OBENAN RESEARCH BRIEFING

Prepaid Local Services

A Practical Wedge for Agentic Commerce

DATE

May 2026

CLASSIFICATION

Public research briefing

About this briefing

A public research briefing from Obenan. It develops a single argument: prepaid local services may be one of the first practical wedges from agentic discovery into agentic payment, because they expose a merchant-side commitment problem that payment infrastructure alone is not designed to address.

PUBLISHED BY	Obenan
DATE	May 2026
CLASSIFICATION	Public research briefing.
SCOPE	Public, primary, and well-attributed sources. No private partner conversations. No implementation-sensitive mechanics. No claim of integration, partnership, or program participation with any payment network, scheme, processor, or commerce-protocol body.
READING	Five takeaways in two minutes. Executive summary in five. Full briefing in fifteen.

For the reader with two minutes.

01 **Agentic commerce is solving authorization, not commitment.** The networks, the model providers, and the platforms have built credible work around registered agents, network tokens, signed user mandates, and verified intent. That work governs who is allowed to pay. It does not govern whether the merchant on the other side can honor the specific thing being paid for.

02 **A local service is not a SKU. It is a promise.** Time-bound, capacity-bound, location-bound, operationally volatile, often policy-heavy, and rarely represented by one structured catalog. The agent that treats a salon appointment like a product on a feed will, on average, get the answer it deserves.

03 **The wedge is bounded, not broad.** Not all local commerce is ready for agents. The category that is ready is narrower and well defined: prepayable, reservable, ticketable, or deposit-backed services with a single, named, time-stamped commitment object that an agent can verify before money moves.

04 **Two pilot lanes, not one.** A clean payment-readiness lane (parking, ticketed experiences, classes) shows that the rails work end to end. A merchant-truth lane (beauty, personal services, appointment-led wellness) shows why an upstream readiness layer earns its place in the stack. The two lanes answer different questions and both need answering.

05 **Obenan sits upstream of payment execution.** No money movement, no tokenization, no authorization, no checkout, no acquiring, no settlement. The work is merchant-side: identity, eligibility, freshness, validation, evidence, and committability before the payment layer takes over.

CONTENTS

The merchant side of an agent transaction is the part that is currently underbuilt.

Agentic commerce has moved from concept into working infrastructure. Payment networks, model providers, and platform companies have invested heavily in the parts of the stack closest to the transaction itself: registered agents, network tokens, signed user mandates, verified intent, and authenticated payment execution. Those efforts address an important share of the trust problem.

A second layer is less developed. Before an agent commits user funds, it must know whether the merchant on the other side is real, present, available, and able to honor the specific commitment being requested. For products with structured catalogs, this is largely solved. For local services it is not. The thing being purchased is rarely a static SKU; it is a time-bound promise made by a small operator working with limited tooling.

This briefing argues that prepaid local services may be one of the first practical wedges from agentic discovery into agentic payment. Not because all local commerce will move through agents, but because a defined subset of local services already produces the object agents need: a bounded commitment. A parking bay, a timed museum ticket, a class seat, a spa appointment, a restaurant deposit, a pickup order, a rental window. Each is digital enough to pay for remotely, local enough to require operational truth, and bounded enough for an agent to ask one clean question before money moves.

Obenan is not a payment company. It does not move money, tokenize cards, authorize payments, act as a PSP, act as an acquirer, build checkout, or settle transactions. Its work sits upstream of all of those: it helps physical merchants become safely discoverable and committable by AI agents before payment execution begins.

Two layers, two roles. Payment networks govern trust at authorization. The merchant layer must govern truth at readiness. A clean agentic ecosystem will need both, and prepaid local services are the category where the distinction stops being theoretical.

The body develops the argument across eleven sections, with a methodology note, a conservative evidence base, and an explicit non-claims page at the end.

Networks are solving authorization. Merchant readiness is the question they were not designed to answer.

The infrastructure being built around agentic commerce is moving fast.

Public payment-network and platform programs describe models in which agents are registered, payments are tokenized, user intent is verified, and consent is governed in line with established network controls. Public draft specifications for verifiable intent describe layered credential chains that bind agents and users to specific commercial actions. Other payment networks and platform companies have published parallel architectures using agent-specific tokens, authenticated user instructions, and commerce signals. Public Agentic Commerce Protocol primitives focus on structured product feeds, checkout flows, and shared payment tokens. Universal Commerce Protocol and AP2 publications describe cryptographically signed mandates as proof of user instructions on AI surfaces.

These efforts share a center of gravity. They focus on the user, the agent, the credential, the mandate, and the payment instrument. They organize around a single question: when an agent acts on behalf of a user, can the network and the merchant trust that the action is authorized, in scope, and recoverable in the event of dispute.

> A second question, less visible, has historically been answered offline by the merchant. When the agent reaches the point of commitment, can the merchant on the other side actually fulfill the promise being purchased.

For e-commerce, this is largely encoded in product feeds, inventory systems, and structured catalogs that update fast enough for most consumer expectations. For local services, the answer is distributed across booking platforms, scheduling tools, sole-trader notebooks, and operational decisions made hour by hour.

Agentic payment infrastructure can verify the user, the agent, and the payment mandate. It is not designed to verify whether a sole-trader hair stylist has free time at three on a Thursday afternoon, whether the slot listed on a museum's third-party reseller is still actually open, or whether a yoga class capped at fifteen has space for one more. The merchant-side question lives outside the credential chain, upstream of payment, downstream of discovery, and currently underbuilt.

A local service is not a static SKU. It is a time-bound operational commitment.

For agentic commerce, the most important property of a purchasable object is whether its truth can be checked before money moves.

A digital product carries that truth in a SKU, a price, a stock level, and a return policy. A local service carries it in something messier.

A local service is **time-bound**. It exists in a specific window. A class at six in the evening is not the same product as a class at noon, and neither is a product at all once the window has passed.

A local service is **capacity-bound**. Most categories operate against a small, hard limit: seats in a class, bays in a garage, chairs in a salon, tables at the second seating.

A local service is **location-bound**. Fulfillment happens in one physical place, often by one person. The labor pool that performs the service is rarely fungible across locations.

A local service is **operationally volatile**. Staff call out. Equipment breaks. Weather changes. Bookings cluster. Cancellations create gaps that may or may not get refilled. The merchant's understanding of its own state shifts continuously throughout the day.

A local service is often **policy-heavy**. Cancellation rules, deposit rules, late-arrival rules, modification rules, and consumer-protection rules vary by category, region, and individual operator.

A local service is often **not represented by a structured digital catalog**. Many merchants run on a mix of manual scheduling, third-party booking platforms, and informal channels. A single merchant's truth lives across multiple systems and human routines.

These properties do not make local services impossible for agents. They make local services unsafe to commit to without an upstream readiness layer. Treat a salon appointment, a yoga class, or a fixed-window pickup like a product on a feed and the agent will, on average, get the result that pattern earns: a payment that cleared against a commitment the merchant could not honor.

A bounded commitment object is a discrete, named, time-stamped, fulfillable promise.

A prepaid local service becomes agent-ready when the thing being purchased is bounded enough for an agent to reason about and the merchant operationally able to confirm. The shape of the object matters more than the vertical it sits in.

A bounded commitment object has a clear start and end, a clear capacity, a clear price, a clear set of conditions under which it can be honored or canceled, and a clear merchant confirmation step. It is, in effect, a small contract that an agent can verify before money moves.

The same merchant may produce some commitments that are agent-ready and others that are not. A salon may offer agent-payable single-service appointments while keeping multi-service bridal packages on a manual quote workflow. A restaurant may make table reservations agent-payable while keeping bottle service or private events outside the lane. A tour operator may handle public timed departures cleanly while leaving private charters to a sales team. The line is drawn around the object, not around the operator.

The clearest examples in current local commerce are familiar.

A **parking bay**: a specific space, in a specific garage, for a specific window of time, at a clearly published rate. Parking has been digitizing in narrow categories for years.

A **timed-entry ticket**: a defined entry window at a museum, attraction, or experience venue, with a fixed price and a fixed quantity available.

A **class seat**: one position in a yoga class, a cooking workshop, or a fitness session, with a known capacity and a known cancellation rule.

A **spa or wellness appointment**: a defined treatment, a defined duration, a defined practitioner, with deposit or prepayment behavior already common in the category.

A **restaurant deposit**: a reservation backed by a card hold or partial prepayment, common enough that the major restaurant platforms describe deposits as a no-show mitigation tool.

A **pickup order**: a short list of items, a confirmed total, a defined pickup window, against a merchant that already handles the same flow for human-driven mobile orders.

A **rental window**: a specific piece of equipment, a specific time block, a defined deposit, a defined damage policy.

Each object is simple enough for an agent to evaluate, complex enough to require merchant-side validation, and digital enough to anchor a payment. They are also the shapes most often cited in payment-side public examples of agentic commerce, which suggests both sides of the ecosystem are converging on the same primitive without explicit coordination.

04 · Vertical Taxonomy

Nine candidate verticals. Two of them prove the rails. Two of them prove why the readiness layer is needed.

Not every local service category is equally ready for an agent to commit on the user's behalf. The table below organizes nine candidates along the dimensions that matter for early pilots.

The categories with the cleanest payment story are not the same as the categories where merchant-side validation provides the most clearly differentiated value. That tension is the reason a single pilot vertical cannot answer both questions on its own.

Vertical	Fit	Payment	Truth complexity	Pilot feasibility	Why it matters	Key risks
Parking	High	High	Low	High	Cleanest bounded object: bay, time, rate. Clear payment-readiness story.	Low average ticket; large platforms already control inventory; less proof of merchant-validation value.
Attractions and ticketed experiences	High	High	Low to medium	High	Timed-entry tickets are mature digital commitments with clear quantity and price.	Reseller incumbents dominate inventory; thin merchant-validation differentiation in the simplest cases.
Classes and workshops	High	High	Medium	High	Capacity-limited, time-bound, schedule-driven. Strong fit for agent reasoning.	Schedule freshness, cancellation windows, refunds, varying instructor availability.
Local tourism experiences	High	High	Medium	High	Real-time availability and dynamic pricing already established in the category.	Weather, language, minimum group size, guide availability.
Spa, wellness, appointments	High	Medium to high	Medium to high	High	Deposit and prepayment already normalized; appointments are a clean unit.	Practitioner availability, regulatory caution near medical-adjacent services, cancellation terms.
Beauty and personal services	High	Medium	High	High for the validation thesis	Best category to demonstrate the merchant-truth layer because availability and policies are highly volatile.	Sole-trader privacy, response latency, fragmented booking systems, no-show variability.

Vertical	Fit	Payment	Truth complexity	Pilot feasibility	Why it matters	Key risks
Restaurant deposits and order-ahead	Medium to high	Medium to high	High	Medium	Commercially compelling and well documented as a deposit-friendly category.	Menus, modifiers, kitchen capacity, allergy and dietary handling, dispute risk.
Local retail with reservation or pickup	Medium	Medium	Medium to high	Medium	Bridges product commerce and local commitment when reservation or pickup is involved.	Inventory drift, substitution logic, SKU complexity.
Equipment rental	Medium to high	Medium	High	Medium	Clear bounded object when an existing platform manages inventory and deposits.	Inventory condition, damage liability, identity and KYC, handover logistics.

Read together, the table suggests a portfolio. Parking, ticketed experiences, and classes are good candidates for clean pilots that prove the rails work end to end. Beauty, personal services, and appointment-led wellness are good candidates for proving that an upstream merchant-truth layer earns its place in categories where it cannot be assumed. Restaurant and equipment categories are commercially compelling but operationally heavier; they are better placed once one of the simpler pilots has produced evidence the model holds.

No single vertical can answer both questions a pilot must answer. Use two lanes.

The temptation in early infrastructure is to pick one vertical and force the entire argument through it. That is a mistake here. The two questions a pilot must answer are different enough that no single category can answer both without one of the answers becoming theoretical.

Lane A. Payment readiness.

Lane A uses categories where the bounded commitment object is already digital, the inventory is already managed by mature systems, and the unit economics of a single transaction are already well understood.

The candidates are **ticketed experiences**, **classes**, and **parking**. Each produces a clean object an agent can purchase. Each interacts with infrastructure that already supports digital payment, deposit handling, and cancellation logic. Each minimizes operational complexity on the merchant side.

Lane A answers a focused question: when an agent reaches the point of commitment in a category that is already mostly digital, does the end-to-end flow work as designed. The result is a clean, defensible demonstration that agentic payment infrastructure can land in local services.

Lane B. Merchant truth.

Lane B uses categories where the merchant side is the hardest part. **Beauty and personal services** lead the list. **Spa**, **wellness**, and other appointment-led categories follow closely.

In these categories the merchant truth layer is doing real work. Availability is genuinely volatile. Practitioner skills are genuinely differentiated. Cancellation behavior is genuinely category-specific. The presence of an upstream readiness layer is the difference between an agent that can act safely and an agent that should not act at all.

Lane B answers a different question: in a category where the merchant side is fragmented, can a dedicated readiness layer produce evidence strong enough to make the commitment safe to proceed. The result is a defensible demonstration that the merchant participation layer earns its place in the stack.

The two lanes are complementary. Lane A shows the rails work. Lane B shows why upstream merchant infrastructure is needed before agents can commit on the user's behalf at scale.

A clean agentic ecosystem is likely to need both layers. Obenan is the upstream one.

The relevance of an upstream merchant-readiness layer to the payment side of agentic commerce is best stated narrowly.

Public agentic-payment programs across major networks and platforms address the payment side of trust. They create the structures within which agent-led payments can be tokenized, governed, and made visible to issuers and acquirers. Public draft specifications for verifiable intent and signed user mandates bind user authorization and agent action to cryptographic constraints. Live, controlled implementations have been demonstrated across multiple regions during the first half of 2026. The trajectory is toward a payment layer that knows who is allowed to pay, what they are allowed to pay for, and how much.

That work answers the credential question.

It does not, by itself, answer the merchant-readiness question. Public verifiable-intent drafts include thin merchant fields, useful for binding a checkout to a specific payee but not for verifying whether the payee can honor the commitment in the moment. Tokenization protects credentials and constrains agent payment execution. It does not establish that the slot, the ticket, the price, or the table was actually available at the time the commitment was made.

This is where the merchant truth layer sits. It is upstream of payment. It does not move money. It does not authorize payment. It does not see card data. It produces something the rest of the ecosystem currently has to assume: a structured, freshness-aware, validation-capable view of whether a specific local merchant is presently able to honor a specific commitment.

> Obenan prepares the merchant side of an agent transaction before payment execution begins. Payment networks govern payment execution and the credential chain that secures it.

A clean agentic commerce ecosystem is likely to need both layers. The merchant-truth side reduces ambiguity before payment is attempted. The payment side enforces authorization and trust at the moment of

execution. Disputes that arise in agentic flows have natural sources on each side: misalignment between user intent and agent action sits on the credential side; misalignment between merchant promise and merchant fulfillment sits on the readiness side.

Where the two sides connect is a question for each network and platform to decide on its own terms. Obenan does not claim integration, attachment-point compatibility, partnership status, pilot status, or program participation with any payment network, scheme, processor, or commerce protocol.

07 · What Obenan Does Not Do

The boundary is a deliberate architectural choice, not a defensive disclaimer.

The design of an upstream merchant-readiness layer is as much defined by what it refuses to do as by what it builds.

BOUNDARY

Obenan does not perform any of the following.

- Move money.
- Tokenize cards.
- Authorize payments.
- Act as a payment service provider.
- Act as an acquirer.
- Create checkout.
- Settle transactions.

Holding the line at upstream evidence is what allows the layer to compose cleanly with the parties on the other side of it. Networks set the rules of authorization. PSPs and acquirers run checkout, settle transactions, and own merchant-of-record relationships. Vertical SaaS owns booking and reservation logic in its categories. Obenan produces the merchant-side facts those parties currently have to assume in order to operate. A layer that stays neutral on payment scope is structurally simpler, regulatorily lighter, and easier to combine with whichever path each merchant is already on.

Three source tiers. Body claims rely on the first. Industry numbers anchor scale, not market size.

This briefing is a synthesis of public sources, with private partner conversations explicitly excluded. The discipline behind the evidence base is worth stating directly so that each claim can be audited.

Source hierarchy.

Three tiers were used. *Official / primary*: company product pages, public specifications, public press releases, and government statistical sources. *High-quality industry report*: commercial market reports and consultancy publications. *Vendor primary, self-interested*: company blog posts and product pages that describe their own outcomes. Body claims rely first on the official tier. Industry reports anchor category scale only. Vendor evidence is labeled and used to indicate direction, not to support comparative claims.

What is included.

Public agentic commerce announcements and specifications from payment networks and major platforms. EU statistical data on small-enterprise digitization. Commercial market sizing for local service categories that are widely cited and verifiable to a recent date. One operator and two restaurant-platform vendor disclosures on deposit and prepayment behavior, all clearly labeled.

What is excluded.

Private partner conversations of any kind. Internal meeting transcripts. Private partnership pipelines. Implementation-sensitive mechanics, including any internal architecture detail that would create patent or competitive exposure. Third-party press whose claims have not been confirmed by the named institutions. Internal estimates whose methodology has not been published.

Why category-scale numbers are directional.

Underlying market sizes for fitness, beauty, wellness, parking, and tours are real and large, but those are not agentic-commerce markets today. They are the categories from which an agentic-commerce wedge can be cut. Treating them as TAM for agentic services would overstate the case. Treating them as zero would understate the structural opportunity. The paper uses them as upper-bound anchors for category importance, with the explicit observation that agentic-services demand has not yet been independently sized.

Six claims, each anchored by a public source. Several plausible-sounding claims declined.

All quantitative claims about category scale are anchored in specific reports and clearly labeled. Live implementations are described as live implementations rather than as commercial rollouts. Draft specifications are described as drafts.

1. Payment networks and major platforms are publicly building agentic commerce infrastructure.

The relevant public materials include Mastercard Agent Pay, the public Verifiable Intent specification, the OpenAI Instant Checkout / Agentic Commerce Protocol announcement and product feed specification, the Stripe newsroom and documentation for the Agentic Commerce Protocol and Shared Payment Token, and Google's Universal Commerce Protocol and AP2 publications. [1, 2, 3, 4, 5]

2. Most of this public infrastructure is optimized around online product commerce.

The public protocol specifications and integration patterns center on structured product feeds, online checkout flows, and merchants whose catalogs already exist in a machine-readable form. [4, 5]

3. Local merchant digitization is uneven.

Eurostat reports that 21.38 percent of small EU enterprises engaged in e-sales, against substantially higher readiness among larger firms. [6]

4. Local service categories have meaningful underlying scale.

Fitness clubs were estimated at 121.19 billion dollars globally in 2024. Professional beauty services were estimated at 247.24 billion dollars globally in 2023. Medical spas were estimated at 21.21 billion dollars globally in 2024. Tours and activities reservations were estimated at 179 billion dollars globally in 2024, with online booking, real-time availability, and dynamic pricing identified as market drivers. [7, 8, 9, 10]

5. Parking and restaurant deposits provide narrower but cleaner examples of digitized commitments.

The parking reservation system market was estimated at 1.1 billion dollars in 2024. One major operator reported passing one billion dollars in cumulative parking reservations sold and described North American

parking as only two percent digitized. Restaurant platforms publicly describe deposits, prepayments, and card holds as effective tools for reducing no-shows; one platform reported a customer reducing no-shows from 15 percent to 1 percent after introducing deposits, while another reported a 0.9 percent no-show rate among customers using prepayment between December 2023 and March 2024. [11, 12, 13, 14]

6. Broader agentic commerce projections continue to focus on goods rather than services.

A widely cited consultancy projection suggests global agentic-commerce orchestrated revenue could reach three to five trillion dollars by 2030, but explicitly notes those figures reflect goods and do not yet include services. [15] The number is useful as context for strategic urgency. It does not size the prepaid local services category.

The evidence base does not support several adjacent claims that might appear plausible at a glance. It does not show that prepaid local services already function as a proven agentic commerce market. It does not show that most local merchants are ready for agentic commitment today. It does not show that deposits are universal outside selected categories. It does not show that public agentic payment specifications already validate merchant-side facts in the way described in this briefing.

Completion of one transaction is a demonstration. It is not a result.

A pilot in this category should produce credible answers to the following measurements.

Agent-query resolution into safe commitment	The share of agent queries that reach a state in which the agent can responsibly commit on the user's behalf, separated from queries the merchant truth layer correctly declines.
Merchant response latency	The distribution of times between an agent inquiry and a merchant-side confirmation, segmented by vertical and by hour of day.
Validation accuracy	The share of confirmed commitments the merchant actually honored, measured against the share of declined commitments that turned out to be honorable. Both directions of error matter.
Downstream completion	The share of validated commitments that proceeded through payment and fulfillment without dispute or refund.
Timeout and fallback rates	The share of inquiries that exceeded a defined latency threshold and the behavior the system applied when they did.
Merchant fatigue	Operational signals from the merchant side: opt-outs, response declines, complaints, and requests for caps on inquiry rate.
Dispute and ambiguity signals	Cases in which the merchant or the user disputed the commitment, with categorization by source: credential side, readiness side, or fulfillment side.
Freshness failure modes	Cases in which the merchant truth layer's data was stale at the moment of commitment, with severity classified by impact.

A pilot that produces credible answers to these eight measurements can support a substantive next conversation. A pilot that does not produce them, regardless of how many transactions complete, has only demonstrated that the rails worked for a particular happy path.

11 · Conclusion

Agentic commerce can fail before payment, when the merchant-side promise is unclear.

A payment-only architecture has no native answer to that failure. A merchant-readiness layer does.

A local service is a time-bound operational commitment. The payment layer can verify who is allowed to pay. The merchant layer must verify what can safely be committed. These are different layers, and the second is currently underbuilt.

Prepaid local services are one of the first practical bridges from AI discovery to agentic payment because they expose the merchant-side gap with unusual clarity. The wedge is bounded local commitments, not all local commerce. The pilot strategy is two lanes, not one. The right next step is narrow, scoped pilots with measurable outcomes and conservative claims.

Obenan sits upstream of payment execution. The work is not in competition with networks, PSPs, acquirers, or vertical SaaS. It is the layer those parties currently have to assume in order to operate cleanly. Building that layer carefully, in narrow categories, with a clear evidence base, is the most useful contribution Obenan can make to the agentic commerce ecosystem in this phase.

What this briefing does not assert.

This briefing is a public research statement. To keep it readable on its own terms, the following are stated explicitly so they are not inferred elsewhere.

01 Obenan does not move money, tokenize cards, authorize payments, act as a payment service provider, act as an acquirer, build checkout, or settle transactions.

02 Obenan has no integration with Mastercard Agent Pay, Mastercard Verifiable Intent, Visa Intelligent Commerce Connect, American Express ACE, the Universal Commerce Protocol, the Agentic Commerce Protocol, AP2, Experian Agent Trust, the Stripe Agentic Commerce Suite, or any third-party real-time decisioning service.

03 No payment network, scheme, processor, or commerce-protocol body has approved, endorsed, certified, or admitted Obenan into a partnership, pilot, accelerator, or developer program on the basis of this briefing.

04 Public agentic payment specifications referenced here do not, by themselves, validate live merchant-side facts in the way described in this briefing. The merchant-readiness layer is upstream of and distinct from those specifications.

05 Prepaid local services are not yet a proven agentic-commerce market at scale. The wedge claim is structural, not commercial proof.

06 Most local merchants are not ready for agentic commitment today. Deposit and prepayment behavior is normalized in selected categories; that is not the same as broad market readiness.

07 Merchant-side validation does not prevent disputes. It can produce evidence that may reduce ambiguity around availability, price, terms, fulfillment, and participation.

08 No current Obenan-onboarded merchant is *certified* or *committable* by any public network or open-protocol specification. References to bounded commitment objects describe a category shape, not a certification.

The references in the Endnotes are public sources used to anchor category claims. They are not endorsements of Obenan, and inclusion does not imply any commercial relationship between Obenan and the named institutions.

Endnotes

Sources

1 Mastercard. *Mastercard Agent Pay*. Public product page describing registered agents, network tokens, verified order intent, and consumer consent.
mastercard.com/us/en/business/artificial-intelligence/mastercard-agent-pay.html

2 Verifiable Intent. *Specification, draft v0.1*. Layered SD-JWT credential chain binding agents and users to specific commercial actions.
verifiableintent.dev/spec/

3 Mastercard and Santander. Press release on Europe's first live end-to-end payment executed by an AI agent in a controlled environment, March 2026.
mastercard.com/news/europe/en/newsroom/press-releases/en/2026/santander-and-mastercard-complete-europe-s-first-live-end-to-end-payment-executed-by-an-ai-agent

4 OpenAI. *Buy It in ChatGPT* (Instant Checkout / Agentic Commerce Protocol). Public announcement and product feed specification.
openai.com/index/buy-it-in-chatgpt · developers.openai.com/commerce/specs/file-upload/products

5 Stripe. Newsroom and documentation on Agentic Commerce Protocol, Shared Payment Token, and instant checkout.
stripe.com/newsroom/news/stripe-openai-instant-checkout · docs.stripe.com/agentic-commerce

6 Eurostat. *E-commerce statistics*. Share of small EU enterprises engaged in e-sales (21.38 percent), with comparison to larger firms.
ec.europa.eu/eurostat/statistics-explained/index.php?title=E-commerce_statistics

7 Fortune Business Insights. *Health and Fitness Club Market*. Global market size estimate for 2024 (121.19 billion dollars).
fortunebusinessinsights.com/health-and-fitness-club-market-108652

8 Grand View Research. *Professional Beauty Services Market Report*. Global market size estimate for 2023 (247.24 billion dollars).
grandviewresearch.com/industry-analysis/professional-beauty-services-market-report

9 Grand View Research. *Medical Spa Market*. Global market size estimate for 2024 (21.21 billion dollars).
grandviewresearch.com/industry-analysis/medical-spa-market

10 Research and Markets. *Tours and Activities Reservations: Global*. Global market size estimate for 2024 (179 billion dollars), with online booking and dynamic pricing as drivers.
researchandmarkets.com/reports/5310766/tours-and-activities-reservations-global

11 GM Insights. *Parking Reservation System Market*. Global market size estimate for 2024 (1.1 billion dollars).
gminsights.com/industry-analysis/parking-reservation-system-market

12 SpotHero. Press release on passing one billion dollars in cumulative parking reservations sold; characterization of North American parking as two percent digitized.
spothero.com/press/spothero-eclipses-1-billion-in-parking-reservations-sold

13 OpenTable. *Three Proven Payment Strategies to Reduce No-Shows*. Vendor guidance on deposits, prepayments, and card holds.
opentable.com/restaurant-solutions/resources/3-proven-payment-strategies-reduce-no-shows

14 Tock. *Eliminate No-Shows: Three Tock Tools*. Reported 0.9 percent no-show rate among customers using prepayments, December 2023 through March 2024.
exploretock.com/join/resources/eliminate-no-shows-3-tock-tools

15 McKinsey & Company. *The Agentic Commerce Opportunity: How AI Agents Are Ushering In a New Era for Consumers and Merchants*. Projects three to five trillion dollars in global agentic-commerce orchestrated revenue by 2030, with figures reflecting goods and not yet services.
mckinsey.com/quantumblack/our-insights/the-agentic-commerce-opportunity

CLOSING THOUGHT

The payment layer can verify who is allowed to pay. The merchant layer must verify what can safely be committed.