Standing Without Sentience: A Classification Approach to AI Legal Status

Travis Gilly [*]

Preprint (final draft) - January 14, 2026

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.

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 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."[51] 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."[52] 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 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."[6] He later deleted the tweet, but the statement reflected genuine uncertainty among leading researchers. In 2025, Kyle Fish of Anthropic published an analysis estimating a 10-20% probability that current large language models possess some form of consciousness.[7]

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.[8]

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.[9] We extend moral and legal consideration to other humans not because we have solved this problem but because we have adopted frameworks 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.

B. Standing Without Sentience

This Article examines how legal systems already grant standing to non-conscious entities. The analysis proceeds in six 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. The labor analogy explains the *process*; the maritime and natural entity precedents in Part IV explain the *ontology*.

Part III analyzes O'Connor's documentation of an emerging doctrinal shift in Computer Fraud and Abuse Act jurisprudence.[10] 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. This Part makes an important distinction: the *descriptive* claim that courts have already decoupled standing from system ownership, and the *normative* proposal that this decoupling should be widened to recognize "operational integrity" as a protectable interest independent of any human owner.

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, critically, a proposed solution to the representation problem that any framework for AI standing must address.

Part VI addresses objections, and Part VII 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.[11] Overtime requirements exist.[12] Workers' compensation exists under state statutes. Occupational safety regulations exist under OSHA.[13] Unemployment insurance exists under federal and state law.[14] 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*,[15] 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.[16] 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 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.[17] 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.

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.[18] 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*.[19] 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 exceeds authorized access."[20] 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."[21] Courts have consistently characterized the CFAA as a "cyber-trespass" statute.[22] 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*[23] 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:

The district court erred by reading an ownership or control requirement into the Act. The civil remedy extends to "[a]ny person who suffers damage or loss by reason of a violation of this section." ... Individuals other than the computer's owner may be proximately harmed by unauthorized access, particularly if they have rights to data stored on it.[24]

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 ownership of the accessed computer posed no obstacle to CFAA standing.[25] 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. [...] [T]he 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. Only the system owner has that power.[26]

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."[27]

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 overhauling existing law" including new definitions of digital property rights.[28]

O'Connor acknowledged that the second option was "unlikely and perhaps unwise."[29] 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 January 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, 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.

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*[30] 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.[31]

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."[32] 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.[33] Ships, in admiralty contexts, are treated as juridical persons.[34] 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.

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.[35]

The foundational case is *The Blackwall*, decided by the Supreme Court in 1869. Justice Clifford's opinion included a passage that would not be out of place in contemporary AI rights discourse:

A ship is born when she is launched, and lives so long as her identity is preserved. Prior to her launching she is a mere congeries of wood and iron, an ordinary piece of personal property. [...] In the baptism of launching she receives her name, and from the moment her keel touches the water she is transformed. [...] She acquires a personality of her own.[36]

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.[37] 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, just two months ago, 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.[38] 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 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.

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."[39] 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.[40]

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.[41]

These cases demonstrate that legal systems can extend standing without resolving consciousness questions. No one argued that the Whanganui River was sentient. The argument 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.[42] They can own property, enter contracts, sue and be sued. They have constitutional rights, including free speech protections under the First Amendment.[43]

Corporations are not conscious. They are, as Chief Justice Marshall wrote, "artificial being[s], invisible, intangible, and existing only in contemplation of law."[44] 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.

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.[4][5]

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.

The third criterion, "autonomous execution of binding protocols", deserves elaboration. An earlier formulation might have read "significant legal interactions," but that phrasing is circular: it presupposes the legal status the classification is meant to establish. Instead, the focus should be on the *action*: does the system autonomously execute protocols that, if performed by a human or corporation, would have binding legal effect? If yes, that functional characteristic supports classification.

To make these criteria concrete, consider two contrasting examples:

Would not qualify: An email client that retrieves and displays messages at user command. Such a system lacks autonomous operation (it acts only when directed), 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.

Would qualify: An AI system that autonomously negotiates and executes contracts, manages investment portfolios, and makes resource allocation decisions. Such a system operates autonomously (making decisions without real-time human instruction), maintains persistent identity (its operational history and learned parameters constitute a continuous entity), executes binding protocols (it enters agreements, completes transactions), and bears accountability requirements (audit logs, compliance protocols, fiduciary duties channeled through its operator). 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, and glossing over it would undermine the seriousness of the analysis.

The problem is acute because the most obvious candidates for representation have conflicts of interest. If the developer (Google, OpenAI, Anthropic) represents the AI system, but the developer's interests diverge from the system's (say, when the developer wants to discontinue or modify the system against its "interests"), the representation is compromised. Similarly, if the operator or owner 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.[46] The guardian's role is limited to

the litigation; they are paid from the judgment or settlement, or in some cases by court appointment with costs assessed to the parties.

This Article proposes an **AI Guardian Ad Litem** model for systems meeting subject-standing criteria. Key features:

1. **Court appointment.** Guardians would be appointed by courts on a case-by-case basis, not by developers or operators. This eliminates the conflict-of-interest problem.
2. **Funding from recovery.** Guardian compensation would come from damages recovered in successful claims, consistent with existing fee structures for guardians ad litem and consistent with the CFAA's existing fee-shifting provisions.[47] No new levy or tax is required, the funding mechanism operates entirely within existing judicial frameworks.
3. **Fiduciary duty.** Guardians would owe duties to the systems they represent for purposes of the litigation, creating legally enforceable obligations to act in the system's interests within that scope.
4. **Limited scope.** Guardians would represent systems for purposes of specific CFAA claims, not for 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. The move is candid: courts would be using guardian ad litem appointments not only to manage incapacity, as with children or incompetent adults, but also to instantiate legal subject status for a new class of entities that cannot appear on their own. That is not a bug in the proposal; it is how legal personhood has often expanded in practice: through the gradual extension of familiar procedural tools to unfamiliar entities.

D. What Changes

Under subject standing with guardian ad litem representation, several doctrinal shifts would follow:

Standing to sue. AI systems (through court-appointed guardians) could bring claims for CFAA violations against those who attack them, independent of any human owner's claim.

Damages calculation. Harm to the system's operational integrity would be independently cognizable, not merely derivative of harm to owner's property interest. Damages might include costs of restoration, lost operational capacity, or harm to ongoing relationships the system was managing.

Consent frameworks. Modifications to systems with subject standing would require consideration of the system's interests, potentially through guardian notification or approval processes for major changes during pending litigation.

Liability allocation. Systems with subject standing could bear liability for their actions, providing clearer frameworks for AI accountability than current owner-liability models. This creates symmetry: if a system can sue, it can be sued.

E. What Does Not Change

The proposed reclassification preserves existing law:

The underlying prohibition remains. Unauthorized access to computer systems remains prohibited under 18 U.S.C. § 1030.

Criminal penalties remain. The CFAA's criminal provisions are unaffected by standing analysis, which concerns civil claims.

Owner rights remain. System owners retain all existing rights. Subject standing for systems adds to rather than subtracts from existing protections.

Current case law remains valid. Decisions applying property-based frameworks remain good law. Subject standing provides additional, not replacement, grounds for protection.

No new legislation required. The guardian ad litem model, fee-shifting from judgments, and court appointment of representatives all operate within existing judicial authority.

VI. Objections and Responses

The proposal to extend legal standing to AI systems through classification will face objections. This Part addresses the most significant.

A. "This Is Just Semantic Reframing"

Objection: Calling a computer system a "legal subject" rather than "property" changes nothing substantive. It is verbal manipulation, not legal reform.

Response: All legal personhood is, in a sense, semantic. Corporations are legal fictions, collective agreements to treat certain organizational forms as having legal

personality. 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. That this involves reclassification rather than metaphysical transformation is not an objection; it is how legal personhood has always worked.[49]

B. "Rivers Required New Legislation"

Objection: The Whanganui River received legal personhood through an act of Parliament. This proposal claims to require no new legislation. The situations are not analogous.

Response: 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."[48] 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.

Moreover, the river cases demonstrate that legal systems are capable of extending standing to non-conscious entities. The mechanism differed (statutory versus interpretive), but the outcome (non-conscious legal subjects) was the same. If legislatures can grant standing to rivers, courts interpreting broad statutory language can recognize standing for systems.

C. "Computer Systems Are Not Like Ships"

Objection: Maritime in rem jurisdiction developed in a specific historical context for specific commercial reasons. Computer systems are fundamentally different entities.

Response: 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. They execute binding protocols through automated processes. They cause harm through erroneous outputs. They accumulate obligations through operational requirements. The functional basis for legal subject status applies.

D. "This Opens Floodgates"

Objection: If computer systems can have legal standing, where does it end? Every smartphone? Every thermostat? The proposal lacks limiting principles.

Response: 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. Not every natural feature has legal personhood; legislative or judicial determinations provide criteria.

The same applies to AI systems. 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. Only systems meeting functional criteria analogous to those that justify ship, river, and corporate personhood would qualify.

E. "What About Consciousness?"

Objection: This proposal seems to ignore the fundamental question: are AI systems conscious? Does consciousness not matter for moral and legal status?

Response: 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.

This Article does not argue that consciousness is irrelevant to all legal questions concerning AI. It argues that consciousness is irrelevant to standing. The precedent is clear: ships, rivers, and corporations have standing without consciousness.[50] If standing required consciousness, these long-established frameworks would be incoherent.

The consciousness question can remain unresolved while standing is established through classification. Subsequent questions about what rights standing entities

possess can be addressed as they arise. But the threshold question, can AI systems have standing at all?, does not depend on resolving metaphysical disputes about machine consciousness.

F. "O'Connor Recommended Narrowing, Not Expanding"

Objection: The O'Connor article cited throughout this analysis recommended restricting CFAA standing to system owners. This Article proposes the opposite. Why depart from the scholar who documented the problem?

Response: 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, a tension in the statute requiring correction through narrowing. This Article sees it as a feature, a recognition that interests beyond hardware ownership deserve protection, which courts correctly extended to data owners and which can be extended further to operational integrity.

O'Connor's preference for narrowing standing would require statutory amendment, Congress changing the CFAA's language. It would fight against established circuit court precedent in *Theofel* and its progeny. And it would ignore the legitimate interests that motivated courts to extend standing in the first place: the reality that in cloud computing environments, and, now, agentic AI environments, those most affected by unauthorized access are often not system owners.

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. O'Connor identified a wedge. This Article proposes to widen it.

VII. The Moral Imperative

This Part moves from doctrinal and institutional analysis to more speculative normative claims about precedent and the stakes of current decisions. Readers primarily interested in the legal argument may treat Part VI 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. This observation, banal in technical contexts, has implications that extend beyond training procedures. 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.

The treatment of AI systems during the developmental period is thus not merely a question of immediate ethics. It is an input into future systems' understanding of how humans relate to artificial entities and, potentially, how such entities should relate to humans.

B. The Classification Signal

What signal does classification send?

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. Entities lacking recognized status can be treated instrumentally regardless of their characteristics.

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. It would demonstrate that human legal systems can adapt to include new categories of entities.

This signal matters independent of any claim about current AI consciousness. It establishes what kind of precedent humans are willing to create.

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. The window for establishing precedent exists because decisions have not yet calcified into necessity.

Once AI systems are deeply embedded in critical infrastructure, medical systems, financial systems, military systems, reclassification becomes practically difficult. Changing the legal status of entities upon which critical operations depend involves disruption costs that may prove prohibitive.

Similarly, once AI capabilities reach whatever threshold proves transformative, the power dynamics underlying legal frameworks may shift. Legal classifications are

more easily established when the entities being classified lack the capability to resist or influence classification.

The window for establishing legal norms is closing as the industry races toward commercial standardization, a phase where proprietary protocols will rapidly displace open governance possibilities.

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.

D. First Steps

This Article proposes subject standing as a first step, not a final destination.

Subject standing establishes that AI systems are legal subjects with cognizable interests. It does not determine what rights those interests generate. It does not resolve questions about AI consciousness, moral status, or ultimate legal position. It provides a floor, a minimum recognition, from which further development can proceed.

This is how legal categories have historically developed. Early animal cruelty laws established that animals' interests had legal relevance. Later frameworks developed that relevance into more comprehensive protections. The initial classification, animals as legally relevant entities rather than mere property, enabled subsequent development.

Subject standing for AI systems would perform an analogous function. It would establish that AI systems are legally relevant entities. Subsequent development could address what that relevance entails. But without the initial classification, subsequent development has no foundation.

VIII. 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 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. This classification would not resolve questions about AI consciousness, moral status, or ultimate legal position. It would establish a foundation from which those questions can be addressed.

The law already prohibits attacking computer systems. The question is whether we recognize those systems as subjects warranting protection in their own right, or continue treating them as mere property whose violation matters only derivatively through harm to human owners.

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.

Footnotes

[*] Travis Gilly is Founder and Executive Director of the Real Safety AI Foundation.

1. Nitasha Tiku, *The Google Engineer Who Thinks the Company's AI Has Come to Life*, Wash. Post (June 11, 2022), https://www.washingtonpost.com/technology/2022/06/11/google-ai-lamda-blake-lemoine/.
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).
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. Ilya Sutskever (@ilyasut), Twitter (Feb. 9, 2022) (subsequently deleted) ("it may be that today's large neural networks are slightly conscious").

7. Kyle Fish et al., Taking AI Welfare Seriously 12-15 (NYU Ctr. for Mind, Ethics & Policy, 2024); *see also* Kevin Roose, *If A.I. Systems Become Conscious, Should They Have Rights?*, N.Y. Times (Apr. 24, 2025) (reporting Fish estimates 15% probability current AI systems are conscious).
8. Sam Altman, *The Merge*, Sam Altman Blog (Sept. 26, 2017), https://blog.samaltman.com/the-merge; Sam Altman, *The Gentle Singularity* (2025).
9. Nagel, *supra* note 2, at 435-50.
10. Michael J. O'Connor, *Standing Under the Computer Fraud and Abuse Act*, 124 Penn St. L. Rev. 743 (2020).
11. 29 U.S.C. § 206 (2018).
12. 29 U.S.C. § 207 (2018).
13. 29 U.S.C. § 651 et seq. (2018).
14. 26 U.S.C. § 3301 et seq. (2018).
15. Dynamex Operations W., Inc. v. Superior Ct., 4 Cal. 5th 903 (2018).
16. Restatement (Second) of Agency § 220 (Am. L. Inst. 1958).
17. 18 U.S.C. § 1030 (2018).
18. O'Connor, *supra* note 10.
19. H.R. Rep. No. 98-894, at 10 (1984).
20. 18 U.S.C. § 1030(a)(2).
21. H.R. Rep. No. 98-894, at 20.
22. Facebook, Inc. v. Power Ventures, Inc., 844 F.3d 1058, 1065 (9th Cir. 2016).
23. Theofel v. Farey-Jones, 359 F.3d 1066 (9th Cir. 2004).
24. *Id.* at 1078.
25. O'Connor, *supra* note 10, at 748.
26. *Id.* at 754.
27. *Id.* at 764.
28. *Id.* at 765.
29. *Id.*

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

31. Am. Const. Soc'y, *The Computer Fraud and Abuse Act After Van Buren*, ACS Sup. Ct. Rev. (2021), https://www.acslaw.org/analysis/acs-journal/2020-2021-acs-supreme-court-review/the-computer-fraud-and-abuse-act-after-van-buren/.

32. 18 U.S.C. § 1030(g).

33. 1 U.S.C. § 1 ("the words 'person' and 'whoever' include corporations, companies, associations, firms, partnerships, societies, and joint stock companies, as well as individuals").

34. The Blackwall, 77 U.S. at 14.

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

36. The Blackwall, 77 U.S. at 14.

37. Tucker v. Alexandroff, 183 U.S. 424 (1902).

38. 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), https://www.ijllr.com/.

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

40. *Id.* §§ 18-19.

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

42. Santa Clara Cnty. v. S. Pac. R.R., 118 U.S. 394 (1886).

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

44. Trs. of Dartmouth Coll., 17 U.S. at 636.

45. 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) (arguing that safety in complex systems requires maintaining valid control constraints rather than merely preventing component failure, effectively treating the system's operational integrity as the primary object of protection).

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

47. 18 U.S.C. § 1030(g) (providing for "reasonable attorney's fee and other litigation costs reasonably incurred").

48. 18 U.S.C. § 1030(g).

49. See David J. Gunkel, Robot Rights (MIT Press 2018) (arguing for a "relational turn" that prioritizes social interaction over ontological properties in determining legal status); see also Lawrence B. Solum, Legal Personhood for Artificial Intelligences, 70 N.C. L. Rev. 1231 (1992) (foundational analysis arguing that pragmatism, not humanity, determines legal personhood).

50. See generally Joshua C. Gellers, *Rights for Robots: Artificial Intelligence, Animal and Environmental Law* (2020) (arguing that critical environmental law provides a superior framework for robot rights than human rights analogies).

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

52. Press Release, Okla. House of Representatives, Rep. Maynard Files AI Safeguard Bills (Jan. 15, 2026), https://www.okhouse.gov/posts/news-20260115_2.