

The technology infrastructure for the agent-to-agent economy is here. The economy is not, and that gap is the opportunity.


80% of chief executives expect AI to overhaul how the work gets done; only 17% expect it to change what customers buy. Agentic systems transact on their owner’s behalf, and systems from different companies can now find each other, agree a price and settle without a person present. Agent-to-agent commerce is barely discussed outside payments and reaches almost no board agenda, yet the exposure is two-sided: a demand assumption that may stop holding, and an advantage to the first company that becomes legible to another’s systems.




FIGURE 1 · WHAT IT IS

The same purchase, before and after

HOW A PURCHASE HAPPENS NOW



HOW IT HAPPENS WHEN BOTH SIDES RUN AGENTS



WHY IT MATTERS

01 Your buyer may be an agent The party deciding to buy may not be a person.	02 Agents do the finding Discovery runs through agents rather than search results.	03 Billable hours reprice first Revenue charged by the hour is what agents reach first.	04 Liability is unassigned No law or standard says who bears a loss an agent causes.
--	--	---	--

Both bands describe the same purchase. The fourth column is the one nobody has settled.

HOW FAST IT WAS BUILT

Published standards now govern every part of an agentic transaction, all agreed inside thirteen months by authors who gave the code away.

Google donated the agent-to-agent protocol to the Linux Foundation in June 2025 and its agent payments protocol to the FIDO Alliance in April 2026, and settlement went to the x402 Foundation, operational since 14 July. Every strategic plan now in force predates all three donations.



13 months

THE BUILD

From the first agent protocol donated to a live settlement body.



3 of 4

WHAT IS SETTLED

Identity, spending limits and settlement. Proof is not.



\$28k a day

THE TRAFFIC

What the main open agent payment standard settled on a day in March.




2027

THE QUESTION

When a buyer starts pricing agent exposure in diligence.

FIGURE 2 · GOVERNANCE

When each part of an agentic transaction was standardized



Unit: the date each piece was handed to an independent body. Sources: Linux Foundation; FIDO Alliance; Stripe; Visa; HM Treasury.



WHY THE INCUMBENTS GAVE IT AWAY

A company keeps a protocol proprietary when it expects to charge rent on it, and donates it when it wants the volume more. The settlement body has 17 premier members, among them Visa, Mastercard, American Express, Stripe, Amazon Web Services and Google, several of whom compete directly with one another.

WHAT A TRANSACTION NEEDS

Identity, spending limits and settlement have owners. Proof, which decides who bears a loss, has none.

Buy anything online and four conditions are satisfied without your noticing: the merchant accepts the card is yours, your issuer has authorised the amount, the money moves irreversibly, and a receipt exists that either side can produce in a dispute.

FIGURE 3 · ANATOMY

The four requirements, what supplies each one, and who governs it

WHAT A TRANSACTION NEEDS	FOR A PERSON	FOR AN AGENT	WHO GOVERNS IT
Identity	A passport, a signature	A signed agent card	Linux Foundation
A spending limit	A card limit somebody set	A limit the owner approved	FIDO Alliance
Settlement	A bank	The x402 standard	x402 Foundation
Proof of delivery	An invoice	Nothing agreed	Nobody

The fourth row is the open one. No standard governs proof of delivery, and no jurisdiction has allocated the loss when an agent gets a purchase wrong.

Systems belonging to different companies have settled orders without a person since the 1970s, through Electronic Data Interchange and the networks built on it, but in every case the two firms negotiated a contract and exchanged credentials before letting the systems talk. The contract was the proof.

What changed this year is that two agents with no prior relationship can find each other, establish identity and settle, and without that contract the evidentiary question arises for the first time. The Linux Foundation, the FIDO Alliance and the x402 Foundation now govern three of the four requirements between them; ERC-8004’s validation registry, written to supply the fourth, had no confirmed mainnet deployment as of May 2026.

 WHY PROOF IS A COMMERCIAL RISK RATHER THAN A TECHNICAL ONE

Without a record that survives a dispute a loss cannot be allocated, and a general counsel who cannot allocate a loss will not authorise an agent to commit the company to anything material.

.... HOW MUCH IS REAL


The identity, spending-limit and settlement standards are live. Adoption is negligible, and some of the measured traffic is manufactured.

Three governed standards against \$28,000 of daily settlement is the strongest counter to everything that follows, and it is also what leaves room to act before agent exposure is priced.

FIGURE 4 · USAGE

Agents registered, against agents that publish a way to reach them

173,000 AGENTS REGISTERED ACROSS THREE NETWORKS



Each square is roughly 1,730 registered agents. Source: Xiong et al., arXiv 2606.26028, 8 July 2026.

Agent payments totalled about \$28,000 on a day in March across 131,000 transactions and Artemis, the analytics firm tracking them, judged roughly half to be self-dealing or wash trading. A June 2026 study of ERC-8004, the leading agent registry, found 173,000 agents registered across Ethereum, BNB Smart Chain and Base, coordinated fake reviewing by a majority of raters, and 3% to 15% publishing a valid service endpoint.

Yet Visa, Mastercard, American Express, Amazon Web Services and Google do not attach their names to a settlement foundation for a market they expect to stay at twenty cents.

WHAT WOULD CHANGE THIS READING

Watch the share of registered agents that publish a working service endpoint. If single digits become a majority, take-up has begun and 2027 is late.


🕒 THE HORIZON

Companies are signing three- to seven-year commitments against liability rules no jurisdiction has written.

The one live consultation is British, closes on 6 October 2026, and asks whether liability for unauthorised payments needs updating for agents. The American Senate draft published in June does not allocate the loss, and the European Union withdrew its AI liability directive last year.

FIGURE 5 · DURATION

What is being committed, against when a rule could exist



Britain’s consultation on agent liability closes 6 October 2026 and the rulemaking that follows runs through 2027. No other jurisdiction has opened one.

Every strategic plan, supply contract and software renewal above assumes a counterparty who is a person, and that assumption is written equally into a listed industrial, a family business and a sponsor-owned platform alike.

The International Monetary Fund named the gap in April, writing that liability regimes assume human intent and direct causation, both of which become ambiguous once an agent decides alone. Thirty-five percent of organisations have no agentic strategy and eleven percent have agents in production, so contractual commitment runs ahead of both the rulebook and the readiness to use it.

🕒 WHAT THIS DOES NOT DEPEND ON

None of this depends on agents arriving quickly. Only one condition has to hold: that contracts signed this quarter outlast the date on which anybody is told who pays when an agent gets a purchase wrong.


WHAT IS PRICED

Markets have already repriced AI that does the work more cheaply.

J.P. Morgan put the software drawdown at 34% over twelve months and about \$2 trillion of market value, the largest outside a recession in thirty years, with a further \$1 trillion off the index in the week after 28 January. The transmission was headcount: customers cut staff and renewed fewer seats.

FIGURE 6 · DISLOCATION

What the price did, against what the company disclosed



Unit: percent. 52-week ranges to 26 August 2026; company figures are reported revenue growth in US dollars. Sources: exchange data; company results.

Goldman Sachs put the medium-term revenue growth implied by software multiples as falling from 15% to 20% down to 5% to 10%. Accenture closed 18% lower in a session on consulting growth of 1% in local currency, and Concentrix and Teleperformance each closed about 11% down on the same June day, on a fear Bloomberg summarised as the sector becoming uninvestible.

Private markets answered by not transacting: North American technology buyout volume fell 78% year over year in the first half of 2026, to \$10.9bn across 133 deals. Bank of America put \$15bn of US broker commissions at risk, the portion it classed as low complexity, doing the work of 20,000 to 30,000 independent agents.

“ Every one of those moves priced the same thing: a company whose costs are people, meeting software that does the work more cheaply.

Q WHAT IS NOT MODELLED

No sell-side model builds agentic buying into a price target, and no annual report names it as a risk.

Every number on the previous page is AI replacing employees. An agent arriving as the buyer is a different question, and no analyst, filing or dataset has touched it.

FIGURE 7 · ABSENCE

What a search of research, filings and market data returns

AGENTIC SOFTWARE AS THE BUYER, AS AT AUGUST 2026

Research notes from major banks modelling it into a price target	None
Annual reports naming it as a disclosed risk	None
Measured data on the share of transactions that agents execute	None
Trading sessions in which it has moved prices	One

Sources: searches of sell-side research, 10-K risk factors and market data, August 2026; Benzinga; Bloomberg; CoinDesk.

There has been exactly one exception. On 22 February 2026 Citrini Research published a scenario piece written from an imagined 2028, in which machine-to-machine commerce makes the 2% to 3% card interchange rate an obvious target, and the next session Visa, Mastercard, American Express and DoorDash fell between 4% and 7%. It was a thought exercise, and the move faded within days.

That is the entire record: one memo, one session, one sector. No research since has modelled a company’s revenue against a buyer that is not a person, and no company has told its shareholders it is exposed to one.

Every adoption forecast published this year is therefore a forecast about the allocation of liability rather than about commercial demand, because the missing evidentiary requirement sets the pace.

THE DEFENSIBLE VERSION OF THIS CLAIM

Agent-to-agent commerce has been priced exactly once, for one session, in payments, and it faded. Everything else priced in 2026 was AI replacing employees.


ACROSS INDUSTRIES

Four business models are exposed. Markets have repriced only the one whose costs are people.

One column below is what 2026 already repriced; the other is what an agent arriving as the buyer would expose, and none of that has been priced.

FIGURE 8 · EXPOSURE

Priced on labour, against exposed when the buyer is an agent



Camel bars are indicative of 2026 repricing already observed. Every teal bar is hollow because no market has priced any of it.

An intermediary earns its position from friction a person experiences: not knowing the market, or lacking the time to compare every counterparty on price and terms. An agent experiences neither. Insurance and freight brokerage, wholesale distribution, travel and real estate services have the most to lose and the least of it priced.

Businesses won through search face a different version, because being findable by a person and being legible to an agent are separate capabilities and almost every company has budgeted only for the first. Asana told investors in March that language-model changes to search and paid media were costing it two points of growth, and would continue to through its next financial year.

WHAT THE TWO LISTS SHOW

Markets have repriced the businesses whose costs are people. They have not touched the businesses whose revenue depends on a person being present at the decision.


HOW BIG IT IS

The exposed models are worth trillions a year. No index provider publishes what share of any market they represent.

US wholesale trade, the business of companies buying from companies, recorded \$11.6 trillion of sales in the last full census, making it the largest sector of the American economy by sales, ahead of manufacturing and retail.

FIGURE 9 · MAGNITUDE

What the exposed models are worth, where a published figure exists



No index provider publishes what share of any market sits in these models.

Unit: US dollars a year, log scale. Wholesale and travel are gross transaction value; advertising is media spend rather than agency fees. Sources: US Census Bureau; Bain; Phocuswright (projected); dentsu (forecast).

The exposure so far has been the sell side, and that \$11.6 trillion is spend rather than revenue. For any company that makes or moves physical goods, procurement is where agentic systems arrive first, because the buying is already systematised through purchase orders and approved-supplier lists, and the counterparties are known.

Index providers classify by sector and intermediation is not one, so these businesses are scattered across financials, industrials, communication services and real estate with nobody aggregating them.

WHERE THE ADVANTAGE SITS


The first company in a category an agent can transact with is the one it transacts with, and on the buy side the capability compounds: an agent comparing every supplier continuously extracts the terms a scale buyer gets, without the scale.

Boards built AI oversight in a single year, and pointed it at deployment rather than demand.

Fifty-four percent of the S&P 100 disclosed board-level AI oversight in their 2025 proxies and 28% disclosed both oversight and a policy, while 23% have assessed how AI might disrupt the business, or from where.

FIGURE 10 · BOARDS

What board AI oversight looked like a year apart



Unit: percent of the 80 Fortune 100 companies that had filed proxies and 10-Ks by 31 July 2025. Source: EY Center for Board Matters.

- 01

What share of our revenue needs a person present at the decision?

Not what share is AI-exposed, but what share depends on a human buyer experiencing the friction our position removes. Both the share and the timetable are knowable.
- 02

Are we still viable in five years, and who has modelled that?

Not whether we win, but whether a business whose revenue depends on a person being present at the decision has a five-year model at all. Nobody has been asked to produce it.
- 03

What would count as a triggering event, and who is watching for it?

On 29 July 2026 the Financial Accounting Standards Board opened a project on how often goodwill is tested, with staff proposing that the annual test give way to one triggered by events, putting far more weight on management’s judgment about when the environment changed.
- 04

Who in this room can test the answers?

Eight percent of directors say their board has strong AI expertise, the lowest score of any skill area in the survey, and the narrower question is whether anyone present can challenge a demand forecast that assumes an agent on the other side.

Chief executives concede AI changes how the work is done. Far fewer concede it changes what customers buy.

The 17% of chief executives who expect a significant change to their customer base compares with 39% who said so during the digital shift, and the leading quarterly survey of finance chiefs does not ask about the revenue model.

FIGURE 11 · CONCESSION

What chief executives concede, and what they do not



Unit: percent of respondents. Source: Gartner, n=469 chief executives and senior executives, published 23 April 2026.

- 01

Is our growth strategy still valid if the buyer is an agentic system?

Every growth plan assumes a customer who can be persuaded, values convenience and does not compare exhaustively. An agent has none of those properties.
- 02

Are agentic systems in our plan as a cost line or as a demand assumption?

Fifty-six percent of 4,454 chief executives surveyed at the end of 2025 reported no significant financial benefit from AI. A cost-line answer is defensible, but an unexamined demand assumption is not.
- 03

What would it cost to make us legible to another company’s agent, and whose budget is that?

Seven percent of organisations told Cloudera and Harvard Business Review their data was completely ready for AI, and no credible costing of agent legibility exists in public, so whoever produces the first sets the terms.
- 04

What can an agent commit us to, and what stops it?


No standard governs how one company’s systems authenticate another’s agent, and no convention covers what it may agree to. Gartner found 15% of application leaders even considering a fully autonomous agent.

Boards and chief executives can wait. A sponsor who must sell a company bought in 2021 cannot.

Private equity held 33,575 unsold companies at the end of June, and Bain puts holding periods at exit near seven years against five to six in the prior decade. The sale therefore comes due after agent exposure reaches diligence.

FIGURE 12 · CONCENTRATION

Who the secondary market serves, by deal count and by money



Unit: percent of all secondary-market transactions in 2025. Source: Jefferies, Global Secondary Market Review.

- 01

Which holdings need a person present at the decision, and on what timetable?

A portfolio answer built company by company, and the share is knowable this quarter.
- 02

What would it cost to make each company legible to another company’s agent?

That number belongs in your model before it appears in a buyer’s, and no published costing exists, so whoever builds the first one sets the terms.
- 03

Who prices agent exposure first?

Representation and warranty underwriters are already scrutinising AI representations, and deal counsel expect AI-specific exclusions to follow. If a secondary buyer gets there next, the mark moves without you.
- 04

Which assets become reachable by buyers who do not exist yet?

A company an agent can buy from is reachable by an acquirer who never enters your process.



IMPLICATION FOR MID-MARKET PE

A \$40m EBITDA company has no traded comparable to signal when this question arrives, so its carrying value rests on an operating assumption that nothing prices.

✧ THE ARGUMENT

The agent-to-agent economy has no incumbents yet, and the companies that move first will set the terms.

80% of chief executives expect AI to change how the work gets done, and those programmes are funded. The second question, whether AI changes what customers buy, has the larger commercial consequence and almost no budget. A company that answers that question now learns which of its revenue assumes a person on the other side of the decision, what legibility to another company’s agent would cost, and where its own procurement gets cheaper. Those are three answers a competitor will not have for another year.

A YEAR-OLD ANSWER IS A WRONG ONE

Commission a diagnostic of the enterprise and its market against a twelve-month horizon and refresh it quarterly. These standards were agreed in thirteen months, faster than any payments infrastructure of the previous decade.

01 · REVENUE

Which revenue assumes a person on the other side of the decision, and on what timetable?

02 · SPEND

Which spend assumes the same, and what would it cost to be legible to another company’s agent?

03 · LANDSCAPE

Who in our market is already legible to another company’s agent, and how would we know?

SOURCES

Standards. Linux Foundation, Agent2Agent, 23 Jun 2025; x402 Foundation launch, 14 Jul 2026 · Google, AP2 to the FIDO Alliance, 28 Apr 2026 · Visa, Trusted Agent Protocol, 14 Oct 2025 · Stripe and Tempo, Machine Payments Protocol, 18 Mar 2026 · ERC-8004; ERC-8183, both draft.
Regulation. HM Treasury, Modernising Payment Services Regulation, closing 6 Oct 2026 · AI AGENT Act discussion draft, Sen. Warner, 29 Jun 2026 · IMF Notes 2026/004 · FASB agenda decision, 29 Jul 2026 · EU AI Liability Directive, withdrawn.
Usage. CoinDesk and Artemis, 11 Mar 2026 · Xiong et al., arXiv 2606.26028.

Markets. J.P. Morgan · Reuters · Goldman Sachs · Bloomberg · BofA Global Research, 3 Mar 2026 · Mergermarket · Citrini Research, 22 Feb 2026 · exchange data · company results and calls.
Surveys. Gartner, n=469 · Glass Lewis, S&P 100 · NACD, n=211 · EY Center for Board Matters · Corporate Board Member · PwC 29th Global CEO Survey, n=4,454 · Deloitte Tech Trends 2026 · Deloitte Digital B2B · Cloudera and Harvard Business Review · Bain · PitchBook · Jefferies.

ABOUT



Alex Kruzel writes on artificial intelligence and private capital, and on the agent-to-agent economy taking shape underneath both. She is founder and chief executive of Telesto, which brings AI-driven market intelligence to private equity sponsors and to the healthcare and industrial companies they own, after two decades advising global businesses at Deloitte and as president of Market Strategy Group. She is the author of *The Courage to Continue* and serves on the board of the Chicago Council on Global Affairs. Her work starts where this column does, at the demand assumption nobody has been asked to defend.

Correspondence: alex.kruzel@telestostrategy.com