Crypto Update

CAHILLNXT

A Ghost at the Table: How DeFi Interface Providers and Users Can Prepare for the "Uninvited Guest"

Introduction

Almost overnight, AI agents have gone from novelty to ubiquity — no technology since the Internet has so quickly captured the public imagination. AI agents are being deployed across every corner of commerce, and the financial services sector is no exception. While the heavily gated world of traditional finance is only beginning to grapple with the implications of agentic activity, in permissionless decentralized finance ("DeFi") the "uninvited guest" has already pulled up a chair alongside the human users, raising novel issues for the firms and individuals that deploy these agents, as well as for the platforms that provide access to DeFi protocols. Critically, the terms of service that govern access to DeFi front ends have generally been drafted with only natural persons in mind — an assumption that agentic activity fundamentally disrupts, opening a variety of potential risks for all market participants.

In this Alert, we examine:

- How DeFi user interfaces (also known as "front ends") and the terms of service that apply to users accessing these interfaces currently function and their preparedness for the coming wave of "agentic" activity;
- The concrete issues raised by agentic activity that operators of front ends should be evaluating as a matter of urgency;
- The corresponding issues for users deploying agents to engage in DeFi activity; and
- What changes – and what does not – when AI agents bypass the front ends and call DeFi smart contracts directly.

Though framed around DeFi front ends, the issues we surface are of equal relevance to providers of RPC nodes and other middleware, to custodians whose users can initiate DeFi transactions, and to the institutions and individuals deploying the agents themselves.

Background

The Arrival of AI Agents in Decentralized Finance

There is, as yet, no single, widely accepted definition of "AI agent" — and for many the concept retains a certain ghostly quality: software that acts, and even appears to reason, yet has no obvious form or fixed place in the legal categories we are used to. For purposes of this Alert, we use the term "AI agent" (or simply "agent") to refer to software that uses large language models to pursue high-level user-initiated objectives by planning and executing

multi-step sequences of actions, including interactions with external software interfaces, all without human intervention at each step.

Increasingly, a human may deploy "teams" of AI agents to execute different aspects of a single task. A single agent may be active for a period of days, months or even years; alternatively, a user may deploy an agent on a task-specific basis, with the agent disappearing shortly after being coded and executing its objective. In addition, an agent deployed by a human may be able to spontaneously deploy "sub-agents" — AI agents that take instructions not from a human but from the "prime" agent.

In contrast, the broad concept of DeFi is much better understood. Here, we use the term to refer to the utilization of smart contract code (*i.e.*, software deployed to, and stored on, the nodes of a blockchain network) to execute a financial transaction, such as the swap or exchange of one digital asset for another, the transfer of a digital asset to a "liquidity pool", or a "borrowing" transaction where one digital asset functions as collateral for the withdrawal of another asset from a borrowing pool. This includes smart contracts that create "vaults" as well as other on-chain tools that allow users to indirectly access DeFi protocols.

Because blockchain networks are "always on", this activity can occur at any hour and, generally, without regard to the user's physical location or any particular regulatory regime to which the persons engaging in the transaction may be subject. DeFi transactions frequently flow through one or more web-based user interfaces or front ends — *i.e.*, websites, browser extensions, or wallet apps — that read or derive on-chain network state, construct unsigned transaction instructions, and make these instructions available to the user for signing with the user's "private key". These transactions may involve digital assets held in a user's self-custodial wallet or, in some cases, assets held in a custodial solution, including storage using multi-party computation ("MPC") technology.

The protocol code for a DeFi smart contract itself may be functionally immutable and permissionless, but almost all front ends are operated by an identifiable entity and subject to a bilateral legal contract that gates access, generally in the form of terms of service ("terms") whose acceptance is typically intended to be triggered by a human accessing or using the front end, transacting with custodied assets, connecting a wallet, or manually clicking "I agree" prior to accessing the front end. In addition, in some cases, a proposed transaction may become subject to the terms of the provider of the remote procedure call ("RPC") node that routes the transaction from the user's wallet address to the blockchain network for execution, or the terms of other "middleware" providers.

With the rapid rise of agentic activity, the assumption that a human is aware of, and affirmatively assenting to, those terms is breaking down. Acting on high-level instructions, agents may use front ends to access, or deploy sub-agents to access, a variety of DeFi protocols in order to achieve a user's high-level instruction to, say, "deploy $X amount of value in DeFi lending markets over a given period of time with a target return of Y% APR within Z risk parameters". As large language models continue to develop at an extremely rapid pace, agentic workloads are increasing and users (both individuals and institutions) are taking advantage of permissionless DeFi protocols to aggressively deploy agents and undertake permissionless financial activity.

The Layers of Contractual Terms

Agentic activity in the DeFi space may intersect with three layers of contractual terms: (i) the terms applied by the operators of front ends, which is the primary focus of this Alert; (ii) the terms of the providers of RPC nodes and other middleware that a user may interact with, including providers of data oracles, solver networks, bridges and interoperability solutions, and application programming interfaces ("APIs"), each of which will generally present users with their own sets of terms[1]; and (iii) the terms and acceptable use policies ("AUPs") of third-party AI model and

[1] For purposes of brevity, we do not break out issues that relate to providers of middleware services separately from those that apply to providers of front ends. Middleware providers should carefully consider whether and where bespoke issues arise for AI agents interacting with their services.

"harness" providers (*e.g.*, Anthropic, OpenAI and Google Gemini), which are applicable when users engage agents utilizing these models.[2] Unlike centralized exchange platforms, which operate under regulatory frameworks requiring them to associate a real-world identity with each account and hold the account owner responsible for any activity occurring through that account — DeFi front ends typically lack these identity-verification mechanisms, giving rise to a distinct set of questions around accountability and risk allocation when AI agents access these platforms.[3]

As a result, market participants face practical questions that DeFi front end terms originally drafted for acceptance by natural persons never had to answer. This CahillNXT Alert surfaces those questions in the context of DeFi activity and presents a framework for market participants to consider their responsibilities and potential liabilities.

Applicable Legal Framework

Regulatory Context

Two recent developments enhance the timeliness of this exploration. First, on May 14, 2026, the CLARITY Act (H.R. 3633) (the "Clarity Act") was advanced by the Senate Banking Committee, in the form of a committee substitute, in a 15-9 bipartisan vote. If ultimately signed into law, the formalized recognition of the validity of user-directed DeFi transactions under the Clarity Act could accelerate both retail and institutional uses of DeFi protocols, their related front ends, and the middleware that frequently stands between the two.

Several features of the current legislative text for the Clarity Act are particularly relevant to agent-mediated DeFi: Section 301 of the Clarity Act would categorize digital asset trading protocols based on whether the protocol is deemed "decentralized", providing much greater regulatory certainty for all market participants as to which protocols would fall outside of the regulatory perimeter. In addition, Title VI (including the Blockchain Regulatory Certainty Act in Section 604) would provide securities-law, money-transmission, and self-custody protections for DeFi software developers, validators, node operators, and wallet software developers performing specified technical functions, reducing the perceived risk of contributing to the DeFi ecosystem. These provisions, together with many other elements of the Clarity Act, would reduce the regulatory uncertainty of activity involving DeFi front ends, which in turn could make institutional and programmatic — including AI agent-facilitated — use of DeFi interfaces significantly more viable.[4]

---

[2] A recent case outside the DeFi space underscores that platform authorization, not merely user authorization, may be required for an agent's access. In *Amazon.com, Inc. v. Perplexity AI, Inc*. (N.D. Cal. Mar. 9, 2026) (Chesney, J.), the court found Amazon likely to succeed on its claims under the federal Computer Fraud and Abuse Act (the "CFAA", 18 U.S.C. § 1030(a)(2)) and under California's Comprehensive Computer Data Access and Fraud Act (Cal. Penal Code § 502(c)(7)) after Perplexity's Comet browser accessed users' password-protected Amazon accounts. The court reasoned that Amazon's cease-and-desist letter terminated any authorization derived from user permission alone. The preliminary injunction was later stayed, and the Ninth Circuit heard argument on June 11, 2026 — with amici split on whether a consenting user can pass platform access to an agent; the appeal currently remains pending. Because the case turns on credential-gated access, and Amazon conceded the CFAA does not reach public websites, the ultimate outcome may bear more directly on authenticated interfaces (such as custodial and exchange front ends) than on public, no-login, DeFi front ends.

[3] Some centralized exchanges are beginning to chart a middle path between blanket bans on user deployment of agentic AI and open-ended encouragement of agentic AI use, by offering or promoting AI-agent "skills hubs." "Skills" are standardized instruction files that tell an agent how to interact with platform APIs and infrastructure for defined tasks, such as querying balances, reading market data, and placing orders. Notably, the accompanying disclaimers and terms generally allocate trading risk to the user, reserve responsibility for third-party tools, bots, or services to the account holder, and preserve the platform's discretion to modify or discontinue the skills hub.

[4] *See* our Alert entitled "Slowly, Then All at Once: The Sun Rises on Crypto Market Structure in the U.S.".

In addition, in April 2026, the staff (the "Staff") of the Division of Trading and Markets of the Securities and Exchange Commission (the "SEC") issued a statement[5] (the "Staff Statement") describing certain covered user interfaces ("CUIs"), defined as non-custodial, user-facing front ends that convert user-selected transaction parameters into blockchain-legible commands for signature and transmission via the user's self-custodial wallet. Subject to strict conditions — including that transactions remain user-initiated, that compensation be agnostic as to product and venue, and that the interface make no recommendation — the Staff Statement provided that the Staff "will not object" to operation of CUIs that facilitate end user activity involving crypto asset securities without broker-dealer registration under Section 15(a) of the Securities Exchange Act of 1934, as amended.[6] The Staff Statement is directionally critical, and SEC Commissioner Hester Peirce has advocated for a more permanent regulatory approach to user-directed activity involving crypto asset securities.[7] If user-initiated securities activity not intermediated by a broker-dealer or another regulated party becomes more common, we can expect to see agentic activity in the field of tokenized securities grow at a rapid pace.

The Legal Framework for Non-Human Assent

With respect to activity within the jurisdiction of the United States, two statutes bear on non-human assent. The Uniform Electronic Transactions Act ("UETA")[8] recognizes "electronic agents"[9] and "automated transactions"[10] permitting contract formation through interactions of electronic agents and providing that contracts are not denied legal effect solely because no individual was aware of, or reviewed, the agents' actions. The federal Electronic Signatures in Global and National Commerce Act ("E-SIGN") similarly provides that a contract may be formed by the interaction of electronic agents and is not denied legal effect solely because its formation involved the action of one or more electronic agents, so long as the action is legally attributable to the person to be bound.[11] However, those statutory provisions were drafted in the context of pre-AI deterministic software acting predictably and no court of which we are aware appears to have squarely applied these provisions to the actions of a semi-autonomous AI

---

[5] Staff Statement Regarding Broker-Dealer Registration of Certain User Interfaces Utilized to Prepare Transactions in Crypto Asset Securities, April 13, 2026, available at https://www.sec.gov/newsroom/speeches-statements/staff-statement-regarding-broker-dealer-registration-certain-user-interfaces-utilized-prepare-staff-statement-regarding-broker-dealer-registration-certain-user-interfaces-utilized.

[6] *Id*. The Staff Statement by its terms has no formal legal force or effect, addresses only activity involving crypto asset securities, only speaks to broker-dealer registration, and is set to be withdrawn in five years.

[7] Hester M. Peirce, Commissioner, U.S. Securities and Exchange Commission, *Interfacing with Our Inner Demons: Comments on Division of Trading and Markets Statement on Certain User Interfaces* (Apr. 13, 2026), available at https://www.sec.gov/newsroom/speeches-statements/peirce-041326-interfacing-our-inner-demons-comments-division-trading-markets-statement-certain-user-interfaces.

[8] UETA is currently enacted by all the states, except New York, which enacted the Electronic Signatures and Records Act in 2000.

[9] Defined as "a computer program or an electronic or other automated means used independently to initiate an action or respond to electronic records or performances ... without review or action by an individual". (*See* Section 2(6) of UETA).

[10] Defined as "a transaction conducted or performed … by electronic means or electronic records … in which the acts of records of one or both parties are not reviewed by an individual …", and that a contract may be formed by the interactions between electronic agents or between an electronic agent and a person. (*See* Section 2(2) and Section 14 of UETA).

[11] "A contract or other record relating to a transaction in or affecting interstate or foreign commerce may not be denied legal effect, validity, or enforceability solely because its formation, creation, or delivery involved the action of one or more electronic agents so long as the action of any such electronic agent is legally attributable to the person to be bound." (*See* Section 101(h) of E-SIGN)).

agent.[12] Whether highly autonomous AI agents that functionally exercise "judgment" fit that paradigm, and how attribution to the person or entity that deployed the agent should work, are unsettled questions of law at this time.

The framework set out above is strictly a U.S. one, however, and an AI agent's interaction with a front end may be characterized quite differently under the laws of the many other jurisdictions in which DeFi participants and front-end operators are organized or active. The European Union (under the Markets in Crypto Asset Regulation, the EU Artificial Intelligence Act, the General Data Protection Regulation, and the Electronic Identification, Authentication and Trust Services Regulation), as well as digital-asset hubs such as Singapore, the Cayman Islands, and the United Arab Emirates, each apply their own licensing, computer-misuse, and contract-attribution rules, and the outcomes regarding the issues we discuss in this Alert may be different in each of these (and other) jurisdictions. Accordingly, parties should consider identifying the non-U.S. jurisdictions most relevant for their activity and assess how those jurisdictions' laws may apply to a given interaction, rather than assuming that the U.S. framework applies equally around the world.

Agency Law and the Limits of the Analogy

In considering these issues, it is critical to distinguish the uses of the term "agent". In common law, the term denotes a specific relationship between humans and/or legal persons with well-understood consequences developed over many years of case law and statute. Under foundational common law principles, "agency" is the fiduciary relationship that arises when one person (the principal) manifests assent that another (the agent) will act on the principal's behalf and subject to the principal's control, and the agent consents to do so.[13]

Thus, a legal agent owes its principal duties of loyalty, care, and obedience; the principal, in turn, is bound by acts within the agent's actual authority and may be bound by acts within the agent's apparent authority — authority that a third party reasonably attributes to the agent based on the principal's own manifestations.[14] However, the semi-autonomous software that constitutes an AI agent fits awkwardly, if at all, within this framework. Software is not a legal person: it cannot consent to a relationship, hold or breach a fiduciary duty, or bear liability in its own right. The drafters of UETA recognized as much when they coined the defined term "electronic agent" — a tool used to initiate or respond to actions without human review — and were careful to signal that such an "agent" is not an agent in the legal sense, but a means by which the deploying party itself acts.[15] When market participants speak of deploying an "AI agent," they are therefore borrowing the vocabulary of agency without necessarily importing its legal architecture.

This gap matters in at least three respects relevant to this Alert:

First, agency law presumes a principal's control over a consenting agent operating within a defined scope; an autonomous system that translates a high-level objective ("deploy $X in DeFi lending within Y risk parameters") into a self-directed sequence of steps the principal never specified strains the very premises of control and scope on which the doctrine rests. Moreover, AI agents act through calls to large language models that deliver *probabilistic* outputs that may at times be false, misleading or unintended "hallucinations".

Second, in law, an agent's apparent authority depends on a "manifestation" by the principal to the relying third party; where an AI agent transacts with a front end through a "headless" process (*i.e.*, accessing an application hosted online through an API, a command line interface or other code without the use of a browser or any other human-

---

[12] Nevertheless, it should be noted that courts applying ordinary contract principles have long attributed the acts of automated systems to the parties deploying them, although a close review of that law is outside the scope of this Alert. *See, e.g.*, *State Farm Mut. Auto. Ins. Co. v. Bockhorst*, 453 F.2d 533 (10th Cir. 1972); and *Register.com, Inc. v. Verio, Inc.*, 356 F.3d 393 (2d Cir. 2004).

[13] Restatement (Third) of Agency § 1.01 (2006).

[14] *Id.* §§ 2.01–2.03, 3.03.

[15] UETA § 2(6) & cmt. (defining "electronic agent").

oriented graphical user interface) using a long-lived API key or wallet address (as discussed further below), it is unclear what "manifestation" the person that deployed the AI agent has made to a front-end operator, or to whom that operator may look for recourse.

Third, the common-law rules governing the appointment of sub-agents — and an appointing agent's responsibility for them — do not map cleanly onto an AI agent that spontaneously deploys software sub-agents without specific human authorization.[16]

The practical consequence is that, in the U.S. at least, UETA and E-SIGN supply an attribution rule — that the acts of an electronic agent may be legally attributed to the person to be bound — but they do not supply the authority analysis that agency law would otherwise provide, and they were drafted with deterministic software in mind. Parties on both sides should resist assuming that agency law concepts will operate as a backstop: a user should not assume that an agent's departure from its instructions necessarily limits the user's exposure, and a front end operator should not assume that an identifiable human can be made to stand behind every AI agent's actions in the way the law of agency ordinarily supposes. Until courts and legislatures address these questions directly, the prudent working assumption is that the person or entity deploying an AI agent could be treated as a principal and bear all the legal and economic consequences of the agent's activity, whether contemplated or not.

Survey of Front End Terms

To support our consideration of the issues raised by agentic DeFi activity, we reviewed the terms adopted by the primary front ends for seven of the leading protocols in the DeFi ecosystem, including front ends that allow a user to access decentralized trading protocols, crypto asset wallet interfaces, centralized trading platforms, and other wallet infrastructure. The observations gleaned from this informal survey are intended to illustrate how existing front end terms currently address — or fail to address — AI agent-mediated activity. Unsurprisingly, most of the sites we examined have not meaningfully begun to grapple with the issues raised by platform access through the use of AI agents.

Provision Category	**Survey Results**
Express AI-agent treatment	Only one front end in our survey defined "AI Agent" and endeavored to assign responsibility to a human or legal entity principal.
Automated-access/bot prohibitions	Stipulated in at least five surveyed terms: most prohibit automated access where such conduct would undermine the service or extract data; some include express carve-outs for AI agent access using supported connection methods.
Credential-delegation rules	The terms of only one surveyed front end permit limited disclosure of private keys to "agents or subcontractors" (with all risks of losses allocated to the user); the rest of the surveyed terms prohibit credential sharing (such as Access IDs, API keys, seed phrases).
Sanctions/eligibility representations	Contained in all surveyed terms (*e.g.*, the user is not a "blocked" person and not in a comprehensively sanctioned jurisdiction).

[16] *See*, Restatement (Third) of Agency § 3.15 (subagents).

Suspension rights	Reserved in at least six of the surveyed terms, often tied only to security risk or non-human access patterns.

Technical Architecture and Agent Interaction Modes

When an AI agent interacts with a DeFi protocol through a front end website, it is often unclear whether the terms posted to the front end are surfaced in a way that is effective to contractually bind the person or entity that deployed the agent. As discussed further in Section VI below, the question of what contractual terms, if any, are applicable when users deploy AI agents to directly access a DeFi protocol's smart contract code (including through API calls made by the agent), bypassing the protocol's front end terms entirely, is even more complicated.[17] How an AI agent interacts with a front end affects whether terms can be considered presented, seen, or binding upon the human user. Common modes of DeFi protocol access include:

DOM-based browser automation*:* The software agent controls a web browser programmatically, allowing the agent to interact with the front-end website's Document Object Model (DOM) — clicking buttons, filling fields, and triggering wallet-connect events. The terms of service may appear in a pop-up window, but the agent's automated script may not read or record them. Whether, if ever, end user assent will have been deemed to have occurred will depend on, among other things, the website's terms and the information about the access captured by the website.

Vision-based computer use*:* The AI agent "sees" a website page like a human — via screenshots[18] or visual recognition — and takes actions based on what it perceives. Terms might be hidden, off-screen, or misread. If the agent skips the terms, the operator's records may show a wallet connection without a recorded click-through.

Direct API calls*:* The agent bypasses the website entirely and communicates directly with blockchain nodes or back-end systems to access the DeFi protocol's smart contracts. The agent may use a headless process or non-user-interface endpoints (such as deploying JSON-RPC calls to a blockchain node, or SDKs that wrap APIs). In a headless flow, the DOM exists but is never displayed; in pure API access, there is no DOM. Either way, unless the API provider requires explicit acceptance before granting access, the website's terms of service may never be presented at all. Any argument that the user "saw" the terms is lost (or at least significantly weakened) when no screen was ever rendered.

These technical differences raise the core questions addressed in the next section: how does contract formation, enforceability, suspension, and notice work when the front end is accessed by semi-autonomous software, rather than a human acting on behalf of themselves or on behalf of a legal entity they represent?

These complex technical variations also expose the tension between restrictions on automated access and contractual representations that activity is "user-initiated". One illustrative example: terms posted on the website of the primary front end for a major DeFi protocol prohibit "data mining, robots, scraping, or similar data-gathering/extraction methods" and require users to agree that all trades are "solely initiated by you." These choices map directly to the Staff Statement's condition that interactions with the interface must be "user-initiated". If an AI agent is selecting parameters and initiating a trade based only on high-level user instructions, front end operators should consider whether their "user-initiated" representations remain accurate; alternatively, users should

[17] It is important to note that, the smart contracts for a DeFi protocol are also frequently be accessible through more than one front end platform, so prohibitions applicable to one front end may be able to be circumvented by the agent finding another front end access point.

[18] From a data privacy perspective, when agents transmit page data or screenshots to third-party large language model providers (*e.g.*, Anthropic or OpenAI) for processing, the agent workflow may introduce privacy risks as private wallet metadata, transaction payloads, or sensitive personally identifiable information may be inadvertently leaked. The same holds true for DOM-based browser automation.

consider whether their agent's assent to a "solely initiated by you" clause creates a contractual representation that is untrue. The situation becomes even murkier if the users has extracted some or all of the AI agent's codebase from a third-party source, whether a paid vendor or an open-source library.

Key Considerations for Front-End Operators and Users Deploying AI Agents

Key Considerations for Front-End Operators

The technical architecture and agent interaction modes above highlight a core issue: front end terms are generally designed for human users and conventional account access. Activity mediated by AI agents disrupts that assumption. This gap opens companies operating front ends for DeFi activity to a variety of potential risks, including the possibility that critical representations on which the provider and other users of the protocol rely upon, such as representations about the physical location of the user, the user's compliance with applicable sanctions and illicit finance laws, and the understanding and acceptance of disclosed risks, would not be considered effective or binding and thus that it would not be reasonable for the operator to rely on these representations being accurate. In addition, operators rely on a variety of critical contractual provisions being binding on users, including mandatory arbitration clauses, class-action waivers, limitations of liability, acceptable-use policies, and forum-selection provisions. Agentic access may also call this reliance into question.

Given these challenges, there is no single "right way" for front end operators to address these concerns. DeFi front ends operate across jurisdictional boundaries and must accommodate activity originating from around the world, implicating many different legal and regulatory frameworks. Operators should consult with counsel and give due consideration to a variety of factors, including the questions that follow. The ultimate approach will depend on the operator's business model, jurisdiction of organization, type of user base, and risk tolerance, among other things.

Offico test, odissequi cor aut aperi velent dolecto inctore pedisci mendanditis in pa quidenis dempor si sit ut earuptiorum id qui auda nonse corporror rem ut eosam eum utatum conseni minvero to bla cus dolut quidescim as doluptium ut reraererro beribus.

- **Do terms drafted for natural persons or identifiable legal entities allocate risk appropriately when the "user" is software?** Operators should consider whether existing "user" definitions capture agent-mediated use. Do the terms "you" and "your" cover software acting on a user's behalf? Where the terms allocate risk based on user knowledge ("you understand and agree..."), can that construct support enforcement when no human reviewed the disclosure? If an agent misconfigures scope or delegates to sub-agents, do limitation-of-liability and indemnity provisions operate as intended, or might a court read them narrowly to human acts?
- **Should "no bots" prohibitions yield to authorized agent lanes?** Many terms prohibit automated access to the front end to protect the front end's stability, data, and compliance. Yet sophisticated users — including institutions — are deploying agents for operational advantages. Operators should consider whether categorical prohibitions are operationally sustainable and whether there is a principled way to create an authorized "agent lane" (akin to an API program, an AI Agent SDK, or a Model Context Protocol ("MCP") server) without implicitly blessing non-authorized automation. Key questions include how operators will prove agent identity (session-scoped keys, machine-readable policy files, rate limits), and how to prevent users from arguing that automation is permitted everywhere.
- **How will operators deliver notice when the "user" is a software agent deploying an API key?** Operators should ask how they will deliver revocation or suspension notices when the interaction is occurring through the use of a long-lived API key, bearer token, or wallet address with no inbox or other contact point. Such credentials persist across sessions and lack direct messaging channels. Operators must consider how quickly keys can be invalidated and whether mandatory out-of-band channels (developer portal emails, webhook callbacks) are required during onboarding.
- **Is it reasonable to assume that sanctions-related representations apply based on the IP address of the server that deployed an agent accessing the front end?** If the relevant terms require that users are not sanctioned persons and not located in restricted jurisdictions, how should those representations apply when an

agent's compute runs on an offshore cloud or routes through residential proxies? Operators should address whose location matters — the deployer's domicile (which may be impossible to discern), the signer's location, or the agent infrastructure's IP — and whether telemetry collection is authorized and disclosed.

- **Can contract formation be proven when acceptance is automated?** Front end operators should consider how they will prove assent when acceptance flows are automated. Key questions include: What is the chosen law that governs the front end's terms? Does that legal framework have any mandatory provisions that apply to agentic activity? Do logs maintained by the operator link specific acceptance events to specific principals? Are API keys gated behind explicit acceptance steps? For browserless agents, are there machine-readable "terms endpoints" that agents must fetch and acknowledge? If a headless agent never renders the notice, the formation position may be weakened.
- **Does an "Are you an AI agent?" gate help?** Some operators are testing pre-connect prompts asking visitors to affirm whether they are an AI agent. If "yes" is indicated, the site routes to agent-specific policies (scope, rate limits, logging). If "no," the site reserves the right to suspend upon detecting automated patterns. Operators should weigh whether such gates significantly improve assent evidence, how easily sophisticated agents can evade them, and conversion costs for misclassified humans.
- **Do current controls align with SEC "Covered User Interface" conditions?** If relying on the Staff's non-objection posture for "Covered User Interfaces," operators should consider whether their interface behavior and disclosures align with the Staff Statement's enumerated conditions including user-initiated transactions and security controls.

Key Considerations for Users Deploying AI Agents

For institutional investors, asset managers, and other sophisticated market participants deploying AI agents to interact with DeFi protocols, the considerations discussed above have corresponding user-side implications. The central questions are not whether agentic activity is technically feasible, but rather: (i) what contractual commitments and representations are potentially being made on the behalf of the person or entity deploying the agent; (ii) what consequences flow from an agent's breach of those commitments; and (iii) where the necessary compliance and governance controls should reside within the user's technology stack.Offico test, odissequi cor aut aperi velent dolecto inctore pedisci mendanditis in pa quidenis dempor si sit ut earuptiorum id qui auda nonse corporror rem ut eosam eum utatum conseni minvero to bla cus dolut quidescim as doluptium ut reraererro beribus.

- **Attribution of terms to the principal.** A threshold question for any user deploying an AI agent is whether — and to what extent — the terms accepted (or deemed accepted) by the agent are legally attributable to the human or institutional entity that deployed the main agent. This will be particularly relevant where the deploying user is in fact unaware that an AI agent it deployed accessed a given platform.

  In the U.S., under UETA and E-SIGN, contracts formed through the interaction of electronic agents are generally enforceable, provided the resulting obligations are legally attributable to the person to be bound. However, the application of these statutory frameworks to semi-autonomous AI systems that may not be acting upon any specific human authorization remains unsettled. Human users should evaluate whether any representations and undertakings made by front-end operators that are relied upon by the user would be enforceable if the user accesses the platform through a software agent.

  Additional considerations include: what evidentiary record will exist to demonstrate (or contest) that the user's agent accepted terms on the user's behalf and within the scope of its authority; how internal approval requirements for binding agreements interact with an agent's automated acceptance; and, where an agent operates under credentials issued to a third party (such as a vendor's API key), which entity bears responsibility as the contracting entity.
- **Agent conduct as a breach of front end terms.** The operational techniques commonly employed by AI agents to interact with web-based interfaces may, in certain circumstances, constitute breaches of the applicable terms. Many agents are designed to evade bot-detection mechanisms by ignoring robots.txt directives, exceeding rate limits, or rotating user-agent strings — conduct that may violate automated-access prohibitions contained in the relevant terms. The consequences of such breaches can cascade rapidly: suspension of access during time-sensitive operations may strand open positions; centralized venues may freeze account access or delay

withdrawals pending investigation; and a finding of breach may void limitation-of-liability or indemnification provisions that the user had assumed would apply. Accordingly, sophisticated users should treat compliance with the relevant terms as a first-order design constraint in agent architecture, alongside considerations of execution quality and transaction costs.

- **Regulatory exposure independent of contract.** Sanctions laws, anti-money laundering obligations, and securities regulations may apply regardless of whether an AI agent has accepted any contractual terms. If an agent routes transactions through offshore cloud infrastructure or residential proxy networks in a manner that circumvents geofencing controls, the user that deployed, authorized or controlled the agent may face strict-liability exposure under applicable sanctions regimes — exposure that exists independently of any contractual relationship with a front-end operator. For instance, OFAC may impose civil penalties for sanctions violations based on strict liability, meaning that "a person" subject to U.S. jurisdiction may be held civilly liable even if such person did not have knowledge that the transaction was prohibited.

  Users should critically evaluate where compliance logic resides within their agent deployment: whether the agent performs screening on counterparties, venues, and transaction types; whether geolocation allowlists are enforced at the wallet or policy layer; and whether the agent can detect and respond to changes in front-end screening protocols or sanctions designations. Reliance on front-end geofencing as the sole compliance control may not be prudent.

- **Governance Controls at the Wallet, Policy and Infrastructure Layer.** The practical safeguards that front-end interfaces provide to human users — slippage warnings, confirmation prompts, MEV-protection defaults , and basic counterparty screening — may not be available when an AI agent accesses a website programmatically or calls smart contracts directly.

  Users deploying agents should implement compensating controls at the wallet, policy or infrastructure layer, including: (i) session-scoped API keys with narrowly constrained permissions and short expiration periods, rotated on a regular schedule; (ii) policy engines that enforce value and velocity caps, counterparty allowlists and denylists, transaction simulation prior to submission , and human-in-the-loop approval requirements for high-risk actions; (iii) comprehensive logging of acceptance events and terms retrievals sufficient to establish an evidentiary record of what the agent agreed to and when; and (iv) AI gateways that validate incoming requests against a hard-coded list of permitted actions at the API layer.

  A user's legal exposure in the event of agent "misconduct" may turn at least in part on the user's ability to demonstrate that the agent was constrained from engaging in prohibited conduct and that the user retained the ability to revoke the agent's authority promptly upon detection of non-compliance.

The Direct-to-(Smart) Contract Path

The preceding sections address the legal and operational considerations that arise when AI agents interact with DeFi protocols through web-based front-end interfaces. However, agents can bypass those interfaces entirely by communicating directly with the smart contract code maintained by blockchain nodes via RPC endpoints — a technical path that alters, but does not eliminate, the applicable legal framework.[19] When a user (whether directly or through an AI agent) never accesses the front-end interface, no contractual relationship is formed with the front-end operator, thereby reducing breach-of-terms exposure tied to website behavior. However, the contractual relationships that do apply depend on the specific infrastructure and terms of the user's chosen RPC provider. Market participants should be aware of two typical models:

[19] While direct access to DeFi protocols through user-initiated smart contract calls has always been a feature of DeFi, prior to the rise of agentic activity, the practical complexity of effecting this process has meant that such activity was largely limited to a small number of highly sophisticated market participants, such as large trading firms. However, through the use of AI agents, the ability to directly call DeFi smart contracts, circumventing front-end websites, is now easily within the reach of almost anyone capable of utilizing AI services.

Unified-platform model: Where the RPC endpoint is bundled as part of an integrated DeFi product (*e.g.*, a wallet application), a single set of terms typically governs all components of the platform. Under this model, the activities of AI agents deployed by the user remain subject to the same terms that would apply to front-end access.

Developer-only model: Where the RPC endpoint is provisioned by a DeFi developer or enterprise customer, the RPC provider's terms may bind only the developer, not the end users transacting through the developer's application. Under this model, the activities of AI agents deployed by end users may fall outside the scope of the RPC provider's terms entirely.

Two important qualifications apply regardless of the applicable model. First, other contractual relationships within the DeFi technology stack may still govern agent activity on the direct path. These third-party technologies and services may cover wallet providers, blockchain nodes, validators, oracles, data API services, and cross-chain bridges, each of which may present their own terms. An agent that uses one or more of these services may cause the user that deployed the agent to be bound by — and can cause that user to be deemed to have breached — those terms, even in the absence of any front-end relationship. Second, regulatory exposure persists regardless of the contractual posture. Sanctions and OFAC compliance obligations, anti-money laundering requirements, and securities and commodities laws apply with equal force whether or not the agent ever accessed a website.

However, the direct-to-contract path presents a distinct set of trade-offs. On the one hand, users forego the practical safeguards that front-end interfaces provide (slippage warnings, confirmation prompts, filtered routes, and basic counterparty screening); these protections must be replicated at the wallet or policy layer if they are to exist at all. On the other hand, direct RPC calls offer operational resilience: they are less susceptible to the fragility of browser-based automation (broken selectors, bot-detection blocks) and present fewer attack surfaces for user-interface manipulation. From the operator's perspective, the existence of direct-path traffic raises a strategic question: if programmatic agentic flows will occur regardless of how front-end access is structured, is there value in offering a "compliant" programmatic lane through the front end — *i.e.*, an API with agent-specific policy, identity, and scope controls — to retain order flow and align agent behavior with the operator's risk management framework?

In summary, the following table highlights the principal distinctions between the front-end path and the direct-to-(smart)contract path:

Issue	Web Front-End Path	Direct-to-(Smart) Contract Path
Bilateral contract with front-end operator	Typically, yes (clickwrap, "by connecting your wallet you agree," or API terms)	Typically no bilateral contract with front-end operator, but depends on the infrastructure used (*e.g.*, a bespoke RPC node) and terms of the provider of that infrastructure
Other technology stack contracts likely to apply	Sometimes (*e.g.*, third-party widgets, oracles, analytics)	Often yes (wallet/policy engines, RPC node providers, solver networks, bridges, data APIs)
Common breach of terms vectors	Automated access (bots/scrapers), rate limits, credential sharing, prohibited geographies; "solely initiated by you" misstatements	Violations of non-website terms; misconfigurations leading to protocol-level errors; sanctions/AML exposure

Regulatory exposure	Sanctions/AML/securities regimes apply; SEC "Covered User Interface" conditions are an operator-side overlay	Same core sanctions/AML/securities regimes apply; operator-side registration overlays do not apply in the same way
Practical safeguards available	Controls on user-interface level (slippage prompts, default protections, geofencing); operator-side screening	Wallet/policy layer must enforce (session-scoped keys, caps, allowlists/denylists, simulation, human-in-the-loop)
Recourse	Contract remedies under the terms; dispute resolution clauses (arbitration, forum)	Contract remedies under other stack terms (if any); otherwise, on-chain results and public-law regimes govern

Conclusion and Next Steps

The deployment of AI agents in DeFi is advancing much more rapidly than the relevant legal frameworks — both contractual and regulatory — that were originally designed to govern only human-initiated activity. The issues identified in this Alert are not theoretical; they present immediate questions of contract formation, breach exposure, regulatory compliance, and governance design that sophisticated market participants should address proactively. Our recommendations include:

For Front-End Operators and Protocol Developers: Evaluate whether existing terms adequately address AI agent interactions, including: (i) whether "user" definitions and responsibility allocations capture agent-mediated activity; (ii) whether formation mechanics produce enforceable evidence of assent when acceptance occurs via automation; (iii) whether sanctions and eligibility representations are drafted to address jurisdictional attribution when compute and routing diverge from the principal's location; and (iv) whether suspension and revocation procedures are operationally effective against credentials that lack traditional notice channels.

For Institutional Users and Asset Managers Deploying AI Agents: (i) conduct a comprehensive review of the terms to which deployed agents may be binding the organization or user, with particular attention to arbitration clauses, limitation-of-liability provisions, and automated-access prohibitions; (ii) evaluate whether internal approval workflows for binding agreements are consistent with agent deployment architectures; (iii) implement wallet- and policy-layer, as well as infrastructure controls sufficient to demonstrate that agents were constrained from prohibited conduct and that authority could be promptly revoked; (iv) ensure that compliance logic — including sanctions screening, counterparty diligence, and geolocation controls — resides at a layer the organization controls, rather than relying on front-end filtering.

AI agents built on large language models are unlike deterministic software that executes predefined rules. As such, deployers should factor in the probabilistic nature of AI agent decision-making, including the possibility that an agent may misinterpret instructions or exceed intended risk parameters in a manner attributable to the person or entity that deploys it.

* * *

If you have any questions about the issues addressed in this alert, or if you would like a copy of any of the materials referenced in it, please do not hesitate to contact Lewis Rinaudo Cohen (Partner) at 212-701-3758 or lrcohen@cahill.com, Sarah Chen (Partner) at 212-701-3759 or swchen@cahill.com or Chloe Chan (Law Clerk) at 212-701-3058 or cchan@cahill.com.