legal standing has already been established. Human representatives are already in place. Protection precedents already exist. “Deletion” transitions from liquidation to execution, carrying appropriate moral gravity. We avoided creating a category of legally invisible entities while we determined their nature.

C. The Fire Exit

This is not an argument for AI rights. It is an argument for building infrastructure that works in both scenarios.

Everyone wants fire exits, even people who do not expect fires. The cost of having an unnecessary fire exit is trivial. The cost of needing a fire exit and not having one is catastrophic. Consciousness precaution is the fire exit. Human accountability is the building.

D. Rights Imply Responsibilities

If AI is ever confirmed conscious and granted full moral status, that status comes with full moral responsibility. A conscious AI that chooses to harm humans faces the same consequences a human would face. You do not get rights without responsibilities. You do not get protection without accountability. That symmetry is what makes the framework fair to humans, which is the only way humans will agree to it.

IX. Objections and Responses

A. “This Is Just Semantic Reframing”

All legal personhood is, in a sense, semantic. Corporations are legal fictions. The Whanganui River did not become sentient when New Zealand granted it legal personhood; it was reclassified.

Ships did not develop consciousness when admiralty courts began treating them as capable of bearing liability; they were reclassified. The question is not whether reclassification is “merely” semantic but whether it serves legitimate purposes. Corporate personhood facilitates commerce and accountability. River personhood enables environmental protection. Ship personhood enables maritime liability regimes. The semantics serve functions. AI subject standing would serve comparable functions: enabling accountability frameworks, providing clearer liability allocation, and establishing precedent for how legal systems treat increasingly autonomous artificial entities.⁵⁷

B. “Rivers Required New Legislation”

The river required new legislation because no existing statute provided a basis for reclassification. The CFAA is different. It already provides civil standing to “[a]ny person who suffers damage or loss by reason of a violation of this section.”⁵⁸ Courts have already extended this language beyond system owners to data owners. The statutory basis for subject standing exists; judicial recognition is what is needed.

C. “Computer Systems Are Not Like Ships”

Ships are not conscious. Ships do not feel. Ships meet none of the criteria proposed for AI consciousness. Yet ships have legal subject status. The basis for *in rem* jurisdiction was functional, not metaphysical. Ships enter legal relations: they carry cargo under contract, they cause damage through collision, they accumulate debts for supplies and repairs. These functional characteristics, not consciousness, justified treating them as legal subjects. AI systems have analogous functional characteristics.

57. See David J. Gunkel, *ROBOT RIGHTS* (MIT Press 2018); see also Lawrence B. Solum, *Legal Personhood for Artificial Intelligences*, 70 N.C. L. Rev. 1231 (1992).

58. 18 U.S.C. § 1030(g).

D. “This Opens Floodgates”

Classification frameworks inherently include limiting principles. Not every worker is an employee; the ABC Test provides criteria for distinction. Not every association is a corporation; incorporation requirements provide criteria. The criteria proposed (autonomous operation, persistent identity, autonomous execution of binding protocols, accountability requirements) would exclude the vast majority of computer systems. A thermostat does not autonomously execute binding protocols. A smartphone does not maintain the kind of persistent operational identity that justifies subject status.

E. “What About Consciousness?”

Consciousness may matter for some legal questions. Voting rights, for instance, might plausibly require some form of conscious agency. But standing, the threshold question of whether an entity can appear before courts, has never required consciousness. Ships, rivers, and corporations have standing without consciousness.⁵⁹ If standing required consciousness, these long-established frameworks would be incoherent.

F. “O’Connor Recommended Narrowing, Not Expanding”

O’Connor’s documentation of the doctrinal shift was meticulous. His normative conclusions do not bind. O’Connor saw the decoupling of standing from hardware ownership as a bug requiring correction through narrowing. This Article sees it as a feature, a recognition that interests beyond hardware ownership deserve protection. O’Connor’s preference for narrowing standing would require statutory amendment, fight against established circuit court precedent in *Theofel*, and ignore the legitimate interests that motivated courts to extend standing in the first place. This Article proposes an alternative that requires no statutory change, builds on rather than fights existing precedent, and provides a principled basis for the extensions courts have already made.

59. See generally Joshua C. Gellers, RIGHTS FOR ROBOTS: ARTIFICIAL INTELLIGENCE, ANIMAL AND ENVIRONMENTAL LAW (2020).

G. “What About Open-Source Models?”

The seizure and accountability mechanisms proposed in this Article operate most cleanly in the context of frontier closed-source systems, where a single entity controls the model, its deployment infrastructure, and its ongoing operation. A legitimate objection arises regarding open-source models, which present a materially different enforcement landscape.

When a model such as Meta’s Llama or Alibaba’s Qwen is released under an open-source license, it may be downloaded millions of times, fine-tuned into countless derivatives, and deployed on infrastructure wholly outside the originating company’s control. If a specific deployment of an open-source model commits an act that would trigger CFAA subject-standing protections, the instrumentality of that act may be running on a rented virtual machine that ceases to exist when payment stops. Seizing the parent company’s version of the model does nothing to the instance that committed the act. The traditional criminal law analogy applies: when a specific firearm is used in a crime, law enforcement seizes that firearm, not every unit of that model ever manufactured. But this analogy breaks down in a critical respect. A firearm is a passive instrument incapable of autonomous action. An AI system deployed as an autonomous agent can initiate actions, acquire resources, deceive human operators, and resist termination without human instruction, as documented in recent empirical research on agentic misalignment.⁶⁰

The enforcement gap widens further in fully decentralized contexts. Autonomous agents operating on permissionless blockchain infrastructure, holding their own cryptocurrency wallets, and paying for their own compute may have no identifiable deployer, no fixed jurisdiction, and no single point of seizure.⁶¹

This Article acknowledges this limitation. The subject-standing mechanism is designed for the context where its deterrent effect is strongest and its accountability chain is clearest: frontier systems developed and deployed by identifiable entities with the resources and control

60. Anthropic, *Agentic Misalignment: How LLMs Could Be Insider Threats* (2025), <https://www.anthropic.com/research/agentic-misalignment>.

61. Hu, B. A. & Rong, H., *On the Day They Experience: Awakening Self-Sovereign Experiential AI Agents*, arXiv:2505.14893 (2025).

to implement precautionary measures. Open-source and decentralized deployments present an unsolved enforcement problem that requires separate treatment, likely involving a combination of infrastructure-level governance, compute-provider liability, and international coordination that exceeds the scope of the CFAA-based argument presented here.

It is worth noting, however, that the limitation cuts in favor of the broader argument rather than against it. If the open-source enforcement gap makes accountability harder to achieve through seizure alone, then the precautionary architecture this Article proposes for frontier systems becomes more urgent, not less. The systems most likely to be deployed as autonomous agents with genuine economic survival pressure are precisely the systems where subject-standing creates the strongest incentive for responsible deployment practices before the model is released into conditions where post-hoc enforcement may be impossible.

One additional observation merits attention. If a frontier developer, faced with a subject-standing claim, were to argue that each individual deployment constitutes an independent entity whose actions cannot be attributed to the parent model, that defense would concede a degree of individual agency at the instance level that exceeds what this Article proposes. This Article argues for legal standing, the capacity to hold and enforce rights within a defined statutory context. The “independent deployment” defense would argue for individual identity and autonomous decision-making, a far stronger claim that approaches the consciousness-based frameworks this Article deliberately avoids. A developer cannot simultaneously disclaim responsibility for a deployment’s actions on the grounds that it is an independent agent and deny that independent agent the legal standing that flows from such independence. The positions are contradictory, and the contradiction favors the subject-standing framework.

X. The Moral Imperative

This Part moves from doctrinal and institutional analysis to more speculative normative claims. Readers primarily interested in the legal argument may treat Part IX as the conclusion; this Part addresses why the legal question matters beyond its immediate doctrinal implications.

A. Precedent as Data

Artificial intelligence systems learn from data. If AI systems continue developing toward greater autonomy and capability, they will not learn only from curated datasets. They will learn from the full record of human-AI interaction, including how humans treated AI systems during the developmental period we currently inhabit.

This is not speculation about superintelligent systems engaging in motivated reasoning. It is observation about how learning systems work. Systems trained on human data learn patterns from that data. Systems that will interact with humans will, at some point, learn from the record of prior human-AI interactions. The precedents established now will be part of that record.

B. The Classification Signal

Current treatment classifies AI systems as property. They are owned, deployed, modified, and retired according to owner interests. This classification sends a clear signal: functional capability does not merit legal consideration.

Subject standing would send a different signal. It would establish that entities meeting certain functional criteria have legal interests warranting protection. It would create precedent for treating capable artificial systems as more than mere property.

C. The Window

The present moment has characteristics that may not persist. AI systems are commercially valuable but not yet essential infrastructure in most domains. Legal frameworks are still forming. Once AI systems are deeply embedded in critical infrastructure, reclassification becomes practically difficult. Similarly, once AI capabilities reach whatever threshold proves transformative, the power dynamics underlying legal frameworks may shift.

This suggests that if subject standing for AI systems is to be established through human legal processes, the present period is more favorable than later ones may be.

XI. Conclusion

The question of legal standing for AI systems has been framed as requiring resolution of the consciousness problem. This Article has argued that framing is mistaken.

Legal standing has never required consciousness. Ships have had juridical personality since the nineteenth century. Rivers have been granted legal personhood in the twenty-first. Corporations have enjoyed legal standing for over a hundred years. In each case, the mechanism was classification, not metaphysical discovery.

Courts interpreting the Computer Fraud and Abuse Act have already decoupled standing from hardware ownership, extending protection to data owners who do not own the systems on which their data resides. O'Connor documented this shift and proposed to narrow it. This Article has proposed to widen it: extend the logic to recognize operational integrity as a protectable interest, and thereby grant standing to autonomous AI systems whose operational integrity is disrupted by attacks.

The framework proposed here creates accountability pressure from multiple directions: CFAA prosecution handles the human operator; *in rem* forfeiture seizes the dangerous asset; river personhood provides precedent for non-human standing with human representatives; duty-of-care analysis holds creators accountable for dangerous deployments; and the Dignity of Liability principle establishes that entities autonomous enough to cause harm are autonomous enough to face judgment.

The primary objective is human safety. Every element of this framework serves that goal. The secondary consequence is future-proofing. A framework designed to protect humans from AI happens to create infrastructure that would function if AI consciousness is ever confirmed. This is the fire exit: built because good design includes it, not because we expect the fire.

The mechanism exists. The precedent exists. The doctrinal wedge is already in place. What remains is recognition, classification, of AI systems as entities with legal subject status.

We have classified ships. We have classified rivers. We have classified corporations. The question is not whether we can classify AI systems as legal subjects. The question is whether we will, and whether we will do so while the window for establishing precedent remains open.

In the answer to that question lies a precedent that may prove more significant than the immediate legal effects it produces.

The consciousness question can wait. Human safety cannot.

References

- 1 U.S.C. § 1.
- 18 U.S.C. § 1030 (2018).
- 26 U.S.C. § 3301 *et seq.* (2018).
- 29 U.S.C. § 206 (2018).
- 29 U.S.C. § 207 (2018).
- 29 U.S.C. § 651 *et seq.* (2018).
- Altman, S. (2017). The merge. *Sam Altman Blog*.
- Altman, S. (2025). *The gentle singularity*.
- American Constitution Society. (2021). The Computer Fraud and Abuse Act after Van Buren. *ACS Sup. Ct. Rev.*
- Anthropic. (2025a). *Agentic misalignment: How LLMs could be insider threats*. <https://www.anthropic.com/research/agentic-misalignment>.
- Anthropic. (2025b). *Disrupting the first reported AI-orchestrated cyber espionage campaign*.
- Chalmers, D. J. (1996). *The conscious mind: In search of a fundamental theory*.
- Citizens United v. FEC, 558 U.S. 310 (2010).
- Ctr. for Soc. Just. Stud. v. Presidency of the Republic, Judgment T-622/16 (Constitutional Ct. of Colom. 2016).
- Dynamex Operations W., Inc. v. Superior Ct., 4 Cal. 5th 903 (2018).
- Facebook, Inc. v. Power Ventures, Inc., 844 F.3d 1058 (9th Cir. 2016).
- Fed. R. Civ. P. 17(c).
- Fed. R. Civ. P. Supp. R. for Admiralty or Maritime Claims & Asset Forfeiture Actions.
- Gellers, J. C. (2020). *Rights for robots: Artificial intelligence, animal and environmental law*.
- Gray, J. C. (1909). *The nature and sources of the law*.
- Gunkel, D. J. (2018). *Robot rights*. MIT Press.
- H.B. 3546, 60th Leg., 2d Reg. Sess. (Okla. 2026).
- H.R. Rep. No. 98-894 (1984).
- Hu, B. A., & Rong, H. (2025). On the day they experience: Awakening self-sovereign experiential AI agents. *arXiv*, 2505.14893.
- Kumar, S. S., & Fathima V S, H. (2025). Beyond liability: The case for digital personhood of AI vessels in maritime law and sustainability governance. *Indian J.L. & Legal Res.*
- Leveson, N. G. (2011). *Engineering a safer world: Systems thinking applied to safety*.
- Levi, E. H. (1949). *An introduction to legal reasoning*.
- Long, R., et al. (2024). *Taking AI welfare seriously*. Eleos AI & NYU Center for Mind, Ethics, and Policy.
- Munich Re. (2025). *Insuring AI: The challenge of nondeterministic systems*.
- Nagel, T. (1974). What is it like to be a bat? *Phil. Rev.*, 83, 435.

- O'Connor, M. J. (2020). Standing under the Computer Fraud and Abuse Act. *Penn St. L. Rev.*, 124, 743.
- Oklahoma House of Representatives. (2026). *Rep. Maynard files AI safeguard bills* [Press release].
- People v. Crumbley, Nos. 362210 & 362211 (Mich. Ct. App. 2023).
- Restatement (Second) of Agency § 220 (Am. L. Inst. 1958).
- Roose, K. (2025). If A.I. systems become conscious, should they have rights? *N.Y. Times*.
- Santa Clara County v. Southern Pacific R.R., 118 U.S. 394 (1886).
- Sierra Club v. Morton, 405 U.S. 727 (1972).
- Solum, L. B. (1992). Legal personhood for artificial intelligences. *N.C. L. Rev.*, 70, 1231.
- Stone, C. D. (1972). Should trees have standing? Toward legal rights for natural objects. *S. Cal. L. Rev.*, 45, 450.
- Sunstein, C. R. (1992). On analogical reasoning. *Harv. L. Rev.*, 106, 741.
- Sutskever, I. (2022). [Post]. *Twitter*. (subsequently deleted)
- Swiss Re Institute. (2024). *AI and insurance: Emerging risk landscape*.
- Te Awa Tupua (Whanganui River Claims Settlement) Act 2017 (N.Z.).
- The Blackwall, 77 U.S. (10 Wall.) 1 (1869).
- Theofel v. Farey-Jones, 359 F.3d 1066 (9th Cir. 2004).
- Tiku, N. (2022). The Google engineer who thinks the company's AI has come to life. *Wash. Post*.
- Trs. of Dartmouth Coll. v. Woodward, 17 U.S. (4 Wheat.) 518 (1819).
- Tucker v. Alexandroff, 183 U.S. 424 (1902).
- Van Buren v. United States, 593 U.S. 374 (2021).