



May 19, 2026

Categories:

[Agentic Commerce](#)

[Series](#)

[Publications](#)

Authors:

[Nilesh \(Neal\) Patel](#)

What Epic Reveals About Agentic Commerce

Platform terms, API rules, and private enforcement may shape agentic commerce disputes before courts or regulators do.

The next major fights in agentic commerce may not turn first on abstract AI doctrine. They may turn on contract stacks, anti-steering rules, platform discretion, and the private terms that decide who gets to transact, on what conditions, and at whose risk.

The Contract Layer Sits Above Everything Else

The next important disputes in agentic commerce may not be resolved first by courts or regulators. They may be resolved first by platform terms, API rules, marketplace agreements, payment-provider contracts, and the unilateral enforcement mechanisms that already govern who gets to transact and on what conditions. Earlier articles in this series mapped the technical, financial, and competitive architecture of agentic commerce: [identity and authentication](#), [delegated authority](#), [contract formation](#), [loss allocation](#), [payment infrastructure](#), [platform gatekeeping](#), [antitrust and surveillance pricing](#). Each of those layers runs on a set of rules. In practice, those rules are contracts.

The agent does not interact with a platform in the abstract. It interacts with a platform's API, and that API is governed by a developer agreement, acceptable-use rules, service commitments, indemnities, change-of-terms mechanics, and termination rights. The merchant does not interact with the platform in the abstract either. The merchant operates under a marketplace agreement, a payments processor agreement, a partner program agreement, and often a separate set of API terms governing automated access. Many of the disputes identified earlier in this series — about identity, authority, loss allocation, or coordinated exclusion — will be resolved first by a clause buried somewhere in that contract stack.

Two structural features matter most. First, these contracts are drafted entirely by the platform. Standard developer and

marketplace agreements are usually posted, not negotiated. Second, the platform usually reserves the right to change them unilaterally on whatever notice it considers sufficient. Taken together, those features mean that the operating law of agentic commerce is often a unilateral product, revised on one party's schedule. That looks less like conventional bargaining than private regulation delivered through contract.

Adhesion at Commercial Scale Is the Operative Problem

These agreements have the familiar features of adhesion contracts: one-sided drafting, take-it-or-leave-it presentation, and limited practical alternatives. The problem is that American contract doctrine often treats those concerns as less serious when both parties are businesses. The law tends to assume sophistication, alternatives, and an ability to walk away.

That assumption weakens quickly in agentic commerce. A small merchant building around Amazon's Selling Partner API likely cannot negotiate Amazon's terms and may not be able to leave without serious commercial harm. A developer building on Apple's App Store has even less leverage. An agent operator dependent on several large-platform APIs faces the same problem several times over. Formally, these are B2B contracts. In practice, many operate as non-negotiable access rules imposed by entities that control critical commercial channels.

That is why the better frame is not ordinary contract doctrine alone, but private regulation. Platform terms increasingly determine when transactions may occur, when access may be withdrawn, and what information merchants may provide to consumers and agents. That framing is contestable. Defenders of platform discretion will characterize these agreements as exercises of property, product-design, and freedom-of-contract rights. But dependence, unilateral modification, and the lack of meaningful exit make the regulatory frame the more accurate description of the disputes likely to matter.

Europe Has Built the Framework the U.S. Has Not

The European Union confronted this problem directly through the Platform-to-Business Regulation, which took effect in 2020. Its premise is straightforward: platform-business relationships can involve enough dependency and asymmetry to justify targeted fair-dealing rules.

For present purposes, the important features are practical. The P2B Regulation requires terms to be intelligible and continuously available. It requires at least fifteen days' advance notice for changes to terms (with longer periods where adaptation is required), thirty days' advance notice with a statement of reasons for full termination, objective grounds for restriction or suspension, statements of reasons for adverse actions at or before they take effect, and internal complaint-handling mechanisms. Changes implemented without the required notice are null and void. The Regulation also imposes disclosure duties around ranking, differentiated treatment, and certain contractual restraints.

The United States has nothing comparable. Federal proposals have stalled, and although state-level intervention is increasing in adjacent areas, no U.S. jurisdiction has enacted a general platform-merchant contract regime of the P2B type. That leaves U.S. merchants with weaker default protections and leaves counsel looking for other doctrinal tools. It also leaves the European framework as the clearest existing benchmark for a more balanced platform relationship.

The Clause Inventory That Determines Outcomes

The policy debate is one thing. The contracts currently governing these relationships are another, and they are more immediate. Eight recurring provisions allocate most of the risk in agentic commerce relationships: termination and suspension rights, unilateral change-of-terms mechanics, indemnification, ownership of transaction and agent-output data, audit rights, agent-identification requirements, scope-of-use restrictions, and overall liability allocation.

There is, however, a structural complication that matters as much as any single clause. In many deployments, the relevant rules are

not in one agreement but in several, and those agreements often conflict. A consumer's contract with an agent operator may authorize automated purchasing. The merchant's customer terms may permit only human-initiated transactions. The merchant's platform terms may prohibit unapproved automation. The payment processor may classify the transaction one way for chargeback purposes while the merchant's records classify it another. The AI provider may disclaim downstream responsibility altogether.

That layered-contract-conflict problem deserves more attention than it usually gets. Many disputes in agentic commerce will look like ordinary contract cases but will actually turn on which contract governs a particular failure and which party is left holding risk created by someone else's terms. The clause-by-clause discussion below is still useful, but it should be read as part of a broader mapping exercise across the full contract stack.

Consider a simple but realistic stack. A merchant may sell through Shopify, rely on Stripe for payment processing, use OpenAI or Anthropic for agent-facing decision support, and reach the customer through Apple or Google-controlled app environments. The consumer may separately authorize an outside shopping agent to search, compare, or even initiate a purchase on the consumer's behalf. If that transaction later fails, the practical legal question is not just what happened in the abstract. It is which contract governs the failure, which intermediary had authority to permit the relevant automation, and whose terms assigned the risk before anyone started arguing about fault.

One further caveat: the negotiation advice below assumes some leverage. For enterprise counterparties, the clause inventory is a negotiation roadmap. For small and mid-sized merchants facing standard terms, it is more likely to function as a risk-documentation tool and a basis for dependency-reduction planning.

Termination and Suspension Rights

Current platform agreements usually let the platform suspend or terminate access immediately, with little or no notice, broad definitions of cause, and no meaningful cure period. Amazon's

Selling Partner API terms, Apple's developer terms, and similar agreements across major platforms generally follow that pattern.

The asymmetry matters more than the precise wording. The platform can usually terminate on short notice or immediately; the merchant often must give lengthy notice and continue performing during the wind-down. Where the platform's definition of cause incorporates all published policies, and those policies can be revised unilaterally, the termination right becomes broader still.

The drafting response is straightforward: narrow cause definitions, require notice for non-emergency actions, provide an opportunity to cure, and preserve data access and export during wind-down. Those protections are not standard in posted terms, but they are the right baseline wherever leverage exists.

Unilateral Change-of-Terms Mechanics

Every major platform reserves a unilateral power to modify terms. The main differences are notice period and whether the platform distinguishes between material and non-material changes. Some terms take effect on posting. Others provide fourteen or thirty days for material changes, with the platform itself defining materiality.

The practical significance is visible in several public episodes. Meta's 2025 update to its WhatsApp Business terms, which restricted AI-provider access, is the clearest recent example in the agentic-commerce context. More broadly, Twitter's abrupt 2023 API access changes and Reddit's 2023 API pricing changes showed how quickly a platform can destroy an existing business model while staying within the formal bounds of its contract.

The lesson is simple: if the platform has reserved unilateral change authority and the merchant has no continuity protection, the merchant's reliance interest may be economically significant but legally fragile. For dominant or near-dominant platforms, antitrust or competition law may become relevant. For many others, it will not. The immediate problem is still contractual discretion.

That makes change-of-terms protection one of the most important negotiation points in any agentic-commerce deployment. At

minimum, merchants should seek real advance notice, a meaningful distinction between emergency and non-emergency changes, termination rights if the merchant does not accept a material change, and transition support where the merchant has built technical infrastructure around the platform's prior rules.

Indemnification Structures

The standard platform indemnity is usually one-way and heavily favors the platform. The merchant indemnifies the platform for claims arising from the merchant's products, use of the platform, breach of the agreement, or violation of third-party rights. The platform's own liability is capped at a nominal amount.

Agent-mediated commerce exposes the weakness in that structure. An agent operated by one party may act through another merchant's credentials on a platform controlled by a third party, causing harm to a fourth participant. Under standard terms, the party with the platform contract may owe indemnity even when it did not design, operate, or meaningfully control the agent that caused the problem.

That is the agent-specific drafting problem. Posted terms generally do not distinguish between harm caused by the merchant's own conduct and harm caused by a third-party agent acting within merchant-authorized parameters. Counsel should push for that distinction, together with the usual protections: carve-outs for platform negligence or breach, meaningful caps, and tighter causal language.

IP Ownership of Transaction Data and Agent Outputs

Agentic transactions generate several commercially valuable categories of data: the transaction record itself, logs showing how the agent evaluated options, performance information across repeated transactions, and derivatives useful for training, optimization, or competitive intelligence.

Current terms address these categories unevenly. Some providers allocate customer inputs and outputs relatively clearly while reserving rights to aggregated analytics. Others are vague about transaction data ownership or reserve broad platform-use rights that matter more in practice than formal title. In many agreements, the most important category — agent reasoning or decision-process data — is not clearly addressed at all.

That gap matters because practical control often follows technical custody. If the contract does not clearly allocate access and use rights, the party holding the logs may effectively control the most valuable byproducts of the relationship.

For agreements involving significant agent activity, counsel should require explicit treatment of raw transaction logs, decision-process records, derived analytics, and training-use derivatives. The agreement should also address anonymization standards, limits on use for competing services, and audit or verification rights around data-use restrictions.

Audit and Inspection Rights

Platform audit rights are typically broad, open-ended, and one-way. The platform may monitor or inspect the merchant's compliance without clear limits on frequency, scope, or use of findings. The merchant usually gets no reciprocal visibility.

Two concerns follow. The first is competitive intelligence. A platform auditing merchants that compete on the platform may gain information with value beyond compliance. The second is selective administration. An audit right exercised unevenly can become a competitive lever even if each individual audit is nominally defensible.

That does not mean audit rights are inherently suspect. Fraud prevention, security enforcement, and ecosystem stability are legitimate interests. The real question is whether audit power is administered symmetrically and whether findings are used only for compliance rather than for platform advantage.

The contractual response should be modest but real: defined scope, reasonable notice where possible, use limitations on audit findings, restrictions on use in competing operations, and reciprocal merchant audit rights for data-use and service-level commitments.

Agent Identification and Registration Requirements

Agent-identification and approved-automation requirements are likely to become more common across platform and API terms as platforms try to distinguish between ordinary user activity, permitted automation, and higher-risk agent-mediated transactions. In the near term, those rules may appear less as a unified legal category than as a patchwork of disclosure obligations, developer-policy conditions, API approval gates, and enforcement discretion.

The contractual significance is twofold. First, enforcement will usually occur through contract before technical standards are mature enough to do the work on their own. Second, identification rules can serve either legitimate trust-and-safety goals or competitive discrimination, depending on how they are designed and applied.

Counsel should focus on objective criteria, symmetric administration, and calibrated remedies. If the standard is undefined, administered selectively, or backed by immediate termination for minor deviations, it becomes a competitive control tool rather than a neutral disclosure regime.

Scope-of-Use and Prohibited-Automation Provisions

Most API agreements restrict how the API may be used. Some restrictions are familiar and legitimate. Others operate much more like competitive restraints. Terms prohibiting developers from building competing services, commercializing certain categories of data, or engaging in unapproved automated actions can become the core contractual weapon in a later dispute.

That is especially true in agentic commerce. A provision barring automated or programmatic activity unless expressly authorized in writing may look like an anti-abuse rule, but it also gives the platform discretion to decide which forms of agent activity count as legitimate and which do not. Once that discretion is paired with a termination right, the platform can decide which automation it will tolerate and on what commercial terms.

Anti-Steering, Price Communication, and Routing Restrictions

The Epic line of cases is most useful here, but only if kept in bounds. Epic did not establish that platform access terms are broadly unlawful, and it did not convert private platforms into regulated utilities. What it did suggest, at least under California's unfair-competition framework, is that a platform's contractual control over merchant communications can become legally vulnerable when it suppresses truthful information about alternative purchasing paths. The Supreme Court declined to stay the Ninth Circuit's mandate on May 6, 2026, leaving that framework operative while Apple's petition for certiorari is pending.

That is the point of the Epic reference in this article. Epic is not the whole story of agentic commerce, and it is not a universal template for every platform dispute. It is a useful preview of how private platform rules can begin to look less like ordinary product design and more like market-shaping restraints when they restrict communication, routing, or commercial alternatives that matter to downstream participants.

That matters in agentic commerce because the most consequential restrictions may not be the headline commission or API fee. They may be the quieter terms governing what merchants may tell consumers and their agents, what data may be exposed to competing agents, whether lower direct prices may be disclosed, whether total-price comparisons may be displayed, whether direct-buy APIs may be made available, and whether alternative checkout paths are technically allowed but commercially disabled.

This is where the distinction between formal permission and functional availability matters. A platform may argue that merchants remain free to use direct channels, alternative payment rails, competing agents, or off-platform links. But if the contract pairs that nominal permission with warning screens, ranking penalties, data throttling, opaque approval requirements, link-format restrictions, off-platform commissions, or similar friction, the practical effect may be to preserve the platform's preferred path while claiming to permit alternatives.

That is the more useful Epic analogy for Article 8. The point is not that every steering restriction is unlawful. It is that one-sided platform terms become more vulnerable when they suppress truthful comparative information, obstruct price competition, penalize lawful alternatives, or preserve platform control through commercially disabling friction. In agentic commerce, those restraints may be written not only for humans but for agents: machine-readable routing restrictions, limits on price-comparison data, direct-buy API prohibitions, and constraints on what an authorized agent may display or execute.

Merchant-Platform Liability Allocation

Taken together, the provisions above place most of the legal and operational risk of agentic commerce on the merchant side. The platform's liability is capped. The merchant's indemnity may be broad or uncapped. The platform may terminate or revise terms unilaterally. It may inspect broadly, control or limit access to data, and determine what forms of agent activity are authorized.

That cumulative imbalance is the real contract problem in agentic commerce. Individual clauses matter, but their combined effect matters more. The platform is not just pricing access. It is defining the conditions under which the merchant may communicate, automate, route, and retain customer relationships.

What Counsel Should Expect to See in the Interim

A comprehensive U.S. platform-contract statute is not imminent. In the near term, counsel should expect a mix of state legislation,

contract doctrine, unfair-trade-practices theories, and competition law to do the work imperfectly.

The implied covenant of good faith and fair dealing remains one likely tool, particularly where a platform uses an express discretion to destroy the central benefit of the bargain. Unconscionability remains available in theory, though courts are often reluctant to use it aggressively in commercial settings. That makes state unfair-trade-practices statutes especially important.

The California Unfair Competition Law and Massachusetts Chapter 93A are obvious examples. These statutes are likely to matter most where the challenged term does more than allocate economic risk and instead suppresses truthful merchant communications, interferes with meaningful comparison, penalizes off-platform alternatives, or obscures the information environment in which consumers and agents make choices. That is where the Epic anti-steering logic has practical value. Even when a platform's broader business model survives antitrust scrutiny, specific restraints on communication, routing, or comparison may still attract unfair-competition or consumer-protection challenge.

Other claims will also appear. Depending on the facts, platform terms may be attacked through borrowed-law theories under state consumer-protection statutes, through state antitrust law, through privacy or fee-disclosure obligations, or through public-policy arguments against enforcement of restraints that suppress lawful competition or truthful information.

Europe remains the practical ceiling for multinational deployments. The P2B Regulation, Digital Markets Act, and Digital Services Act together impose transparency and fair-dealing constraints beyond what U.S. law currently requires. For clients operating across jurisdictions, that makes the European standard the clearest operational benchmark even where U.S. law remains underdeveloped.

The strategic point is simple. Where leverage exists, the contract architecture should be negotiated directly. Where leverage does not exist, merchants should assume that any single platform's continued cooperation is a planning risk and build accordingly.

Action Items

- **Map the full contract stack for any material agentic-commerce workflow, not just the merchant's platform agreement.** Platform terms, payment terms, customer-facing terms, and AI-provider terms should be reviewed together so conflicts in authorization, chargeback treatment, automation limits, and data ownership are identified before they become disputes.
- **Audit each relevant platform agreement against the core clause inventory:** termination, change-of-terms, indemnity, data rights, audit rights, agent-identification rules, scope-of-use restrictions, and liability caps. The goal is not just a redline list. It is a written record of where the merchant is structurally exposed and which risks require operational mitigation if they cannot be negotiated away.
- **Treat change-of-terms power as a material business risk.** Where a client depends on continued API access, transition planning, migration capacity, data export, and multi-platform redundancy are not optional housekeeping. They are basic resilience measures.
- **Preserve merchant communication and routing rights wherever possible.** Review whether the agreement restricts lower-price disclosures, direct-purchase links, machine-readable checkout instructions, total-price comparisons, fee disclosures, access by authorized outside agents, or communications that help customers move to direct channels. Those provisions may matter more than they first appear because they determine whether the merchant can maintain a meaningful customer relationship outside the platform's preferred path.
- **Insist on agent-specific limits in indemnity and data provisions where negotiation is possible.** Merchants should not absorb open-ended liability for third-party agent conduct without clear causal limits, and they should not discover too late that the most commercially valuable decision-process data is effectively controlled by the platform because the

contract never addressed it.

- **For multinational deployments, draft toward the European standard once** rather than trying to maintain one architecture for Europe and a weaker one for the United States. In many cases, the stricter standard will be the cheaper operational choice.

Please contact the author if you have questions or comments on this article. You can also reach out to any member of the firm's [Data, Digital Assets & Technology](#) practice for help navigating AI deployments, integrations, and governance.

This article was prepared with the assistance of generative AI tools. The analysis, conclusions, and legal positions are the author's own.

Agentic Commerce Series

- [Part 1 — AI Agents Are Starting to Act Inside the Transaction, and Commerce Law Is Not Ready](#)
- [Part 2 — The Identity Problem: Authentication, Fraud, and Who's Actually Buying](#)
- [Part 3 — The Authority Problem: When Does an Authorized Agent Become an Unauthorized Buyer?](#)
- [Part 4 — Contracting by Agent: When the Agent Clicks, Who Assents?](#)
- [Part 5 — When the AI Purchase Goes Wrong, Who Pays, and Who Can Prove It?](#)
- [Part 6 — The Gatekeepers of Agentic Commerce](#)
- [Part 6 \(A\) — The Payment Infrastructure Layer: How Network Rules Already Govern Agentic Commerce](#)
- [Part 7 — Agentic Commerce, Price Control, and the New Antitrust Questions](#)
- [Part 8 — What Epic Reveals About Agentic Commerce](#)

© 2026 FBT Gibbons LLP. May be considered attorney advertising in some jurisdictions.