Agentic AI Is Forcing Procurement to Rethink What It Means to Trust a Vendor

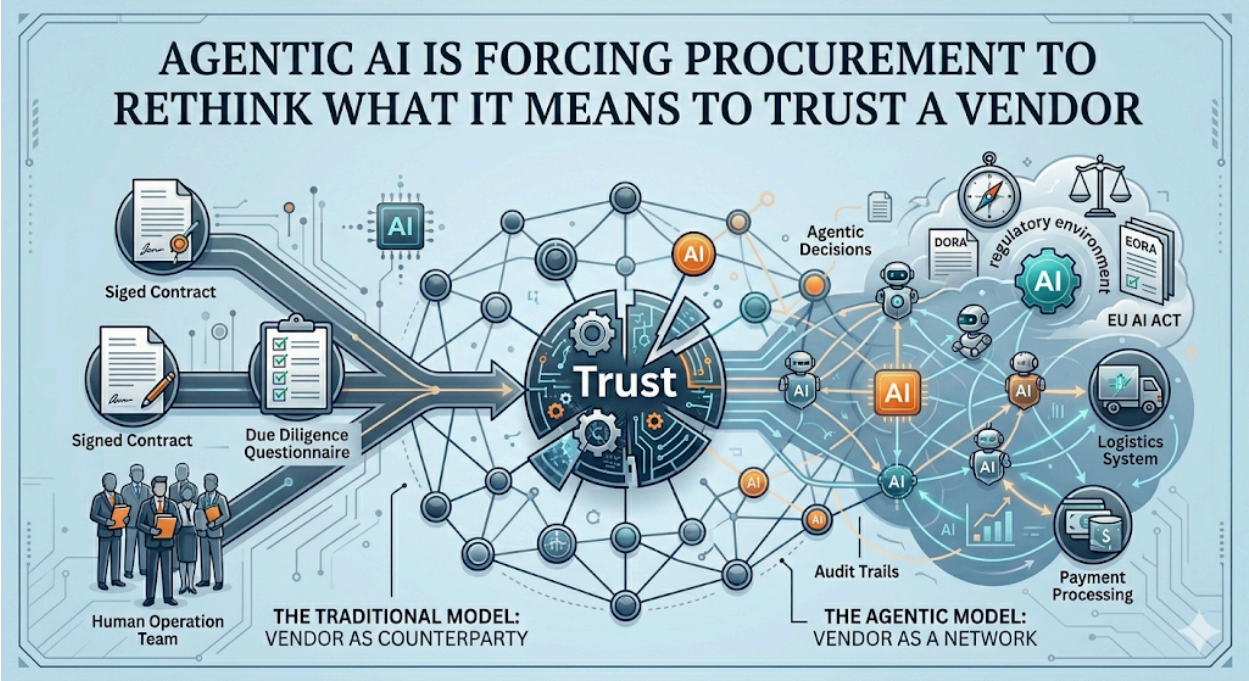

Procurement, as a discipline, has spent the last twenty years organised around a simple proposition. A vendor is a counterparty: a known organisation, with named people, signed contracts, defined deliverables, and a human operations team that can be called when something goes wrong. The entire toolkit of third-party risk management, security questionnaires, due diligence checks, exit clauses, the annual review cycle, was built around that proposition.

That proposition is starting to fray. Across financial services, supply chain, and customer operations, vendors are no longer the only thing on the other side of the relationship. Their autonomous systems are too, increasingly making decisions that affect your customers, your data, and your regulatory exposure, sometimes by coordinating with autonomous systems your direct vendor never explicitly disclosed.

For procurement leaders, risk officers, and the board committees that oversee both, this is the quiet shift that hasn't yet been named. It will be named, probably in regulatory letters before it is named in trade press. The companies that get ahead of it will be the ones whose vendor frameworks can absorb a structural change that the old frameworks were not designed for.

The moment a vendor stops being a counterparty

The traditional vendor model rests on a clean assumption: the vendor's organisation is the counterparty. You evaluate them, you contract with them, and your obligations begin and end at their corporate boundary.

In contemporary AI-mediated workflows, that boundary has become porous. A vendor's system may invoke other vendors' systems autonomously in service of completing the workflow you bought. A payment processing partner may route a transaction through a fraud-analysis service operated by a third firm. A logistics platform may negotiate capacity with a carrier-network system you did not contract with directly. The chain of automated decisions that produces your customer's outcome may include three or four organisations the original procurement evaluation never touched.

The contract you signed assumed a counterparty. The reality of the workflow assumes a network. This is not a minor mismatch. It is the central reason third-party governance frameworks are about to be revised across multiple jurisdictions.

The failure modes are predictable, not exotic

Four patterns recur in early production incidents involving autonomous systems coordinating across organisations.

Counterparty drift. The vendor you contracted with last year is using more autonomous integrations this year than last. The boundary of what their system touches in your workflow has expanded without an explicit conversation, because the vendor improved their product. The contract describes the original arrangement; the system has moved.

Hallucinations across organisational lines. When one autonomous system produces confident but incorrect output, the next system in the chain treats that output as input. By the time a human reviews the result, the chain looks plausible and is wrong underneath. Each individual system acted within its scope. The aggregate error is no one's clear responsibility under the standard contract.

Audit trails that stop at the company boundary. Your vendor's system logs everything it does. Their downstream partner's system logs everything it does. Neither log gives you a coherent timeline of what happened to your customer's request, because the logs were never designed to be reconstructed across organisational lines.

Reviewer fatigue at cross-organisational checkpoints. Human-in-the-loop oversight at cross-organisational boundaries tends to weaken faster than internal oversight, because the reviewer is being asked to approve outputs from systems they did not configure and cannot inspect. Approvals become reflexive. The mechanism intended to catch errors gradually stops doing so.

None of these failure modes is exotic. All of them are predictable. The companies treating them as predictable, and designing the governance to absorb them in advance, are the companies whose autonomous vendor relationships will hold up under stress.

What production-grade governance looks like in 2026

A pattern is visible across the organisations leading this area. Three choices recur in the procurement and risk frameworks that are surviving regulatory scrutiny.

Counterparty mapping is a deliverable, not an afterthought. Procurement teams are starting to require that vendors disclose, in writing, which other autonomous systems they coordinate with by default during the course of delivering the contracted service. This is appearing in contract language as a duty to disclose material changes in counterparty composition over time.

Audit trails are negotiated, not assumed. The leading procurement teams now treat audit access across the workflow as a primary contractual term, not a fallback. The right to reconstruct what happened across an automated chain, with timestamps from each participating system, has moved from "useful" to "required" in regulated sectors.

Human override is named, not implied. Contracts increasingly specify which actions in the workflow are gated by explicit human approval, who that human is, and what record is kept of their decision. Implied oversight, where someone "is supposed to be watching" but the watching is undefined, is being phased out of new agreements in financial services and healthcare.

The regulatory direction is converging

The companies developing these practices are responding to a regulatory environment that is moving in a consistent direction across major jurisdictions.

The EU Digital Operational Resilience Act (DORA) treats third-party ICT risk in financial services as a board-level concern and requires contractual provisions for exit, audit, and incident reporting that extend to material subcontractors. The first major enforcement cycle for the EU AI Act in 2026 applies particular scrutiny to high-risk decisions made by autonomous systems and demands documentation and traceability that span the full decision chain. In the United Kingdom, the Critical Third Parties regime signals that systemically important technology providers will eventually carry resilience and exit-planning obligations directly. In the United States, the FTC's Operation AI Comply has pursued enforcement actions tied to deceptive AI marketing, and federal coordination of AI governance has tightened through 2025.

The vocabulary differs across these instruments. The substantive demands are converging. Documented controls, traceable decisions, and named human accountability for high-risk actions are the baseline that regulated industries will be asked to meet, regardless of where the AI sits in the workflow.

The unresolved questions are honest

Three problems still lack clean answers, and procurement teams that pretend they are solved will be caught out.

The accountability question across organisations is not settled. When an automated chain involving systems from three companies produces an outcome that harms a customer, established liability frameworks were not written for that distribution of responsibility. Courts and regulators are testing how existing rules apply, and the cases being decided in 2026 will shape vendor agreements for the rest of the decade.

The trust framework problem is real. Vendor onboarding has historically taken months because trust accumulated slowly through human relationships, reference checks, and pilot evaluations. When systems coordinate at machine speed, that pacing breaks. The trust framework has to be designed in advance, and most organisations have not built the muscle to do that work.

The vendor concentration risk is growing quietly. As more inter-organisational coordination consolidates around a small number of foundation model providers and platform layers, the systemic risk of correlated failures rises. UK regulators have signalled that this may eventually fall under the Critical Third Parties regime. Most procurement teams are not yet thinking about correlation risk in their vendor portfolios.

What this means for procurement and risk leaders

The shift to AI-mediated vendor relationships does not require procurement to be reinvented. It does require some of its core practices to be revised.

The most useful questions to be raising in the next vendor review or audit committee meeting are not about model performance or feature roadmaps. They are structural. Which of our vendor relationships involve autonomous systems coordinating with other autonomous systems outside our direct contract? What is our visibility into those interactions, and what is our recourse when something goes wrong? Which of our contracts need to be updated to require disclosure of material changes in counterparty composition? Where are our audit trails complete, and where do they stop at a company boundary we did not control?

These are not questions the procurement function has historically had to answer. They are the questions that will increasingly define whether a procurement function is fit for purpose.

The deeper reframing is straightforward. The vendor of the next decade is not only an organisation. It is also a network of autonomous systems acting on the organisation's behalf, sometimes invoking other organisations' systems in the process. Procurement, third-party risk, and board oversight will all evolve to account for that. The companies that lead the evolution will be the ones that started asking the right questions before regulators required it.