



A THOUGHTWORKS POSITION PAPER · SIBOS
2026, MIAMI BEACH · 28 SEP - 1 OCT

When the buyer is an algorithm

**Re-engineering the payments
architecture for AI-driven economies:
on both sides of the transaction.**

Authors:

Rav Hayer, Head of UK & Ireland; Head of BFSI, EMEA


Sid Sengupta, Head of BFSI, UK & Ireland

Nathan Hilt, Financial Services & Payments Leader



Design. Engineering. AI.

Summary: algorithms do not care about your checkout button	3
1. Two rebuilds, one control plane	6
2. Five things that break, and what replaces them	8
4. The standards land grab is already under way	13
5. How big, really, and what could slow it down	14
6. Two geographies, two clocks	16
7. Three imperatives, and the order to do them in	18
8. Eight questions for your next payments steering group	19
Who we are: Thoughtworks	20
If the buyer is an algorithm, what does it read about you?	21
Sources & notes	22



Summary: algorithms do not care about your checkout button

For 30 years, the payments industry has optimized one thing: persuading a human eye to press a button on a glass screen. Card art, one-click wallets, biometric friction-reducers, A/B-tested checkout flows: all of it is addressed to an audience that is leaving the room.

When a machine takes the buyer's seat, brand attachment stops being a moat and becomes a rounding error. An agent does not recognize your logo; it reads your terms. It does not abandon a basket; it resolves for a cheaper path. And it doesn't transact once: it transacts continuously, at a density and a price point your core banking platform was never sized for.

At Sibos 2026 in Miami, under the theme Digital finance for AI-driven economies, that's the uncomfortable structural fact behind every panel on speed, intelligence and trust. This paper sets out what has to change in the architecture, and in what order.

1. The demand layer has inverted, not upgraded. Commerce is moving from search-then-checkout to intent-then-execution. That's a change of buyer, not a change of channel, and it invalidates the assumptions underneath conversion, loyalty and fraud.
2. Selection becomes a solved-for outcome. Agents compare interchange, FX spread, settlement speed, reward yield and refund terms programmatically. If your product's terms are not machine-legible in milliseconds, you are not losing the comparison: you are absent from it.
3. Authority, not identity, is the hard problem. Rails assume a human is present at the moment of purchase. Agents act on delegated intent, hours or days later. The gap is closed with cryptographically signed mandates, scoped and time-bound, not with another one-time passcode.
4. There's a second market behind the first. Every autonomous purchase sits on a cascade of machine-to-machine payments: inference, tool calls, data queries, compute. Card economics cannot price a \$0.00007 transaction. Something else will, and it won't be a card network by default.
5. Trust is now an infrastructure product, not a compliance artifact. Speed without verifiable governance is an operational risk with a marketing budget. The institutions that win issue trust as a programmable credential, and sell it.

The numbers, and why they disagree

- **\$3-5T:** global transaction value agents could orchestrate by 2030.[1]
- **~90%:** of B2B buying expected to be agent-mediated by 2028.[2]
- **\$0.50-0.80:** per-transaction cost floor on a credit rail, before interchange.[3]
- **4,700%:** year-on-year growth in generative-AI referral traffic to US retail.[4]

Sources are listed in full at the end of this paper. Forecasts differ by an order of magnitude because they measure different things. See figure 08.

What this paper covers

- Two rebuilds, one control plane (Fig 01-02)
- Five things that break, and what replaces them (Fig 03-05)
- The market behind the market: the AI stack (Fig 06)
- The standards land grab is already under way (Fig 07)
- How big, really, and what could slow it down (Fig 08)
- Two geographies, two clocks (Fig 09)
- Three imperatives, and the order to do them in (Fig 10)
- Eight questions for your steering group (Checklist)





1. Two rebuilds, one control plane

Most agentic commerce strategy decks stop at the buyer-facing agent. That's half the rebuild, and the easier half.

The front end is visible and commercially urgent: autonomous agents discovering, evaluating, negotiating and executing purchases on a human's behalf. The back end is where the unit economics are actually decided. Behind each autonomous purchase runs a high-frequency cascade of machine-to-machine transactions: paying for model inference, querying tool APIs, licensing data, routing compute. That traffic has to be metered, priced, cleared and reconciled in real time, at values four to six orders of magnitude below a normal card payment.


These two problems are usually owned by different parts of a bank, or by nobody. They resolve to the same architecture: identity and authority that a machine can prove, terms a machine can read, limits a machine cannot exceed, and settlement that clears continuously rather than in batches.

The front end: machines spending human money. An agent does not browse your storefront. It queries it, ranks it against 10 others, and buys, or skips you without ever rendering a page.

FRONT END: AGENTIC COMMERCE	BACK END: COMMERCE FOR THE AI STACK
Machines spending human money • Intent to execution A goal replaces a basket, the click disappears • Algorithmic rail selection Fees, FX, yield and speed compared per line item • Delegated authority Signed, scoped mandates replace human-present checks • Identity retention Merchant and issuer context survives the agent layer	The stack paying for itself • Sub-cent payments Tool calls, inference and data priced per request • Metered billing Per token, per call, per GPU-second, not per month • Agent-to-agent settlement No human in the loop, at machine frequency • Continuous clearing Ledgers that never close for the night

ONE CONTROL PLANE HAS TO SERVE BOTH SIDES			
Machine-legible endpoints	Cryptographic mandates	Usage metering & sub-cent clearing	Policy-as-code & real-time

The buyer-facing rebuild and the machine-settlement rebuild are commonly treated as separate programs. They depend on the same four capabilities, which is the strongest argument for funding them as one.



Three intermediaries, every relationship. Customer context stops at the dashed line, unless the payment carries identity across it.






2. Five things that break, and what replaces them

Human payment systems run on latency and persuasion. Agentic systems run on throughput and arithmetic. Each of the following is already observable in production somewhere; none is fully solved anywhere.

What a human sees



Card art. A trust badge. A button that 30 years of design effort went into, for three seconds of human attention.

Everything on the left is invisible to the buyer that now matters. A single missing field on the right removes you from the set entirely.

What an agent sees

```
{
  "apr_pct": 21.9,
  "interchange_bps": 165,
  "fx_markup_bps": 45,
  "settlement": "T+2",
  "reward_yield_bps": 175,
  "refund_window_days": 30,
  "mandate_support": false
}
```

→ Ranked 9 of 11 on net yield
→ Excluded: cannot accept a mandate

No brand attachment. No hesitation. No second look.

01. Breaks: brand-led selection. An agent buying a flight or replenishing inventory evaluates payment options programmatically: interchange, FX spread, settlement speed, net reward yield, refund rights, per line item. In card-heavy markets it will route to whichever credential maximizes net yield, and consumer attachment to your brand does not enter the calculation.

Build: machine-legible commercial terms. Fee schedules, reward rules, eligibility, dispute rights and settlement times published as structured, queryable endpoints with versioning and SLAs. Treat the terms API as a distribution channel, because that's now what it is.

02. Breaks: the human-present assumption. CVV entry, fingerprints and SMS passcodes all assume the buyer is at the keyboard. An agent acts on delegated intent, often days after the human set the goal, sometimes across dozens of merchants the human never named.

Build: signed, scoped mandates. Cryptographically signed authority carrying merchant allowlists, amount caps, time bounds, revocation and an audit trail, verifiable by the merchant and the network, not re-asserted by the agent.

03. Breaks: your customer relationship. If a generalist model becomes the permanent proxy between buyer and seller, merchants become interchangeable fulfillment endpoints and issuers lose the context that makes servicing, credit decisions and loyalty work. Intent and identity get absorbed into someone else's agent layer.


Build: payment as an identity-carrying protocol. Bind merchant and issuer identity into the transaction and the whole post-purchase lifecycle: receipts, returns, warranty, disputes, re-order. That way the relationship survives disintermediation of the shop window.

04. Breaks: your core banking throughput. Humans transact linearly. An agent evaluating hundreds of permutations triggers thousands of quotes, micro-authorizations, balance checks and fraud evaluations in seconds. Instant rails, including FedNow, SEPA Instant and Faster Payments, can settle it; batch-oriented cores, nightly reconciliation and per-call fraud stacks cannot absorb it.

Build: burst-tolerant, idempotent services. Event-driven ledgers, idempotency by design, quote/reserve/commit separation, rate-shaping at the agent boundary, and fraud scoring that costs microseconds rather than a vendor call.

05. Breaks: fixed-fee economics. Agents pay for inference, tool calls via MCP, data queries and compute: transactions worth fractions of a cent, millions of times a day. A rail with a per-item cost floor near \$0.50 cannot price a \$0.00007 call. The margin is negative before the value is delivered.

Build: metered micro-settlement. Streaming, usage-based metering with net-out settlement, pre-authorized spend sessions, and rails whose marginal cost is below the value of a single call. This is where stablecoin and ledger-native settlement stop being ideological and start being arithmetic.



The question is no longer “is this the cardholder?” It’s “did the cardholder authorize this, to this agent, for this merchant, up to this amount, until when, and can you prove it later?”

That's a different product. Authentication answers a question about a person at an instant. Agentic commerce needs a durable, revocable, auditable grant of authority that outlives the moment it was created, and that a counterparty can verify without trusting the agent that presents it.

Today: human present

INTENT + PAYMENT, SAME INSTANT


$$\Delta = 0s$$

The check works because the person is there to be checked: a CVV, a fingerprint, a passcode. Presence does the work, and presence is what disappears.

What an agent sees

HUMAN SIGNS

AGENT PAYS


$$\Delta = 3d\ 14h$$

THE AUTHORIZATION DEFICIT

MANDATE: WHAT CLOSES THE GAP

```
payee_scope allowlisted merchants / categories
amount_cap per item, per day, aggregate
valid_until expiry + instant revocation
signature verifiable by merchant and rail
```



Human-present checks collapse the gap between intent and payment to zero. Delegated execution reopens it, and what closes the gap again is the mandate.




3. The market behind the market: commerce for the AI stack

Agentic commerce decides how machines spend human money. The AI stack has a separate problem: it has to pay for itself, continuously, in fractions of a cent.

Every autonomous purchase is the visible tip of a much larger transaction count. An agent resolving one instruction may make thousands of billable calls: model inference, retrieval over a vector store, a tool invoked through the Model Context Protocol, a fraud score, an FX quote, a seat map, a compute reservation. Each has a marginal cost. Today, most of that cost is absorbed inside monthly platform invoices, because nobody can economically bill for a single call.

That's a temporary state, and it's the largest unclaimed position in payments. Card economics carry a per-item cost floor in the region of \$0.50 to \$0.80 on credit transactions, plus interchange of roughly 1.5-3.5%.[3] Against a \$0.00007 tool call, the rail costs four orders of magnitude more than the thing being bought. The traffic will therefore settle somewhere else: on metered ledgers, pre-authorized spend sessions, and stablecoin-native rails built for exactly this shape of traffic. Coinbase's x402 protocol, which revives the dormant HTTP 402 status code as a payment negotiation channel, has already processed hundreds of millions of agent transactions and now sits under Linux Foundation stewardship.[5]

The strategic point is not which rail wins. It's that a high-volume, high-margin clearing business is being built in the open, and incumbents are largely absent from it.



» 4. The standards land grab is already under way

Between 2025 and 2026, the industry did not converge on one agentic payments standard. It produced a stack of them, layer by layer, each with a different sponsor and a different theory of who should be trusted.

Model context and agent-to-agent protocols handle discovery and tool access. Checkout protocols standardize how an order reaches a merchant. A trust layer carries proof of what the human actually authorized. A settlement layer moves the money, through scoped card credentials or through ledger-native rails. Most of these specifications are young: versions below 1.0, pilots rather than production at network scale. And the identity of the agent itself is still not solved by any of them.

Two conclusions follow, and they're less contradictory than they look. Do not bet the architecture on a single protocol winning. Do make sure your rails can speak all four layers, because the cost of fluency is an integration and the cost of illiteracy is exclusion from the selection matrix.

THE AGENTIC PAYMENTS STACK, MID-2026: NO SETTLED WINNER AT ANY LAYER

Reasoning & tools

How does it find and call things?

MCP

A2A

Checkout & order

How is the order actually placed?

ACP: OpenAI + Stripe

UCP

Trust & intent proof

Who authorized what, and can you prove it?

AP2: signed mandates

Visa TAP

Web Bot Auth

THE LAYER BANKS SHOULD OWN

Settlement

How does the money move?

Agentic Tokens

scoped card credential

x402

stablecoin/HTTP

MPP

streamed micropay

FedNow, SEPA Inst, FPS

instant fiat clearing

Agent identity remains unsolved at every layer: treat it as a design constraint, not a dependency.




5. How big, really, and what could slow it down

Published forecasts for agentic commerce span two orders of magnitude. That's not analyst incompetence; it's a definitional disagreement about what counts.

The narrow measures count only transactions where an agent took meaningful autonomous action. The broad measures count the whole orchestrated flow: the purchase, plus the logistics, payment and pricing decisions the agent influenced along the way. Both are defensible. Quoting them interchangeably is not, and any board paper that does will fall apart under scrutiny.

Read them as scope, not as precision. Restricted to comparable definitions, a consensus does exist: by 2030, somewhere between 10% and 25% of US e-commerce will involve agents in a meaningful capacity.[6] The debate has moved from whether to how fast.

THE FORECASTS ARE NOT MEASURING THE SAME THING



Four things that would slow this down

Agent identity is unsolved. None of the current specifications establishes durable, portable identity for the agent itself. Merchants are being asked to trust a signature without a settled way to know who is holding the pen.

Liability is unallocated. When an agent buys the wrong thing within a valid mandate, the loss sits somewhere between the human, the agent operator, the merchant and the issuer. Until scheme rules and regulation answer that, risk teams will cap exposure, and caps suppress volume.

The new rails are attracting adversaries. Security researchers have already published systematic attacks on machine-payment protocols, including free-riding against x402 deployments.[7] Machine-speed payments imply machine-speed fraud, and controls tuned to human tempo will not see it.

Regulators are tightening delegation, not loosening it. In Europe, delegating strong customer authentication to a wallet or gateway is explicitly treated as outsourcing: the delegating institution keeps the liability and must retain audit rights.[8] Delegating authority to an agent will be read through the same lens.


Design for provable delegation from the start; retrofitting it is a multi-quarter program.

Every one of these is an argument for building the trust layer first, and for building it as a product you can price, not a control you have to absorb.

⌚ 6. Two geographies, two clocks

Agentic payments are maturing along two structurally different tracks. In the US, commercial standards set the pace. In Europe and the UK, statutory dates do.

The practical consequence for any institution operating in both: the US roadmap is driven by whichever scheme, platform or protocol acquires distribution first, and can change in a quarter. The European roadmap is anchored to fixed compliance windows that are already published. That makes it slower, and considerably easier to plan against. Build the capability once; sequence the rollout by clock.



US milestones are commercial and can be re-dated by one vendor announcement. European and UK milestones are statutory: published, dated, unforgiving.

Dimension	United States	European Union & UK
What sets the pace	Commercial standards and platform distribution. Card networks, model providers and processors are defining the interfaces.	Statutory deadlines. PSD3/PSR, the Instant Payments Regulation and eIDAS 2.0 are converging on one trust infrastructure.
Dominant credential	Scoped network tokens and shared payment tokens bound to a specific agent, merchant and amount.	Account-to-account authority: commercial VRPs, Pay-by-Bank, plus a state-issued digital identity wallet.
Settlement infrastructure	Network tokenization extensions, FedNow, commercial card APIs, and ledger-native rails for machine traffic.	SEPA Instant, UK Faster Payments, open banking APIs, with a digital euro still on a legislative track.
Trust anchor	Scheme and issuer token frameworks; agent identity asserted at the merchant edge and verified against a directory.	Strong customer authentication evolving toward wallet-based, phishing-resistant factors, with delegation treated as outsourcing.
Where value concentrates first	Autonomous procurement, SaaS and license renewal, reward-yield optimization on consumer cards.	Travel and retail agents executing bank transfers inside pre-agreed limits; recurring and variable billing.
The main risk	Betting on a protocol that loses, or waiting for a winner and arriving without an integration.	Treating fixed compliance dates as the ceiling of ambition, and shipping conformance instead of capability.

— 7. Three imperatives, and the — order to do them in

Sibos 2026 is a conversation about moving from pilots to governed, execution-ready infrastructure. Here is the shortest credible path.

One: make your products legible to machines. Fee schedules, reward rules, eligibility criteria, settlement times, refund rights and merchant terms have to exist as structured, versioned, queryable endpoints with performance guarantees. An agent that cannot evaluate your credit product inside its decision window does not rank it lower; it excludes it. Agent engine optimization is the new distribution, and unlike search rankings it's deterministic: you can read the criteria and satisfy them.

Two: treat cryptographic trust as a product line. Issue programmable credentials that bind identity, authority, spend limits, merchant scope, expiry and revocation to a named agent, and make them verifiable by counterparties without trusting the agent. This is the layer where regulated institutions have an advantage no protocol can copy: licenses, balance sheets, dispute machinery and a supervisory relationship. Price it accordingly.

Three: become the clearinghouse for the AI stack. Extend beyond merchant acquiring into banking-as-a-service for machine workloads: token-billing engines, low-latency settlement ledgers, spend sessions with pre-authorized caps, and yield-aware liquidity routing for high-frequency traffic. The AI stack currently settles its own accounts with platform

invoices and crypto rails, because nobody regulated has offered it anything better.

The sequence matters. Legibility without trust produces a fast route to unbounded liability. Trust without legibility produces an excellent product nobody can find. Clearing without either is a compliance incident with a roadmap.

NEXT 90 DAYS: PROVE IT	NEXT 4 QUARTERS: SHIP IT	18-24 MONTHS: SELL IT
• Publish one product's full terms as a machine-readable endpoint. One. This quarter. • Instrument agent traffic: can you tell an agent from a human in your logs today? • Meter one internal AI workload per call, not per month. • Name an accountable owner for agentic commerce. It is currently nobody.	• Issue scoped, revocable mandates in one live corridor. • Separate quote, reserve and commit in the authorization path; make it idempotent. • Sub-cent metering with net-out settlement, live on internal traffic first. • Fraud scoring at machine tempo; rate-shaping at the agent boundary.	• Trust as a product: sell verifiable authority to agent operators and merchants. • Clearing for the AI stack: billing engine, spend sessions, liquidity routing. • Own the post-purchase lifecycle: receipts, returns, disputes, re-order. • Policy-as-code control plane over every agent, internal and third-party.

MEASURE: TIME-TO-QUOTE FOR AN AGENT • COST PER MACHINE TRANSACTION • % OF TERMS MACHINE-READABLE • REVOCATION LATENCY



8. Eight questions for your next payments steering group

If more than three answers are “not yet,” the roadmap is a slide, not a plan.

1. Can an agent retrieve our full pricing, fees and refund terms programmatically, without a human, in under a second?
2. Can we distinguish agent traffic from human traffic in production today, and do we price the two differently?
3. Can a customer grant an agent authority that is scoped, capped, time-bound and instantly revocable, and can a counterparty verify it?

4. If an agent transacts inside a valid mandate and the customer disputes it, do we know who bears the loss?
5. What happens to our core when authorization volume rises by 100x for 11 seconds?
6. Can we bill a counterparty for a single \$0.0001 API call and make money on it?
7. Are we in at least two protocol pilots at different layers of the stack: trust and settlement, not just checkout?
8. Who, by name, owns agentic commerce P&L in this institution?

Who we are: Thoughtworks

Thoughtworks is a global technology consultancy that integrates strategy, design and engineering to help financial institutions rebuild for an AI-driven economy. We're usually brought in when the architecture, not the ambition, is the constraint.

Our AI/works and Agent/works offerings provide the governed runtimes, policy-as-code orchestration, usage metering and real-time control planes needed to run autonomous agents and AI-stack workloads safely at production scale. In payments, that means machine-legible product endpoints, cryptographic mandate frameworks, burst-tolerant authorization paths and metered settlement: the difference between an agentic pilot and a business with defensible unit economics.

If the buyer is an algorithm, what does it read about you?

We're running working sessions at Sibos on machine-legible product design, mandate architecture and settlement for the AI stack. Bring one product line and we'll trace its agentic path from discovery to settlement: where it gets excluded, where the authority breaks, and what it costs to fix.

Closing thought: no card number is typed. No form is filled. The agent presents authority a counterparty can verify, which is the product most institutions have not built yet.

Dates: 28 Sep - 1 Oct 2026

Theme: Digital finance for AI-driven economies

[Arrange a session.](#)



Sources & notes

1. McKinsey & Company: agentic commerce could orchestrate as much as \$1 trillion in US B2C retail revenue and \$3-5 trillion globally by 2030 (research reported October 2025).
2. Gartner: approximately 90% of B2B buying expected to be mediated by AI agents by 2028, representing more than \$15 trillion of spend (November 2025).
3. Per-transaction cost on credit rails of roughly \$0.50-0.80, based on Federal Reserve Bank of Kansas City estimates cited in industry analysis, alongside interchange in the region of 1.5-3.5%.
4. Adobe Analytics: generative-AI referral traffic to US retail sites grew roughly 4,700% year on year between July 2024 and July 2025.
5. x402 (originated by Coinbase) uses the HTTP 402 status code as a payment negotiation channel; it has processed hundreds of millions of agent transactions and moved to Linux Foundation stewardship in 2026. Related machine-payment specifications include AP2 (Google, v0.2 April 2026), ACP (OpenAI and Stripe), UCP, Visa's Trusted Agent Protocol, Mastercard's Agentic Tokens under Agent Pay, and the Machine Payments Protocol (Stripe and Tempo, mainnet March 2026).
6. Comparable-scope forecasts for 2030 US e-commerce: Morgan Stanley 10-20% (\$190-385bn), Bain 15-25% (\$300-500bn). Wider-scope estimates: Juniper Research \$1.5 trillion global transaction value by 2030 (from roughly \$8 billion in 2026); eMarketer approximately \$144 billion by 2029.

7. Published security analyses of x402 and related agentic payment protocols, including systematic free-riding attacks (2026 preprints).
8. EU and UK regulatory dates: Verification of Payee mandatory for euro instant credit transfers from 9 October 2025 under the Instant Payments Regulation; PSD3/PSR politically agreed 27 November 2025 with final compromise texts published April 2026, Official Journal publication expected in the second half of 2026 and most obligations applying approximately 21 months later; EU member states required to make a certified European Digital Identity Wallet available by November 2026; UK Payments Initiative commercial VRP scheme announced 2 June 2026. Delegation of strong customer authentication to third parties is explicitly classified as outsourcing, with the delegating institution retaining liability and audit rights.

Transaction values used in the cover figure and figure 06 are illustrative and rounded to show orders of magnitude, not vendor pricing. Forecast figures are quoted with the scope on which they were published; readers should not aggregate across scopes. Protocol status is as of mid-2026 and moving quickly; verify before citing.

Rav Hayer is Head of UK & Ireland and Head of BFSI, EMEA at Thoughtworks. Sid Sengupta is Head of BFSI, UK & Ireland. Nathan Hilt is Financial Services & Payments Leader. © 2026 Thoughtworks. This paper is for information only and is not legal, regulatory or investment advice.

For more than 30 years, Thoughtworks has helped shape how the world builds technology.

Today, by harnessing AI to scale what is possible, guided by human judgment and experience, we aim to be the force behind the world's most consequential technology transformations.

We combine our engineering and design expertise with AI to modernize systems, unlock data and create AI-powered products.

AI/works™ is our agentic development platform and Agent/works™ is our enterprise platform for governing and running autonomous AI agents. Together they bring decades of our engineering know-how to every stage of AI software delivery, from development through to deployment and governance.

thoughtworks.com