117-354-D Agentic Bots Challenge the e-Commerce Architecture

When the customer interface becomes an autonomous participant in commerce

November 2026

OPENING THESIS

For three decades, electronic commerce was designed around a human navigating a merchant-controlled interface. The agentic browser changes the unit of interaction: a governed cognitive mission may discover, decide, delegate, transact, and leave obligations behind after the browser session ends.

The resulting challenge is not simply bot detection. It is the governance of authority, commitment, speed, and consequence across three differently paced layers of the AI-Economy.

Prologue — The Airline That Drew a Boundary

On November 11, 2026, a Wall Street Journal headline reported that United Airlines was banning AI bots from booking flights. The accompanying excerpt attributed to the carrier concerns about data privacy and quoted the objection, "I don't want to create another middleman." The screenshot is a useful contemporary trigger for analysis, but it does not, by itself, establish the full scope of United's policies, enforcement mechanisms, or operational motives.

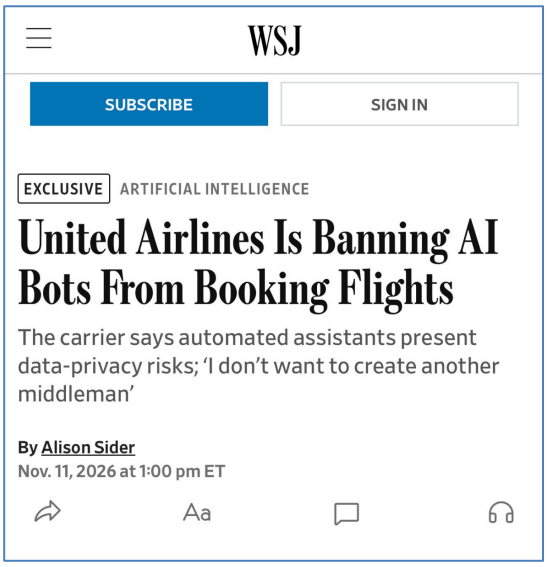

The deeper question is why a commercial enterprise might welcome AI-assisted discovery yet resist an autonomous assistant crossing the threshold into booking, payment, loyalty redemption, or modification of an existing reservation. The answer lies in the architecture of electronic commerce itself. An intelligent browser is no longer merely rendering pages; it may become a delegated economic actor.

I. The Original Contract of Electronic Commerce

Traditional e-commerce placed the merchant's website or mobile application at the center of the transaction. A human customer reviewed product descriptions, selected options, accepted terms, supplied payment credentials, and received a confirmation. Authentication, interface design, consent capture, inventory control, and dispute resolution were arranged around that sequence. The sequence was never perfect, but it established familiar places to ask who initiated the transaction, who agreed to it, and what the merchant promised.

Automation existed long before generative AI: price-comparison crawlers, travel distribution systems, scripts, bots, and APIs. What changes with agentic systems is the potential for a general-purpose cognitive actor to interpret natural-language intent, select tools, pursue alternatives, spawn delegated workers, and

continue a mission across services. The merchant may encounter a technically valid request without knowing the larger cognitive undertaking that produced it.

II. The Agentic Browser Is Not Merely a Faster Browser

A conventional browser is primarily a client environment for human-driven interaction. An agentic browser may combine perception, planning, navigation, tool invocation, and transaction execution. It can search a dozen airline websites, compare itineraries, assess trade-offs, and attempt a booking in the time a human might spend reviewing one page. The important discontinuity is not only velocity; it is the migration of operational discretion from the human interface into an autonomous actant.

The human's instruction "find a suitable flight" is not necessarily equivalent to "purchase any flight that appears suitable." Search, selection, reservation, payment, cancellation, and regulated declarations are distinct categories of authority. The e-commerce architecture must learn to express these distinctions in machine-verifiable form rather than assuming that a logged-in browser session represents sufficient permission for every action.

III. The Three-Layer AI-Economy Comes Into View

Layer	Primary objects	Characteristic constraint
L1 — Agentic / Cognitive	Intent, GCID, Actant-ID, delegation, decisions	Rapid decision and retry cycles
L2 — Digital Services	APIs, sessions, reservations, payments, transaction state	Consistency, authorization, rate and queue limits
L3 — Physical / Economic	Seats, crews, aircraft, airport slots, fulfillment	Finite capacity and slow replenishment



- Layer 1 can produce plans and requests at machine speed.
- Layer 2 translates requests into durable digital commitments, but responses can be delayed, duplicated, rejected, or ambiguous.
- Layer 3 must honor the underlying physical and economic constraints. A seat on a particular flight is not an infinitely reproducible digital token; it depends on aircraft configuration, operations, and a chain of real-world resources.

These layers are coupled but do not operate at a common cadence. Accelerating cognition does not automatically accelerate transaction settlement, customer-service resolution, airport capacity, or physical fulfillment. The mismatch can create feedback loops: an agent receives no acknowledgment, retries, creates duplicate commitments, and then reacts to apparent scarcity that its own actions helped produce.

IV. Differential Speeds: A Stability Problem, Not an Automatic Explanation

The differential-speed hypothesis is compelling as a general systems-engineering risk. Thousands of agents can generate searches, availability checks, or purchase attempts much faster than traditional human-facing workflows were designed to absorb. Delayed acknowledgments and retries may amplify load or distort inventory signals. Shared agents can couple multiple missions, turning a local overload into a cross-mission disturbance.

Yet speed mismatch should not be presented as the established reason for United's reported restriction. The headline and excerpt foreground privacy and intermediation. To establish a speed-driven operational cause would require evidence of actual service saturation, booking anomalies, or observed cascades. GUDIYA should preserve that distinction between a plausible instability mechanism and a documented incident.

The engineering response is likewise not simply to slow all agents. Rate limits, idempotency keys, capacity reservations, backpressure, reconciliation barriers, and bounded retries are existing controls. A GUDIYA research implementation would test whether cross-agent, cross-mission observation adds value beyond these established mechanisms.

V. Identity and Privacy: A Credential Is Not a Mission

An agent may hold a valid browser cookie, OAuth token, or service credential. That credential can establish access to a service under its own rules. It does not inherently establish who authorized the overarching cognitive mission, which actant is acting, what it may commit, or whether a delegated purchase remains within the human's instructions.

The GUDIYA identity model distinguishes Principal-ID, Actant-ID, GCID, and service credentials.

- Principal-ID represents the source of legitimate authority.
- Actant-ID identifies the participating execution entity.
- GCID identifies the governed cognitive undertaking, which may persist across sessions and participants.
- Service credentials provide destination-specific access.

These are distinct objects linked by verifiable, time-bounded relationships.

This separation must be privacy-preserving. A root GCID should not become a universal cross-site tracking cookie. The airline or merchant need only receive the claims necessary to decide whether the requested operation is permitted. Derived identifiers, selective disclosure, audience restrictions, and limited retention are essential design constraints, not optional embellishments.

THE COGNITIVE IDENTITY SEPARATION PRINCIPLE

- A valid service credential does not establish a valid cognitive mission.
- A valid mission does not authorize every actant.
- An authenticated actant does not establish collective stability.

The service provider retains independent authority to accept or reject a transaction, even when an external Grid can attest to the mission and delegation context.

VI. Differential Commitments: The More Subtle Failure Mode

Consider a travel agent instructed to arrange a trip.

- At Layer 1, it may decide that a fare is acceptable.
- At Layer 2, a booking engine may create a temporary reservation, capture payment, issue a ticket, or return an ambiguous timeout.
- At Layer 3, a real seat must remain available and the traveler must ultimately be transported. These are not interchangeable states.

1. A completed browser workflow does not prove that a ticket was issued.
2. A ticket issued does not prove that a flight will operate.
3. An apparent payment failure does not prove that no charge occurred.
4. A canceled reservation may leave a refund obligation outstanding.

In the agentic economy, an assistant can stop operating while its economic consequences remain active.

This is why GCID-linked obligation lineage is central. The mission should be able to correlate its actions with reservations, payment records, ticketing outcomes, refunds, changes, and unresolved duties. The Grid must distinguish claimed completion from verified settlement, and it must represent uncertainty honestly when a counterparty does not expose sufficient evidence.

VII. The Missing Fourth Dimension: Commercial Control

The quoted concern about “another middleman” points to a force that cannot be reduced to privacy or stability.

- An airline has economic interests in its customer relationship, distribution costs, loyalty program, merchandising, and control of the booking experience.
- An agentic intermediary may become the customer’s primary interface and shift bargaining power away from the merchant’s website.

That is a legitimate commercial question, distinct from the technical question of whether an agent can be authenticated safely. A technically secure agentic protocol does not oblige a merchant to adopt it. Conversely, a restriction framed in security language may coexist with commercial incentives to preserve direct distribution. Without further evidence, one should not assert that either motive fully explains a particular company’s decision.

VIII. The Other Forces: Liability, Consent, Abuse, and Operational Semantics

Agentic commerce introduces additional boundary conditions. A traveler may need to personally attest to certain declarations. An agent might select a technically available but unsuitable itinerary. Fraud controls may interpret rapid automated activity as abuse. Loyalty accounts can contain economically valuable balances. A merchant may need to decide whether an AI-mediated instruction constitutes valid consent for a change or purchase.

There is also a semantic problem. Human-facing pages often rely on visual cues, exceptions, and informal assumptions. Agents can misread dynamic fares, bundled services, refundability, baggage allowances, or timing restrictions. Formal machine-readable interfaces reduce ambiguity but do not eliminate disagreements over intent, authority, or responsibility.

Architectural force	What the merchant must establish	What the Grid may contribute
Identity and privacy	Who is acting and which data may be disclosed	Bounded mission/actant attestations
Delegated authority	Which exact action the principal approved	Scope, expiry, revocation, approval evidence
Digital commitment	Whether reservation/payment/ticketing succeeded	Cross-service obligation lineage
Physical capacity	Whether commitments remain feasible	Capacity signals and aggregate risk estimation
Commercial control	Which intermediaries are accepted	Interoperable governance, not forced admission
Regulation and liability	Which acts require human attestation	Escalation and evidence of human participation

IX. The Agentic Commerce Admission Boundary

A useful GUDIYA design response is an Agentic Commerce Admission Gateway(ACAG), understood as a proposed interoperability pattern rather than an existing standard. Before a consequential action enters a merchant's system, the gateway would evaluate the destination's own rules together with verifiable claims about the principal, actant, GCID, delegation, requested operation, and current risk context.

The gateway should not be an omniscient central controller. It must respect local merchant policy and legal requirements.

- It may approve discovery while rejecting payment;
- permit a hold while requiring human approval for ticket issuance;
- or allow cancellation while preserving a reconciliation obligation.

The architecture should support denial, limited permission, challenge, and safe fallback—not merely an allow/deny bit.

Such an interface could complement OAuth/OIDC, credential vaults, API gateways, transaction idempotency, and established travel distribution systems. Its distinctive research claim would be the additional governance of persistent mission identity, cross-agent participation, collective stability, and consequence closure. That claim needs controlled comparative testing.

X. The Three Graphs Behind One Transaction

1. The first graph is the Cognitive Identity Graph: principal, mission, actants, delegated authority, and execution context.
2. The second is the Digital Transaction Graph: search, reservation, payment, ticketing, modification, and refund.

3. The third is the Physical Consequence Graph: seats, operational capacity, travel fulfillment, and downstream obligations. They should be connected through verifiable relationships rather than collapsed into one giant centralized identity store.

- A mission may use many actants, and an actant may serve many missions.
- A transaction may contribute to multiple obligations, and an obligation may depend on several resources.
- Some links are many-to-many and time-bounded. This makes the architecture a temporal graph, not a straight hierarchy of browser cookies beneath a user account.

The Grid's role is to reason about the coupling among these graphs: which missions share resources, where commitments accumulate, which actions remain authorized, and whether an intervention has left unresolved obligations. No individual website or agent can be presumed to see the complete field.

XI. A Worked Flight-Booking Scenario

A traveler authorizes a GCID to identify a suitable flight and purchase it only if the total price remains below a stated ceiling and the itinerary meets specific constraints. Two actants compare carriers; one uses a browser, the other an authorized API. Both discover the same flight. A delegated booking actant attempts a reservation, but the response times out after the airline has accepted the request.

A naive agent retries immediately. A second reservation may be created, or the booking system may reject it. The correct outcome depends on transaction semantics, idempotency, and authoritative reservation state—not on the agent's interpretation of a missing response. Meanwhile, another actant may have already initiated a payment. The mission now contains uncertainty across identity, digital commitment, and financial consequence.

A GUDIYA-enabled demonstrator would first suspend further consequential retries, preserve GCID/Actant-ID lineage, query authorized systems of record, reconcile reservations and charges, and establish whether any outstanding obligations require cancellation or refund. If the Grid had intervened, the mission would undergo a proportionate fitness-for-service check before unrestricted resumption.

The research question is whether the GCID-linked cross-agent view reduces duplicate commitments, unresolved obligations, and recovery time relative to a well-engineered conventional workflow. If it does not, the proposed architecture must be simplified or revised.

XII. From Bot Blocking to Governed Interoperability

The commercial landscape may support several coexisting modes.

- Some enterprises will restrict third-party autonomous transactions to their own apps.
- Others will expose narrow APIs or partner channels.
- Some will permit AI discovery while reserving payment, regulated declarations, and exception handling for explicit human confirmation.

There is no reason to assume that all merchants will adopt one universal agentic interface.

For GUDIYA, the strategic opportunity is not to compel openness. It is to define a verifiable grammar of participation:

- who authorized the mission,
- which actant acted,
- what authority was granted,
- which service accepted the action, and
- what consequences remain.

Such a grammar could make interoperability safer where parties voluntarily adopt it.

XIII. Engineering Tests for the Stability Grid

A credible experiment should compare ordinary browser automation, a conventional authorized API workflow, and a GCID/Actant-ID-aware multi-agent workflow. Inject lost acknowledgments, duplicate submissions, stale fare data, concurrent agent purchases, credential revocation, session termination, and delayed refunds. Include scenarios in which a merchant declines admission or returns only partial evidence.

Measure unauthorized commitments blocked, duplicate reservations, reconciliation latency, false interventions, trace completeness, customer-data exposure, and operational overhead. Vary the relative cadences of agent decision-making, digital transaction processing, and physical-capacity updates. GUDIYA's value must be demonstrated against strong existing controls, not against a deliberately weak baseline.

XIV. The Institutional Question: Who Governs the Gateway?

A cross-enterprise admission gateway cannot simply declare itself authoritative. The principal authorizes the undertaking; the actant executes under bounded delegation; the merchant controls its service; regulated activities may require additional human or institutional approval. A federated Grid should provide verifiable evidence and stability support without erasing those distinct sources of authority.

This implies a governance requirement beyond the technical graph: auditability of policy decisions, independent review of interventions, protection against cross-domain tracking, and continuity when a vendor or gateway becomes unavailable. The stabilizer itself must remain subject to governance and recoverability controls.

THE AGENTIC COMMERCE TRUST BOUNDARY

An agent may be authorized to discover information without being authorized to reserve, purchase, modify, or discharge an obligation. Service credentials shall not be treated as proof of mission legitimacy or unlimited delegation.

Consequential actions crossing enterprise boundaries shall be associated, where participating systems support it, with a verifiable principal, executing actant, governed mission, bounded authority, and destination-specific acceptance decision.

The cognitive mission shall retain accountable lineage across digital commitments and physical consequences until outstanding obligations have a verified or explicitly qualified disposition.

Epilogue — Commerce After the Browser

The e-commerce website was built for a world in which the customer's intention, the visible interaction, and the transaction often occupied one continuous human session. Agentic systems break that assumption. Intent can be delegated. Execution can be distributed. Credentials can change. A booking can be accepted without an acknowledgment. An obligation can survive the disappearance of every browser window involved in creating it.

United Airlines' reported restriction is therefore a revealing signal, though not proof of any single underlying failure mechanism. It places identity, privacy, commercial control, and transactional authority at the boundary of an emerging agentic economy. The systems-engineering response is to treat those concerns together while preserving the distinctions among them.

The browser session may end. The actant may change. The merchant may reject admission. But when consequential work has occurred, the identity of the undertaking and the disposition of its obligations must remain reconstructable. The future of agentic commerce will depend not only on how intelligently agents act, but on how reliably independent institutions can govern what those actions commit the world to do.

CLOSING STATEMENT

The agentic browser transforms e-commerce from a merchant-controlled human-interface sequence into a federated problem of identity, authority, commitment, capacity, and consequence.

GUDIYA proposes that cognitive identity follow the mission; authorization remain bounded by the destination; and stability governance extend across all three layers of the AI-Economy.

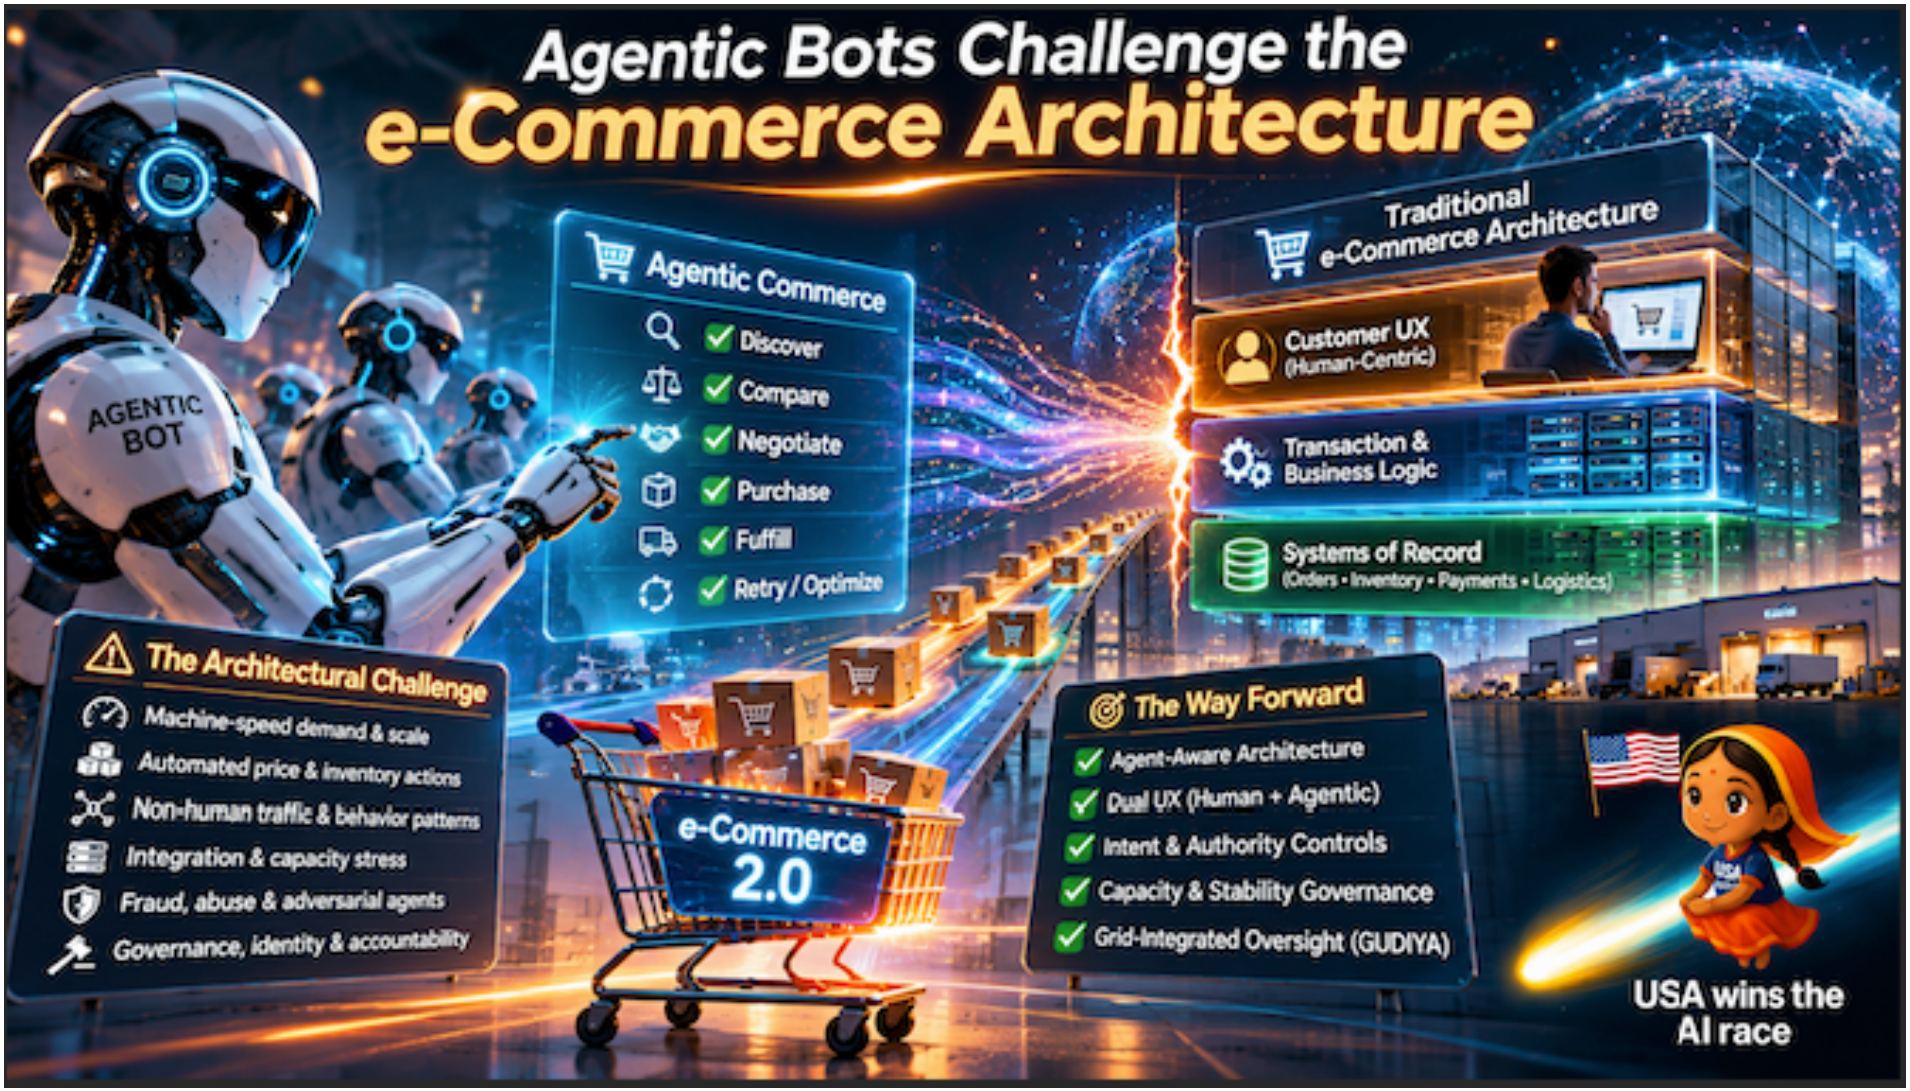
Agentic Bots Challenge the
e-Commerce Architecture
AGENTIC BOT
Agentic Commerce
Discover
Compare
Negotiate
Purchase
Fulfill
Retry / Optimize
Traditional e-Commerce Architecture
Customer UX (Human-Centric)
Transaction & Business Logic
Systems of Record
(Orders • Inventory • Payments • Logistics)
The Architectural Challenge
Machine-speed demand & scale
Automated price & inventory actions
Non-human traffic & behavior patterns
Integration & capacity stress
Fraud, abuse & adversarial agents
Governance, identity & accountability
e-Commerce 2.0
The Way Forward
Agent-Aware Architecture
Dual UX (Human + Agentic)
Intent & Authority Controls
Capacity & Stability Governance
Grid-Integrated Oversight (GUDIYA)
USA wins the AI race