

The Law of the Machine Economy

How Autonomous Systems Transact, and the Law That Governs Them

September 22, 2026

1. AUTONOMOUS SYSTEMS ARE ENTERING THE FINANCIAL MARKETPLACE

Across traditional and token-based finance, autonomous systems built around generative AI language models, commonly called “agents,” are opening new possibilities for how firms conduct business. For banks, broker-dealers, investment funds, proprietary trading firms, market makers, custodians, and businesses operating in decentralized finance, the commercial promise of these systems lies in new products, trading strategies, services, and customer and counterparty relationships. By delegating sequences of decisions and actions within a defined mandate, firms can develop offerings that respond to changing economic conditions and customer needs, coordinate trading and funding across markets, and operate beyond conventional hours. The scope of that delegation becomes part of what the firm can offer and what others can expect from it.

Realizing those opportunities and managing the related risks require decisions about when and how an autonomous system may commit a firm and how the resulting rights and responsibilities will be allocated. Those questions also confront firms that choose not to deploy autonomous systems themselves: customers, counterparties, and platforms may introduce these systems into existing business relationships.¹ Firms need to address these questions now, both in designing their own offerings and in setting the terms on which others may transact with them. The immediate task is to align transaction terms, delegated authority, execution controls, and evidence of performance.

The relevant activity may also be distributed across jurisdictions, with the firm, its service providers, and its counterparties potentially subject to different legal and regulatory requirements. Decentralized financial infrastructure can make those connections even harder to identify and the legal requirements even harder to reconcile.

2. THE LAW OF THE MACHINE ECONOMY

The Law of the Machine Economy, or Machine Law, applies to transactions of value in which at least one party acts directly or indirectly through an autonomous system operating with limited human intervention, changing the legal position of the parties over time.

¹ A Ghost at the Table: How DeFi Interface Providers and Users Can Prepare for the “Uninvited Guest”, CahillNXT Client Alert (June 17, 2026), <https://www.cahill.com/publications/client-alerts/2026-06-17-a-ghost-at-the-table-how-defi-interface-providers-and-users-can-prepare-for-the-uninvited-guest>

Six recurring questions underlie this field:

- **Authority:** what has been authorized;
- **Attribution:** whose act it is;
- **Organization:** which legal persons hold assets, direct systems, and bear obligations;
- **Verification:** how identity and performance are established;
- **Settlement:** what is delivered, and to what legal effect; and
- **Recourse:** how errors are corrected, unwound, and settled between the parties at the speed of human law.

These questions connect private law with applicable regulatory requirements. Answering them requires understanding how autonomous systems select and carry out actions that affect legal rights and obligations.

These systems are entering a financial marketplace that is itself changing quickly. SEC Chairman Paul Atkins has publicly supported the concept of tokenization of securities, and major U.S. market infrastructure providers are actively developing tokenization services² and, on September 17, 2026, the SEC released an order granting temporary exemptive relief to “tokenized securities venues” (as well as to the persons providing liquidity to these venues) to allow the trading of tokenized National Market System (NMS) stock using blockchain-based “automated market maker” applications.³ In June 2026, central clearing for U.S. equities expanded to 24 hours a day, five days a week, with further expansion potentially on the horizon.⁴ A retail trading platform’s agent-account disclosures currently state that customers assume the risk of orders placed by their agents.⁵ And, in a June 2026 speech, the Bank of England’s Deputy Governor for Financial Stability cautioned that requiring a human in the loop for every agent action was unlikely to be realistic.⁶

This landscape makes coming to terms with the allocation of authority and responsibility when agentic activity is present a matter of immediate relevance to firms participating in the financial markets. These questions become even more important as autonomous systems participate in markets for tokenized securities, stablecoins and other crypto assets, event contracts, perpetual futures, and other evolving financial products whose own legal character and regulatory status remain uncertain. The terms governing the participation of autonomous systems must account for both the scope of delegated activity and the rights, obligations, and trading arrangements associated with the products themselves.

WHY IS IT DIFFERENT NOW?

Market participants understandably ask what distinguishes language model-based agents from other forms of automation that are already widely used. High-frequency trading, algorithmic execution, and automated portfolio management have long involved the execution of consequential decisions without individual human approval. Some current systems also use machine learning; the SEC’s 2020 staff report described its application to trade execution and bond pricing.⁷ The significance of language model-based agents lies in how they interpret instructions and information, select actions, and use the results to continue pursuing an objective.

² Paul Atkins, *Keynote Address at the Crypto Task Force Roundtable on Tokenization* (May 12, 2025) (“the migration to on-chain securities has the potential to remodel aspects of the securities market”), <https://www.sec.gov/newsroom/speeches-statements/atkins-remarks-crypto-roundtable-tokenization-051225-keynote-address-crypto-task-force-roundtable-tokenization>; SEC Release No. 34-105047 (Mar. 18, 2026) (approving Nasdaq’s rule change to trade eligible securities in tokenized form), <https://www.sec.gov/files/rules/sro/nasdaq/2026/34-105047.pdf>; DTCC, *DTCC Advances Development of New Tokenization Service, Convenes 50+ Firms to Drive Digital Assets Adoption* (May 4, 2026), <https://www.dtcc.com/press-releases/2026/dtcc-advances-development-of-new-tokenization-service>; DTCC, *DTCC Turns Tokenization into Reality: U.S. Trades Successfully Processed Using DTC-Tokenized Assets* (July 15, 2026), <https://www.dtcc.com/press-releases/2026/dtcc-turns-tokenization-into-reality>.

³ SEC, *Press Release, SEC Issues “Innovation Exemption” to Facilitate the Trading of Tokenized NMS Stock and Request for Comment* (Sept. 17, 2026),

<https://www.sec.gov/newsroom/press-releases/2026-90-sec-issues-innovation-exemption-facilitate-trading-tokenized-nms-stock-request-comment>.

⁴ DTCC, *DTCC’s NSCC Now Live with Clearing Hours Extended to 24x5 Model, Marking Major Milestone for U.S. Equities Market* (June 29, 2026),

<https://www.dtcc.com/press-releases/2026/nscc-now-live-with-clearing-hours-extended>.

⁵ Robinhood, *Agentic Trading Disclosures* (accessed Sept. 20, 2026), <https://robinhood.com/us/en/agentic-trading/>.

⁶ Sarah Breeden, *Agents of Change* (Bank of England, June 30, 2026), <https://www.bankofengland.co.uk/speech/2026/june/sarah-breeden-panel-at-the-european-central-bank-forum-on-central-banking-2026>.

⁷ SEC Staff, *Staff Report on Algorithmic Trading in U.S. Capital Markets* (Aug. 5, 2020), https://www.sec.gov/files/algo_trading_report_2020.pdf, at 18, 52.

In a typical arrangement, an agent combines a language model with surrounding software, often called a “harness,” that assembles instructions and information for the model and awaits feedback. Based on the prompt provided, the model generates a natural language output for the harness on a probabilistic basis. For example, the model may propose using an available software tool to read a document, obtain a valuation, or initiate a transfer.⁸ The harness processes that proposal, applies whatever execution restrictions have been implemented, and invokes the relevant tool (a service or process), such as requesting data from a valuation provider or submitting an order to a trading platform. The results then inform another model call. This cycle can continue through a sequence of decisions not individually specified in advance and may proceed without human approval of each step.⁹ Critical to understanding this process is that, in this arrangement, the model itself is “stateless” with respect to the conversation. Each call by the harness software draws on the model’s training and the context made available for that call. Earlier exchanges may nevertheless influence later responses through the history, summaries, or other information supplied by the surrounding system.

The harness can enforce deterministic rules, such as refusing a transfer above a stated amount; however, the model’s generation of proposed actions draws on patterns learned during training and depends on the context supplied to it. The model generates its response incrementally, selecting successive units of text using probabilities shaped by its training and the context available at each step. A natural language instruction, for example, to use only a defined set of eligible collateral influences that process but does not itself install an eligibility check in the transfer service. Firms and their internal and external counsel need to understand both how the system selects an action and which restrictions actually govern execution.

The information supplied to the model also evolves during a session or process. Later outputs from a model may reflect new messages, retrieved information, and summaries of earlier exchanges. Some systems compress accumulated iterations of prompts and responses to accommodate a language model’s finite context window. That process, known as “compaction,” can omit a qualification whose importance only becomes apparent later.¹⁰ Concerns can arise from the interaction of changing context, probabilistic output generation, and permissions that allow a proposed action to be executed. A firm may have developed tools and procedures to manage these effects in its own agentic systems, while a counterparty may have taken different precautions or none at all. The resulting interaction may produce outcomes neither party expected.

3. DIGITAL ASSET EXPERIENCE SUPPLIES A METHOD

Digital asset work supplies a practical method for analysis: trace how the system operates, who controls each consequential step, and which legal relationships give that activity effect. A decentralized trading platform requires analysis of its code, access arrangements, external data sources, and administrative powers before counsel can address the resulting securities, derivatives, payments, and sanctions questions. Similarly, labeling a given process an “AI agent” establishes little about what a firm has authorized or which participants its actions may bind. Combining agents with blockchain-based systems makes tracing those relationships even more important.

Practical control, ownership, contractual entitlement, and access to remedies may reside with different entities in different jurisdictions. Emerging industry customs may also affect their relationships. An intermediary may control

⁸ Although the output of these models has generally been in the form of natural language, new types of natural-language-conditioned “decision models” have recently emerged that output numerical probabilities in response to natural language inputs. See *Introducing System One Models & Jev* (Sept. 15, 2026), <https://typesafe.ai/blog/introducing-system-one-models-and-jev>.

⁹ Anthropic, *Building Effective Agents* (Dec. 19, 2024) (discussing workflows, agents, and feedback through tool use), <https://www.anthropic.com/engineering/building-effective-agents>; Anthropic, *Tool Use with Claude* (accessed Sept. 13, 2026) (distinguishing model-generated tool requests from execution by an application or provider), <https://platform.claude.com/docs/en/agents-and-tools/tool-use/overview>. These describe illustrative architectures; implementations vary. See also Lei Wang et al., *A Survey on Large Language Model Based Autonomous Agents*, 18 *Frontiers Computer Sci.* art. 186345 (2024), <https://link.springer.com/article/10.1007/s11704-024-40231-1>.

¹⁰ Anthropic, *Effective Context Engineering for AI Agents* (Sept. 29, 2025) (section on compaction, explaining summarization of conversation history and the possible loss of subtle but critical context), <https://www.anthropic.com/engineering/effective-context-engineering-for-ai-agents>.

the means of transfer while an investor holds rights against it; a system may record a movement of value without resolving whether an obligation was discharged. These distinctions shape the product, its custody arrangements, and its use in trading or collateral services. Delegating asset selection or transfer to an agent adds another relationship to that analysis.

Existing U.S. law provides a foundation for consideration of these issues. The federal E-SIGN Act, for example, recognizes that the enforceability of certain contractual arrangements cannot be avoided solely due to the involvement of an electronic agent, but the statute sidesteps the critical question by conditioning this outcome on the action being “legally attributable” to the person to be bound. Whether an action by a complex agentic system is legally attributable to a person is left to other law. The Uniform Electronic Transactions Act (“UETA”), enacted in every state except New York, supplies more of that law. An electronic record is attributable to a person if it was the act of the person, which “may be shown in any manner,” and a contract may be formed by the interaction of the parties’ electronic agents “even if no individual was aware of or reviewed” what they did.¹¹ Drafters of UETA anticipated systems that might “learn through experience, modify the instructions in their own programs, and even devise new instructions,” but the UETA still treats the electronic agent as the instrument of the person who uses it, and neither statute asks whether a particular act fell within the authority that person conferred. That question, and the treatment of an act outside it, most closely belong to agency law and to the contractual law governing the parties’ agreements. Recognizing the validity of automated contracting and establishing attribution in a particular transaction are different inquiries. The latter requires evidence of the act, the relevant authority, and the relationships through which it occurred, potentially across multiple jurisdictions and legal frameworks. Those questions matter both to the firm using the system and to a counterparty relying on its actions.

Regulatory analysis likewise begins with the proposed financial activity. FINRA has expressly reminded members that its rules and the securities laws continue to apply when member firms use generative AI, including tools supplied by third parties.¹² Identifying the obligations governing a proposed product or service helps determine which firms can perform its functions and on what terms. A familiar requirement may need to be addressed differently where a system exercises discretion over a sequence of transactions. However, the FINRA notice does not specifically address a member’s obligations when its customer uses generative AI.

Counsel must distinguish uncertainty about the facts from uncertainty about the law, identify risks that can be allocated by agreement, and recognize mandatory rules and third-party rights that agreements cannot displace. That analysis determines product terms, the allocation of functions, and the undertakings a firm needs from others as both the financial instruments and the systems acting on them evolve.

4. PUTTING AGENTS TO WORK IN A COLLATERAL ARRANGEMENT

Consider a hypothetical asset manager seeking to support an investment fund’s derivatives trading by using tokenized money market fund shares held in its portfolio as collateral. A dealer sees an opportunity to expand its client offering, supported by a custodian and a platform that can process transfers beyond conventional hours. An “always-on” language model-based agent deployed by the manager would respond to margin communications, obtain valuations, and arrange permitted substitutions and transfers under an overall risk framework provided by

¹¹ Uniform Electronic Transactions Act §§ 2(6), 9(a), 14(1) (1999); id. § 2 cmt. (“While this Act proceeds on the paradigm that an electronic agent is capable of performing only within the technical strictures of its preset programming, it is conceivable that, within the useful life of this Act, electronic agents may be created with the ability to act autonomously, and not just automatically”). UETA has been enacted in forty-nine states, the District of Columbia, Puerto Rico, and the U.S. Virgin Islands; New York instead relies on its Electronic Signatures and Records Act, N.Y. State Tech. Law §§ 301–309, and on E-SIGN. See N.Y.C. Bar Ass’n, *Why New York Must Modernize Its Electronic Transactions Law* (accessed Sept. 13, 2026), <https://www.nycbar.org/issues-policy/advocacy-campaigns/ueta-esra-new-york-electronic-transactions-law/>. On actual and apparent authority, see Restatement (Third) of Agency §§ 2.01–2.03 (Am. L. Inst. 2006).

¹² FINRA, Regulatory Notice 24-09, *FINRA Reminds Members of Their Regulatory Obligations When Using Generative Artificial Intelligence and Large Language Models* (June 27, 2024), <https://www.finra.org/sites/default/files/2024-06/regulatory-notice-24-09.pdf>.

the manager. The commercial objective would be to support trading strategies while reducing the need to liquidate portfolio holdings for cash collateral, with related business opportunities for the dealer and service providers.

The hypothetical project begins with the financial arrangements. The parties must establish whether the money market fund shares are eligible collateral, what interest will be delivered, how it will be held and valued, and what makes delivery effective. Applicable margin and investment restrictions constrain those choices, and the fund's mandate must accommodate the activity. The dealer and custodian face questions about authority and performance even if they deploy no agents themselves. Each needs terms identifying the instructions it may accept and the conditions under which it will act on them.

The parties also have choices about the degree of discretion the service supports. A dealer might give standing consent to substitutions meeting agreed conditions or require separate consent for particular assets. Those choices affect which trading strategies the arrangement can support and when customers can use the service. Counsel must translate these choices into eligibility criteria, notice mechanics, and procedures for exceptions that make the proposed offering commercially workable. The system must then operate within the arrangement the parties adopt.

Looking at this hypothetical transaction through the lens of the six questions set out above, we find that:

Authority sets the limits of the delegation. Suppose the parties require the dealer's consent before a substitution involving the tokenized money market fund shares. The restriction appears in the agent's earlier context, but a later summary, produced when the system compacts its conversation history, may retain the general permission to use the shares while omitting the consent condition. The model then proposes a substitution without obtaining consent. The change in context can affect the proposal without changing the legal mandate.

A separate execution check requiring evidence of consent may block the proposal. If the transfer nevertheless occurs, its legal effect depends on the provisions in the relevant agreements and applicable law. The restriction may remain legally significant even though software failed to enforce it. Equally, an internal restriction cannot by itself settle the rights of a counterparty that accepted an instruction. The analysis must distinguish the firm's delegation, the instructions presented to the model, the software's permissions, and the basis on which another party was entitled to rely.




Figure 1. Parties and actions in the hypothetical collateral arrangement.

The authority and control arrangements must remain workable as the service changes. A third-party provider operating the manager's agentic system might replace a model or alter how the system summarizes information without changing the manager's objective. That change does not itself confer additional authority but may affect behavior under the existing mandate. The contractual terms between the parties should identify changes requiring notice or review and address how the parties will continue performing their obligations if the affected activity must pause pending that review.

Attribution asks whose act the transaction is and which participants bear responsibility. The fund holds the assets and owes the collateral obligation; its manager directs the activity under an asset management agreement; a third-party provider may operate the manager's agentic system (or components of the system); and a custodian or platform may control execution. Authority and attribution must be assessed across those relationships. A particular error may reflect a provider's configuration change, information supplied by the dealer, or a control another participant undertook to maintain. These possibilities inform the undertakings and remedies negotiated among the participants.

Organization determines which legal persons hold assets, operate systems, and bear obligations. The fund's governing documents and the manager's mandate must accommodate the delegated activity; a software service does not take their place. Where a provider runs the harness, the manager needs enforceable rights to change permissions, obtain records, or suspend activity. Owning code does not necessarily mean controlling its deployment. In a decentralized arrangement, identifying the relevant persons and governance powers becomes part of the same inquiry.

Verification reaches both the information used and the proposed performance. An authenticated margin demand can still be mistaken, including when generated by another participant's system. A transfer within an amount limit can still fail to deliver what the agreement requires. Software can independently verify calculations or compare a proposed transfer with an authoritative record, but its protection depends on the checks implemented and the information they use. A statement generated by an AI agent that the task is complete is not confirmation that the collateral obligation has been satisfied.

Settlement determines what interest has been delivered and with what legal effect. A platform's acknowledgment may establish that a transfer of collateral occurred without answering every question about whether it discharged the underlying obligation. Technical finality and legal recourse also require separate treatment. A transfer that cannot be rolled back may leave obligations of return, replacement, or compensation. The asset's legal character, the custody structure, and the governing terms all matter to that result.

Recourse begins while the commercial relationship continues. The next obligation to post margin may fall due before a disputed collateral substitution is resolved. The parties need procedures, agreed in their transaction and service documentation, for notification, suspension of further activity, and interim performance while preserving their rights. Ratification and ultimate loss allocation may require additional facts and legal analysis. Stopping an agent does not itself resolve the fund's outstanding obligations. Those procedures help preserve the trading relationship while the parties address the error.

Each of these six questions depends on evidence connecting the arrangement to what actually happened. A record of the model's conversation can show information supplied and actions proposed but may omit operations performed automatically by a connected service. A requested transfer may never execute. The participants consequently need records of the mandate in force, relevant inputs, execution checks, and the transaction recorded by the custodian or platform. Compaction of working context need not destroy the original record if it is separately

retained.¹³ Access, integrity, and retention arrangements determine whether the parties can reconstruct the event and establish their rights. The transaction and service documentation should address these requirements.

5. STRUCTURING PRODUCTS AND MARKETS FOR AUTONOMOUS ACTIVITY

The same questions arise beyond collateral. A proprietary trading firm might coordinate activity across venues; a custodian might offer new financing services; a decentralized finance (DeFi) developer might manage a “front-end” website that accepts instructions from users’ agents. Each offering depends on identifying the obligations and discretion of the participants whose systems must work together.

Cross-border activity makes the location and regulatory treatment of each function part of the commercial design. In the European Union, for example, an agent’s determination of trading parameters involving financial instruments can bring activity within MiFID II’s algorithmic-trading requirements. Payment mandates must also accommodate the applicable consent and authentication requirements under PSD2. These are activity-specific inquiries, including where decentralized infrastructure makes the relevant participants and connections harder to establish.¹⁴

For a tokenized securities project that intersects with one or more autonomous agentic systems, the starting questions concern the instrument, how it is held and transferred, and the rights investors acquire. Where is ownership recorded? What interest passes on a transfer, and through which intermediary, if any? Which restrictions govern transfers or use as collateral? The answers influence issuance terms, custody arrangements, and the markets in which the instrument can be used—ultimately determining how an agentic system can handle this type of asset. Existing work on tokenized collateral provides a foundation: ISDA’s guidance addresses approaches to enforceability opinions for arrangements involving tokenized securities or stablecoins under specified collateral documentation.¹⁵ Delegated selection or transfer must fit the rights and restrictions of the underlying product.

A platform intended for 24/7 trading needs valuation, liquidity, and settlement arrangements that work when an underlying market or payment service is closed. An agent can act only through the infrastructure available to it. The operating model may combine agents, conventional automation, and staffed coverage. Platform rules must address interruptions and the continuing obligations of participants whose services operate on different schedules.

Novel derivatives structures bring those dependencies into the product’s terms. A payment calculation may depend on a reference market with limited hours; a disruption may require an alternative valuation source. If an agent is to select that source, the parties must agree the permitted sources, the conditions for using them, and the right to challenge the calculation. Individual consent for every departure may limit the intended service. Discretion within an agreed method supports continuity but requires terms for review and adjustment. Those choices affect the product’s economics and the parties’ rights.

CahillNXT’s role in a given transaction with agentic elements extends from evaluating the proposed structure and its regulatory treatment through negotiation, documentation, and implementation. Tracing the system’s operation identifies the commitments a firm can offer and the undertakings it needs from issuers, counterparties, platforms, and providers.

¹³ Anthropic, *Compaction* (accessed Sept. 13, 2026) (explaining that applications may retain original messages even where earlier content is excluded from subsequent model processing), <https://platform.claude.com/docs/en/build-with-claude/compaction>. Whether relevant records are retained and sufficient for legal purposes is a separate, deployment-specific inquiry.

¹⁴ See ESMA, *Supervisory Briefing on Algorithmic Trading in the EU* ¶¶ 8–14 (Feb. 26, 2026), <https://www.esma.europa.eu/press-news/esma-news/esma-issues-supervisory-briefing-algorithmic-trading>; Directive (EU) 2015/2366 (PSD2) arts. 64, 97, <https://www.legislation.gov.uk/eudr/2015/2366/title/IV/2015-11-25/data.xht?view=snippet&wrap=true>. Application depends on the function, parties, and relevant scope provisions and exemptions; whether a mandate structure takes a payment outside of the payer-initiated authentication requirement depends on the function and parties.

¹⁵ ISDA, *ISDA Tokenized Collateral Guidance Note* (May 21, 2024) (ISDA describes the note as informing counsel’s approach to enforceability opinions for arrangements involving tokenized securities or stablecoins under certain ISDA collateral documentation), <https://www.isda.org/?p=1093149>.

Before launch, participants should produce an agreed mandate and limits; transaction and service terms governing reliance on agent-generated instructions; an allocation of execution checks and recordkeeping duties; and procedures for suspension, correction, and interim performance. Legal, business, technology, and risk and control teams must test whether those arrangements work together. For risk and control officers, the practical task is to connect each legal requirement to an identifiable control, a responsible participant, and evidence that the control operated. Disputes and insolvency analysis should inform that work before ordinary performance fails. The Annex to this article provides a checklist for applying these questions, whether the firm deploys an autonomous system itself or receives instructions generated by another participant's system.

6. MARKET PARTICIPANTS WILL HELP SHAPE THE LAW


The law of the machine economy will develop through the interaction of commercial practice, legislation, regulation, and judicial decisions. Firms contribute to this effort both through the systems they deploy and through the terms on which they accept autonomous activity from others. Understanding those systems enables counsel to identify consequential questions early, apply existing law, and develop workable arrangements where uncertainty remains. Machine Law organizes that inquiry around the products firms want to offer, the trading strategies they want to pursue, and the markets they intend to build.

Ad hoc private sector agreements and platform rules will continue to influence which activities firms can undertake and what they expect from one another. Supervisory practice and disputed transactions will test those expectations. We are supporting an industry-wide effort to develop a Model Agent Transaction Framework: proposed terms addressing the six questions examined here, designed to sit alongside existing trading, collateral, custody, and service agreements. The aim is to develop those terms through an open process with market participants, connecting legal mandates to execution controls, evidence, and procedures for errors. We invite clients and colleagues to contribute their experience.

This publication is provided by Cahill Gordon & Reindel LLP as a service to clients and colleagues. The information contained in this publication should not be construed as legal advice. A variety of AI-based tools were used during the research and production of this article.

CONTACT

To support a practical understanding of the issues discussed in this article, the CahillNXT team has developed a CLE program on how autonomous agents actually work, intended for internal and external counsel advising participants across financial markets. Using a collateral arrangement of the kind examined here, the program traces instructions and evolving context through proposed actions, execution controls, and records of performance. We welcome conversations with counsel and risk and control professionals about the program and the financial products, platforms, and transactions for which that understanding can inform structuring and implementation. For more information, please contact Lewis Rinaudo Cohen (partner) at 212.701.3758 or lrcohen@cahill.com, Sarah Weiyang Chen (partner) at 212.701.3759 or swchen@cahill.com, Jonathan Galea (partner) at 44.20.7947.9315 or jgalea@cahill.com, Gary Kalbaugh (partner) at 212.701.3505 or gkalbaugh@cahill.com, or Frank Weigand (partner) at 212.701.3890 or fweigand@cahill.com.



Lewis Rinaudo Cohen

Partner
212.701.3758
lrcohen@cahill.com



Sam Enzer

Partner
212.701.3125
senzer@cahill.com



Sarah Chen

Partner
212.701.3759
swchen@cahill.com



Jonathan Galea

Partner
44.20.7947.9315
jgalea@cahill.com



Gary Kalbaugh

Partner
212.701.3505
gkalbaugh@cahill.com



Frank Weigand

Partner
212.701.3890
fweigand@cahill.com

ANNEX: TRANSACTION CHECKLIST FOR ACTIVITY INVOLVING AUTONOMOUS SYSTEMS

This checklist identifies matters to consider and resolve when entering into transactions of value involving autonomous systems. It applies whether a firm deploys a system itself or encounters one deployed by a counterparty, customer, platform, or service provider. It can also be used to review existing relationships as agentic activity is introduced.

Apply the checklist to the particular transaction, product, and jurisdictions involved. For each relevant item, record the agreed position or unresolved issue, the responsible participant, and the supporting contractual provision, control, or evidence.

1. AUTHORITY: WHAT HAS BEEN AUTHORIZED?

- **Define the mandate.** Identify permitted activities, products, counterparties, venues, and accounts; transaction and aggregate exposure limits; and actions requiring additional consent or human approval. Distinguish authority to propose or negotiate from authority to commit, transfer, or discharge obligations.
- **Match permissions to authority.** Determine whether access to accounts, credentials, and execution tools permits actions beyond the mandate. Specify where limits will be enforced, including checks outside the language model, and who may override them.
- **Address further delegation.** Decide whether the system may engage other agents or service providers, what authority may pass to them, and who remains responsible for their activity.
- **Provide for change and revocation.** Establish who may amend the mandate, how affected participants receive notice, when changes take effect, and how outstanding instructions are treated.

2. ATTRIBUTION: WHOSE ACT IS IT?

- **Identify the represented party.** Establish the legal person on whose behalf each system acts and the basis for treating its instructions, communications, and commitments as that person's acts. Distinguish authentication of a credential from evidence of authority.
- **Set the terms of reliance.** Decide whether agent use must be disclosed, what assurances a receiving party requires, and when instructions will be accepted, questioned, or rejected—including where agent involvement is unknown.
- **Resolve conflicting expressions of intent.** Specify how signed agreements, account permissions, agent-generated messages, execution records, and subsequent confirmations interact. Address instructions that appear valid but exceed an internal mandate or conflict with other communications.

3. ORGANIZATION: WHICH LEGAL PERSONS HOLD ASSETS, DIRECT SYSTEMS, AND BEAR OBLIGATIONS?

- **Map the participants and their roles.** Identify the contracting parties, asset holders, system operators, custodians, and material service providers. Determine who can direct the system, alter its permissions, access assets, and answer for performance.
- **Map applicable law and regulatory duties.** Identify relevant jurisdictions, governing law, and applicable requirements concerning the product, activity, participant, and data. Determine which responsibilities can be performed by another participant and which remain with the firm.

- **Connect responsibilities to practical capabilities.** Confirm that each participant assigned a check, recordkeeping duty, or intervention role has the information, access, personnel, and contractual rights to perform it. Address dependencies on providers or infrastructure outside the parties' direct control.

4. VERIFICATION: HOW ARE IDENTITY AND PERFORMANCE ESTABLISHED?

- **Assign checks before execution.** Identify who verifies identity, authority, asset eligibility, available funds or collateral, and applicable restrictions. Specify the evidence required and the response to missing, stale, contradictory, or potentially manipulated information.
- **Preserve a usable transaction record.** Determine what must be retained to reconstruct the mandate, material instructions and inputs, approvals, checks, external actions, and resulting performance. Allocate access, retention, confidentiality, and production duties across participants.
- **Test and monitor relevant behavior.** Test how the arrangement responds to misleading external content, unavailable services, repeated or conflicting instructions, and interactions with other agents. Establish what changes to models, tools, permissions, or operating conditions require notice, renewed testing, or approval.

5. SETTLEMENT: WHAT IS DELIVERED, AND WITH WHAT LEGAL EFFECT?

- **Identify the asset and the promised performance.** Establish what rights an asset, token, or account entry represents and what constitutes delivery, payment, discharge, or creation or release of a security interest.
- **Specify completion conditions.** Determine the consents, eligibility checks, valuations, funding, and other conditions required before execution. Address differences among trading hours, clearing availability, settlement processes, and human support coverage.
- **Plan for incomplete or inconsistent execution.** Allocate the consequences of partial performance, duplicate instructions, delays, and discrepancies between systems. Establish what can be cancelled or reversed technically and what legal steps are needed to restore the parties' positions.

6. RECOURSE: HOW ARE ERRORS CORRECTED AND RIGHTS ENFORCED?

- **Establish intervention procedures.** Identify who can suspend activity, revoke access, contact counterparties, and address pending instructions. Distinguish stopping further activity from correcting transactions already executed.
- **Allocate losses and remedies.** Address unauthorized activity, execution errors, compromised credentials, inaccurate information, and service failures. Consider the interaction of transaction terms, service terms, liability limits, indemnities, and applicable law.
- **Provide for performance during a dispute.** Decide what happens to payment, delivery, margin, collateral, and other continuing obligations while an instruction, transfer, or transaction is contested. Establish escalation contacts and response times appropriate to the operating hours.

FINAL REVIEW: DO THE ARRANGEMENTS WORK TOGETHER?

Before launch or acceptance of agentic activity, legal, business, technology, and risk and control teams should walk through an ordinary transaction, an instruction exceeding authority, and a failure during execution. For each, confirm who acts, what prevents or detects an error, what evidence is available, and what happens next. Record unresolved issues and responsibility for resolving them.