RESEARCH & INSIGHTS

THE HIDDEN LAYER SERIES | ARTICLE 8

By Joe Batista | November 2026 | 18-minute read

To subscribe to the series, complete the form

Governance Is the New Supply Chain

Why proof of authorization is beginning to determine where agents can act—and who will stand behind them

The Hidden Layer series maps where value is created and captured across the AI stack, one chokepoint at a time. Articles 1 and 2 drew the map and showed how bottleneck power migrates through it. Articles 3 to 7 followed constraints on building AI: packaging capacity, the energized megawatt, the rules that ration it, the cost of money and code already written. Article 8 is the first to move into the top half of the map, where the input in short supply is not physical at all. It examines the constraints on putting AI to work.

CORE INSIGHTS

1. **Capability is not the limit. Accountability is.** American Express has said it will protect Card Members from errors made by registered AI agents that send Amex the customer's authenticated purchase intent. The offer is not yet live, and it depends on proof. An agent goes only as far as someone will stand behind it.
2. **Identity is not authorization.** A log shows which agent acted, not who allowed it. Every consequential agent action that crosses an organizational or financial boundary must clear four gates: access, authority, evidence and accountability. The slowest one sets how far the agent can go.
3. **Creating proof is not the same as having it accepted.** The parties that rely on proof records—and the parties that absorb losses when an agent is wrong—help decide which are good enough. That is where leverage could form. It has not yet been shown to.
4. **The buying question is changing.** From "How do we govern AI?" to "What must we prove before we allow AI to act on our behalf?" The enterprise's own risk, legal and finance teams are gatekeepers too.
5. **Two precedents point the same way.** Web certificates became free while browsers kept the power to decide which issuers to trust. The chip-card liability shift showed that changing who pays changes what gets adopted. Whether agent proof follows is our hypothesis, tested separately.

Proof Before Purchase

On October 13, 2026, American Express published what it calls its Intent Intelligence framework: a way to capture and verify what a customer wants before an AI agent buys anything on the customer's behalf. The customer's instructions and limits are recorded and approved, and the system issues an Intent ID and a Proof of Intent Token, which can be used to obtain a payment credential scoped to what was approved. Later, the record is used to check whether a purchase matches and to support dispute resolution, as The Paypers described it. TechCrunch gave the practical version: if a customer asks for white Nike sneakers in size 9 and the agent tries to buy Adidas, Amex could block the transaction.

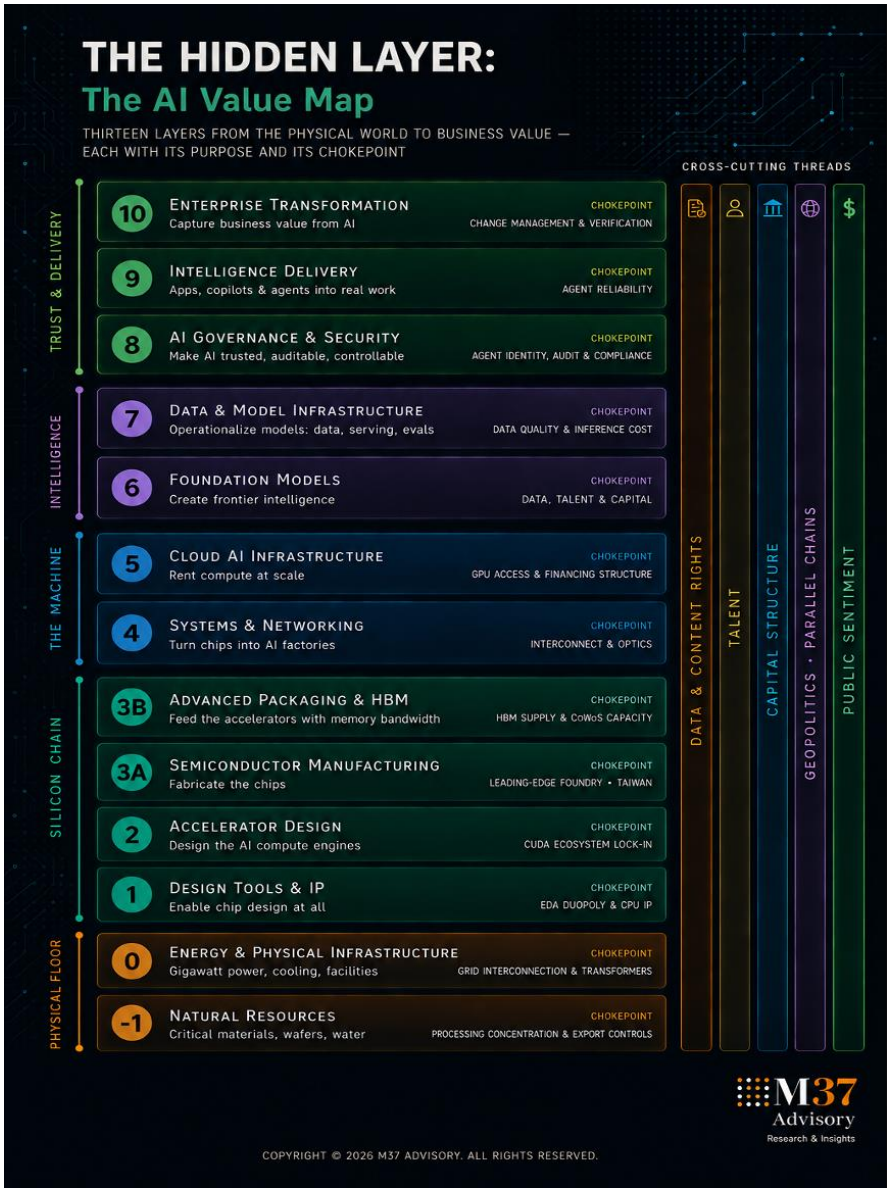
The framework is the second half of a promise. In April, Amex announced what it called “an industry-first commitment to protect Card Members from registered agent error.” It is conditional, and it is not yet live: Amex said it would offer it “in the future,” to Card Members whose registered agent sends Amex the customer’s authenticated purchase intent, and that “terms and conditions will apply.” The October framework does not change that. Luke Gebb, Amex’s executive vice president and global head of innovation, described the condition to Digital Commerce 360: “if the merchant has made no error and the cardmember has made no error, but the agent has made the error,” then “we will stand behind that transaction and credit the consumer so they’re not out of pocket.”

Most commentary on AI agents asks what they can do. This answers a different question: how far will anyone let them go? A card network that settles disputes has said it intends to absorb the cost of an agent’s mistakes, and the offer attaches to proof. TechCrunch also reported that Amex has not yet seen an increase in disputes tied to agentic commerce. The terms are being written before the losses arrive.

Amex was not alone. In the five weeks to October 14, Google gave its new coworker agents their own accounts, NIST summarized responses on agent identity, members of Congress proposed agent security standards, the UK’s data regulator opened a call for evidence on agentic AI, and Velotix described AI that carries security policies into data structures it designs. Each was working on part of the same question: when an agent acts, who let it, and who answers for it?

> *“Your AI agent can only go as far as someone is willing to stand behind it.”*

Where the Constraint Sits: Four Gates



The Hidden Layer: The AI Value Map. 13 layers, each with its purpose and its chokepoint.

ON THE MAP: Layer 8 — AI Governance & Security | Chokepoint: Agent Identity, Audit & Compliance | This article: who creates proof of authorization, and who decides whether it counts

Layer 8 is where AI is made trusted, auditable and controllable. It sits above the models and their infrastructure (Layers 6 and 7) and below the applications and agents that put AI into real work (Layer 9). Everything beneath it on the map makes AI possible. Layer 8 decides where AI is allowed.

Article 1 put the idea in one line: “Governance is the new supply chain problem. You can deploy AI anywhere. Regulators decide where you can defend deploying it.” This article does not argue that all of governance becomes a supply chain. It argues something narrower:

> *Proof of authorization is becoming a necessary input for autonomous agents. It becomes a supply-chain constraint when an agent acts on assets it does not own, and the owner of those assets needs evidence and accountability before it will let the agent act.*
>
> — The thesis of this article

An agent drafting a memo for its own user rarely meets that condition. An agent moving money, changing a customer's account or committing a company to a contract does. It must clear four gates. The first three establish what an agent can access, who authorized its actions, and what evidence remains afterward. The fourth is where the financial consequences land.



The four gates. "What exists" reflects the evidence in this article as of October 2026. Source: M37 Advisory.

A supply chain moves at the speed of its slowest input, and so does an agent: excellent data access does not help if nobody can show who approved an action, or accept the loss if it goes wrong.

Gate 3 has two parts, and the difference between them is the economic center of this article. **Evidence creation:** can you produce an independently verifiable record of who authorized the action, and within what limits? **Evidence acceptance:** will the party relying on that record recognize it? Consider an agent with its own identity and access to a customer's account. The customer authorizes a $500 purchase, and the system records the instruction and every step the agent takes. Everything looks correct, and the payment network or the merchant still disputes whether the record establishes authorization. The problem is not the absence of evidence. It is its acceptance. Creation can become cheap and open. Acceptance sits with whoever relies on the record.

Proof makes the decision to stand behind an agent possible. It does not make anyone willing to.

Two of the map's cross-cutting threads run through this layer. **Public Sentiment**, because agents will be allowed to do only what people and their regulators trust them to do. **Geopolitics · Parallel Chains**, because the same agent may have to satisfy a proposed US standard, binding EU logging duties and UK data protection law at once.

> *"The company that creates the proof may not be the one that decides whether it counts."*

Three Examples

① Who Is Responsible When an Agent Acts? The Payment Networks

Gates 3 and 4 are furthest along in payments, which already has a system for deciding who pays when something goes wrong. In just under thirteen months, four major players in digital payments published different pieces of the trust infrastructure for agentic commerce. Google and Mastercard focused on verifiable records of user intent. Visa focused on helping merchants recognize trusted agents. American Express combined agent registration and authenticated intent with a conditional promise to stand behind certain agent errors. These initiatives are related, but they are not the same: identifying an agent is not the same as proving what a customer authorized.

- **Google, September 2025.** The Agent Payments Protocol (AP2), developed with "more than 60 organizations," is built on Mandates: "tamper-proof, cryptographically-signed digital contracts that serve as verifiable proof of a user's instructions."
- **Visa, October 2025.** The Trusted Agent Protocol lets merchants recognize trusted agents and distinguish them "from malicious automation and rogue bots."
- **Mastercard, March 2026.** Verifiable Intent, co-developed with Google, is "a tamper-resistant record of what a user authorized when an AI agent acts on their behalf." Mastercard said it is open-sourcing the specification.
- **American Express, April and October 2026.** Agent registration, the intent token, and a conditional promise, not yet live, to stand behind registered agents.

Three questions sit underneath these designs, and they can have different answers. Was the agent technically permitted to act? Did the customer authorize this transaction? Who bears the loss if it is disputed? Technical permission, customer authorization and financial liability are related, but they are not interchangeable. An agent can have permission to act without establishing that a particular transaction reflected the customer's intent. And proving intent does not, by itself, determine who bears the loss. A cryptographically signed record does not automatically settle a legal or financial dispute.

That is why gate 4 remains open. Reshmi Suresh, head of agentic commerce at the payments processor Worldpay, wrote in July: "No liability shift exists yet." Without a signed mandate, "a disagreement between agent and shopper becomes a coin flip, and issuers tend to side with the cardholder."

Observed: of the four, only Amex has announced it will absorb losses from agent error, and only where the proof exists; that protection is not yet live. **Inferred:** a specification can be open to all, but the parties that shape dispute outcomes help decide whether a given record is enough.

IN BRIEF: WHAT TO WATCH IN PAYMENTS

- **The signal that matters:** not another protocol, but the first live network rule that makes protection or dispute rights depend on a proof record.
- **The open question:** whether Visa and Mastercard follow Amex in absorbing agent error, or leave it to issuers, merchants and agent platforms to settle by contract.
- **Ask your payments partner:** which agent proof formats will you accept, and what happens to a disputed agent purchase that arrives without one?

"In payments, the parties that shape dispute outcomes also help determine what counts as acceptable proof."

② Who Authorized the Agent to Act? The Identity Platforms

An agent's identity is not the same as its authority to execute a specific transaction. Inside companies, that gap arrives through the identity system.

On October 8, Google introduced coworker agents that, as The Next Web reported, each get their own Workspace account, email address and directory entry and act under their own identity. According to the report, the announcement describes a "cryptographically attested identity" stamped into each agent's logs, and administrators approve role-based permissions. The Next Web summed up the audit design: "Every action is logged to the agent rather than to a person."

Microsoft took a different route. Its documentation for Entra Agent ID generally requires at least one sponsor when an agent identity or agent identity blueprint is created, with documented exceptions, such as blueprint principals at creation. Sponsors are "business representatives accountable for the agent's purpose and lifecycle decisions"; they can disable an agent. Microsoft also sells the controls: a June 17, 2026 notice (MC1395007, as archived) said Conditional Access and ID Protection for agents are included in Microsoft Agent 365 and Microsoft 365 E7.

Google gives the agent a standing identity. Microsoft attaches a standing, accountable person to it. Both address who stands behind the agent in general. Neither, in the material we reviewed, records task-level authority: who approved a specific action, and within what limits. NIST's commenters saw the same gap. Its summary reports that they wanted delegation chains traceable to the responsible human or organization, and task-level authorization over standing permissions. They also named the cost: "Consent Fatigue: At machine speed, asking for human approval for every action is impossible." Someone still has to decide, in advance and in writing, what an agent may do without asking.

IN BRIEF: WHAT TO ASK YOUR IDENTITY VENDOR

- **Not** "does the agent have an identity?" **but** "can you show, for one specific action, who authorized it and within what limits?"
- **Who is the sponsor of record** for each agent today, and what happens to its permissions when that person changes roles or leaves?
- **Will the record leave your platform** in a form a customer, auditor or card network would accept, or does it only make sense inside your console?

"A log shows which agent acted, not who allowed it."

③ What Can It See, and What Must Be Kept? Governed Data and the Rules

Gate 1, data access, is where the most products are arriving. Immuta's Agentic Data Access (March 2026) gives an agent "its own identity, attributes, intent, and audit trail," grants access temporarily for each task, and, Immuta says, makes "every action … traceable across the human, the agent, the purpose, and the data accessed": a gate 1 product reaching toward gate 3.

AI is also beginning to surface relationships nobody labeled. On October 14, Velotix described its Atlas research. It describes "multiple specialized AI agents" that aim to "discover, structure, and optimize knowledge without relying on human-defined schemas or labels," and says "existing security policies, sensitivity classifications, and masking rules are carried into" the resulting structures. Atlas is, in the company's words, "currently in rapid research and development." These are the company's descriptions of its own research, not claims we could verify. (Disclosure: Velotix is a former M37 Advisory client.)

Three things need to be kept apart here. **Discovery** is finding that two pieces of data are related. **Inference** is what an agent can conclude from that relationship. **Access** is whether it may see the underlying data at all. Discovering a relationship does not give an agent access to the data behind it. The governance challenge is narrower and harder: access policies were written for structures people designed, so they may not anticipate what an agent can infer from relationships nobody modeled. And even a fully governed gate 1 settles only what an agent can see. Permission to read a customer's record is not permission to change the account, approve a payment or make a commitment.

WHAT EVIDENCE MUST BE KEPT? THE RULES

Washington. The Stop Rogue AI Act, announced by Representatives Josh Gottheimer and Mike Lawler and introduced on September 14, 2026 as H.R. 10362, would direct NIST to develop standards for discovering, verifying and controlling AI agents. Axios reported the standards would be voluntary for most organizations. The bill is a proposal, not law.

NIST. Its concept paper on agent identity (February 5, 2026) drew more than 600 responses. "The overwhelming majority of commenters favored an evolutionary approach," building on existing standards.

Brussels. The EU AI Act (Regulation (EU) 2024/1689) divides the logging duty for high-risk AI systems three ways. Article 12: the system must "technically allow for the automatic recording of events (logs)." Article 19: providers keep the logs under their control for at least six months. Article 26(6): deployers keep the logs under their control for "at least six months, unless provided otherwise in applicable Union or national law." Separately, Article 26(2) requires deployers to assign human oversight to "natural persons who have the necessary competence, training and authority." These duties apply where the AI system's intended use is classified as high-risk. For systems classified as high-risk under Annex III, the requirements apply from December 2, 2027, White & Case reported.

London. The ICO's call for evidence, with accountability among its listed areas, closes November 20.

The pattern: the rules ask organizations to keep a record and name a person who answers for it. They say evidence must exist; they do not yet say whose evidence counts.

What This Changes for the Business

The company that makes an agent capable of acting may not be the company that has the authority to let it act, the evidence to prove it was authorized, or the balance sheet to absorb the consequences. That is the business problem. If the thesis holds, the buying question shifts from "How do we govern AI?" to **"What must we prove before we allow AI to act on our behalf?"** Answering it means solving seven problems that today sit in different budgets:

- Discovering and inventorying agents.
- Establishing agent identities.
- Managing delegation and authorization.
- Recording and verifying agent actions.
- Enforcing permissions and limits.
- Establishing accountability and liability.
- Making evidence interoperable across platforms.

Those problems divide into two economic roles:

Role	What it does	Illustrative providers and initiatives	How its economics may work (our inference)

Role	What it does	Illustrative providers and initiatives	How its economics may work (our inference)
Proof producers	Supply identity, authorization records, logs and the tools that create them.	Google, Microsoft, Immuta, Velotix; AP2, Verifiable Intent	Many suppliers and open specifications. Likely to face price pressure over time, as certificates did.
Proof gatekeepers	Decide which records are accepted, under what conditions, and how responsibility is allocated when they fall short.	Card networks and large counterparties (external acceptance and allocation); the enterprise's own risk, legal and finance teams (internal authorization); regulators (baseline requirements)	Few parties, close to where losses are settled. Positioned to set terms, if urgency arrives.

The roles overlap: an identity platform can create proof, enforce policy, or both. And the existence of a gatekeeper does not automatically create a profitable bottleneck. That requires evidence of limited alternatives, meaningful switching costs, and the ability to impose terms. We have identified the potential position, not yet demonstrated the economic power.

The enterprise is a gatekeeper too. Card networks are the clearest external example, but most agents will be stopped, or allowed, inside the company. An agent that approves supplier payments, changes customer pricing, issues purchase orders or commits the company to a contract needs no payment network to block it. The company's own risk, legal, finance and security teams can. External acceptance asks whether another party will recognize the evidence. Internal authorization asks whether the enterprise will allow the agent to act in the first place. The former is primarily a gate 3 question; the latter is gate 2. Both ultimately connect to gate 4: who stands behind the consequences.

WHERE THE MONEY MOVES: ONE PROCUREMENT AGENT

An enterprise deploys an agent to approve supplier purchases. It has access to procurement data and a spending limit. It places an order outside the intended scope. To resolve it, the company must establish whether the agent had permission; whether the person who authorized it had the authority to; whether the recorded evidence is reliable; whether the supplier or payment provider accepts that evidence; and who absorbs the loss.

The commercial consequence (our illustration): before deploying, the enterprise may pay for identity infrastructure, authorization controls, evidence retention, independent verification, insurance and integration services. The party that can establish acceptable proof may influence the deployment decision. The party that absorbs the risk may influence the contract terms. That is where value and liability move.

The following business implications are M37 inferences, not reported findings.

- **For a technology CEO:** customers may start asking you to stand behind what your agent does, not just license it. Are you creating proof, or proof that someone who pays will accept? Only the second lets you price the risk.
- **For a chief revenue officer:** the size of an agent deal may be set less by what the agent can do than by what the customer's risk, legal and finance teams will let it do. If gates 3 and 4 are shut, the deal stays a pilot.
- **For a channel founder:** installing agents and configuring access is gates 1 and 2. The opening is in gates 3 and 4: designing what gets recorded, making it acceptable to the customer's counterparties, and helping decide who carries the risk.
- **For a board member:** who in your value chain is absorbing agent risk today without knowing it? Your company, your vendor, your customer, or nobody, until the first dispute?

"Identity and access products are arriving. Acceptance and accountability will determine how far the deals can go."

The Pattern: Two Precedents, Two Lessons

Two earlier markets faced versions of this question. They teach different things, and the argument needs both.

Lesson 1, the padlock: creating proof versus accepting it

How does a browser know a site is who it says it is? The answer was the certificate, issued by a certificate authority the browser chose to trust. **Proof became a business.** On May 19, 2010, VeriSign agreed to sell its Authentication Services business, which included its SSL certificate services, to Symantec for about $1.28 billion in cash.

The gatekeeper, not the issuer, decided which proof counted. In January 2017, Symantec-issued certificates that broke industry rules were disclosed. Google's Chrome team wrote in September 2017 that Symantec had allowed several organizations to issue certificates "without the appropriate or necessary oversight," and that a longer pattern had led Chrome to "lose confidence" in its infrastructure. Chrome removed trust in stages, ending with Chrome 70 in 2018. Symantec sold the business to DigiCert, announced August 2, 2017, for about $950 million in cash and about a 30% stake in DigiCert.

Basic proof got cheap; the decision about which proof counts did not. On December 3, 2015, Let's Encrypt, "a free, automated, and open Certificate Authority," entered public beta. In February 2018 Google announced that, beginning in July with Chrome 68, "Chrome will mark all HTTP sites as 'not secure'." Basic certificates became close to free. The power to decide which issuers were trusted, and to label the absence of proof in front of every user, stayed with the browsers.

Lesson 2, the chip card: changing who pays changes what gets adopted

On August 9, 2011, Visa announced that, effective October 1, 2015, liability for "domestic and cross-border counterfeit card-present point-of-sale (POS) transactions" could move to the merchant's acquirer when a chip card was presented to a merchant without chip terminals. The rule did not require anyone to adopt chip technology. It changed who paid without it. Amex's agent offer has the same shape, turned around: protection for those with proof.

Where both precedents fall short. A certificate vouches for a website, a bounded identity relationship. An agent's authorization record may have to establish who initiated the request, which agent acted, what it was authorized to do and within what limits, whether that authorization was still valid, whether the agent exceeded it, and whether the result created a financial obligation. NIST's commenters said the intent behind it is not yet well defined. And the chip-card shift had a date; as of October 2026, agent liability has none. If the pattern repeats, value could migrate from creating proof to controlling its acceptance. That is our hypothesis, not a market fact.

"Basic proof got cheap; the decision about which proof counts did not."

The M37 Lens: Article 2's Four Questions, Applied to Proof

Article 2 asks four questions of any bottleneck: Is there a shortage? Can suppliers exploit it? Can buyers wait? Can buyers pay? Agent proof does not fit its measurable 0–10 bands, so we describe each pillar in words. The absence of measurable scarcity does not mean the absence of a constraint. Here the constraint is institutional acceptance, contractual risk and the inability, so far, to allocate liability, which is different from a capacity shortage. The Market Structure pillar splits in two: for producers we keep Article 2's question; for gatekeepers, who are not suppliers, we ask whether they can capture leverage.

Pillar	Article 2's question	What the evidence shows (October 2026)	Scored?
Constraint	Is there a shortage?	No measurable capacity shortage. Acceptance remains unsettled: competing proof designs, one conditional offer to absorb losses, and no settled liability framework (Worldpay).	No. There is no supply headroom or spot price to measure.
Market Structure: proof producers	Can suppliers exploit it?	Unlikely to have much room. Many suppliers; AP2 involves more than 60 organizations, and Mastercard is open-sourcing its specification.	No. No source places producers in the rubric's bands.
Market Structure: proof gatekeepers	Can gatekeepers capture leverage?	Possibly. Few parties decide which proof is accepted, close to where losses are settled. No evidence yet of pricing power or of buyers unable to switch proof systems.	No. Described qualitatively.
Customer Urgency	Can buyers wait?	In payments, for now: Amex reports no rise in agent-related disputes. Inside companies, less so: agents are being issued accounts and sponsors now.	No.
Capital Capacity (enabler)	Can buyers pay?	For proof, the cost is integration and process. For accountability, it is balance sheet: whoever stands behind agents must be able to absorb losses.	No.

M37 VIEW: the infrastructure is being established before the economic pain becomes widespread. On Article 2's matrix, that is a latent bottleneck, not a demonstrated one. A latent bottleneck is a position. A demonstrated one shows limited alternatives, meaningful switching costs and the ability to impose terms. We have evidence of the first, not the second. That placement is our judgment, not a score.

WHAT WE MEAN BY "GOVERNANCE IS THE NEW SUPPLY CHAIN"

Observed: the authorization problem is real, and the infrastructure addressing it is emerging. No liability shift exists, by Worldpay's account, and Amex's proof-linked protection is not yet live.

Inferred, the proof-of-authorization thesis: where agents act on assets they do not own, proof of authorization becomes a necessary input, and agents go only as far as all four gates allow. The economic importance of acceptance is plausible.

Hypothesized, the gatekeeper hypothesis (unproven): creating proof becomes cheap; the power to decide which proof counts concentrates with those who control access or bear the consequences. The concentration of economic power remains unproven.

What would weaken both: protection for agent actions offered without proof requirements; agents moving into high-stakes work while gates 3 and 4 stay unresolved; or acceptance decided by a neutral body rather than by those who pay.

The Four-Gate Test: Four Questions You Should Be Asking

Each question is also a buying requirement, and each has an answer you can measure.

1. **Access: how many agents are running in your systems, and for what share can you list what each can see, including relationships AI has surfaced that nobody labeled?**
2. **Authority: what share of your agents have a named sponsor and written limits, and is each sponsor still in role today?**
3. **Evidence: for one agent action chosen at random, how long would it take to show who authorized it, and would your customer, card network or regulator accept that record?**
4. **Accountability: for your highest-value agent action, do your contracts, your vendor's terms, the network rules and your insurance name the same party as bearing the loss?** In payments, according to Worldpay, they do not yet.

Then ask the question on the graphic: which gate is your slowest?

The Predictions

Same rules as Articles 2 to 7: checkable by anyone, including where we were wrong. Every counting window runs from the date of publication, November 17, 2026, to June 30, 2028, when measurement stops. October 2028 is the reporting date, when we mark the results in public. The tests measure observable signals; on their own they cannot show that gatekeepers capture economic value. In plain terms, here is what each test asks:

Prediction (reported Oct 2028)	What we are really testing	Tests
Networks tie protection to proof	Does proof become commercially necessary?	Thesis
Others start standing behind agents	Does anyone beyond the networks accept agent risk?	Thesis
Open proof, gated acceptance	**Does economic power concentrate with those who decide which proof counts?**	**Hypothesis**
Standards move to guidance	Does institutional standardization develop?	Supporting
Federal rules arrive	Does regulation accelerate it?	Supporting

THE GATEKEEPER HYPOTHESIS: THE LEAD TEST

Prediction (window ends Jun 30, 2028)	Supports the hypothesis	Weakens this pattern	Status
Open proof, gated acceptance	**Supports:** both of the following. (a) At least one agent-authorization specification published under an open-source or royalty-free license is accepted in live programs by at least two of Visa, Mastercard and American Express. (b) Each network that accepts it still requires agents to be registered, approved or listed by the network, or by a party it designates, before their transactions qualify	Either no open specification is accepted by two networks (the open-proof, gated-acceptance pattern has not emerged; control may instead remain with proprietary proof providers), or a network accepts open-specification proof with no network-run registration (no gatekeeper). Untestable if fewer than two networks have live agent programs	Open

This test measures one specific form of gatekeeping: network-controlled acceptance. Others are possible, such as an enterprise consortium, an insurer, a regulator or a standards body deciding which proof counts. We will report any we observe, but they do not count toward this test.

THE PROOF-OF-AUTHORIZATION THESIS

Prediction (window ends Jun 30, 2028)	Supports the thesis	Weakens it	Status
Networks tie protection to proof	**Supports:** at least two of Visa, Mastercard and American Express have live terms under which cardholder protection, dispute rights or liability for agent-initiated purchases depend on a registered or verified agent and an authenticated record of what the cardholder authorized. Amex's April 2026 commitment counts only once it is live	None or one of the three has such terms live	Open
Others start standing behind agents	**Supports:** at least three companies outside the card networks, such as software vendors, cloud providers, IT services or channel firms, or insurers, publicly offer to compensate customers for losses caused by AI agents' actions, through a warranty, indemnity, guarantee or insurance product	One or none	Open

SUPPORTING INDICATORS (THEY INFORM THE READING BUT DON'T DECIDE IT)

Prediction (window ends Jun 30, 2028)	What it would show	What would weaken it	Status
Standards move to guidance	NIST publishes at least a draft practice guide or other draft guidance on agent identity and authorization, beyond a project description	Only a project description, or nothing	Open
Federal rules arrive	The Stop Rogue AI Act, or another federal statute directing agent identity, inventory or logging standards, is enacted	No such statute is enacted. We report this either way	Open

MEASUREMENT RULES, FIXED NOW

- **Sources:** official documents of the company, network or agency (rules, program terms, product terms, insurance policy wordings, filings, congress.gov, NIST publications). Where no official text exists, two independent major news reports.
- **Live:** generally available to US customers. Pilots, betas, waitlists and commitments described as "in the future" do not count.
- **Open-source or royalty-free:** a specification released under an OSI-approved license, or published by a standards body under its royalty-free patent policy.
- **Offer to stand behind agents:** a published commitment to compensate a customer for a loss caused by an AI agent's action, available to customers generally, not negotiated one-off terms. Insurance counts if offered by a licensed insurer or under a named insurance program.
- **Deduplication:** each company counts once per test, however many products, regions or announcements it has. Subsidiaries and acquired companies count as their parent. A renamed product counts once. A specification and its later versions count as one. A joint offer counts once, credited to the party that pays the claim, and an offer counted under the networks test is not counted again under the others test.
- **Expansions:** an existing program or offer extended to new regions, products or customers counts as the same offer, not a new one.
- **Insufficient evidence:** if no official text or two independent reports can settle whether an item qualifies, it does not count, and we say so when we report. A test with too little qualifying activity to decide is reported as untestable, not as supported or weakened.

- **Judgments:** the thresholds and named lists are M37 judgments, fixed on publication and not revised afterwards.

The Bottom Line

Articles 2 through 7 examined constraints on building AI. This article examines the constraints on putting AI to work: who has the authority to let an agent act, who can prove that authority existed, and who will absorb the consequences. In October 2026 the first two gates are filling with products. Proof formats compete, liability is unsettled, and by Amex's account the losses have not arrived. We have identified a potential chokepoint, not yet a demonstrated one.

The next competition in AI may not be about who builds the most capable agent. It may be about who has the authority to let it act, the evidence to prove it was allowed, and the ability to stand behind what it does. Capability gets an agent to the door. Authorization lets it through. Accountability determines who is willing to open it.

> *"You can only defend AI where you can prove what it did."*

How This Article Was Tested

Sources. Facts come from primary documents wherever they exist: company releases and documentation from American Express, Google, Visa, Mastercard, Microsoft, Immuta and Velotix; NIST publications; the sponsors' release and congressional record for the Stop Rogue AI Act; the text of the EU AI Act; and, for the historical pattern, releases and filings from VeriSign, Symantec and Visa and posts from Google and Let's Encrypt. Where a fact rests on reporting, the outlet is named in the text. Company descriptions of their own products are reported as such, not as evidence that the products work. Every statistic and quotation is listed in a separate fact-check record.

Three levels of claim. Observations, inferences and the hypothesis are labeled separately, and the hypothesis has its own test. The business implications and the producer and gatekeeper economics are labeled as inferences. Article 2's scoring rubric was not applied, because no evidence places agent proof in its bands.

What we left out. Reported incidents in mid-2026 in which AI agents allegedly took unauthorized actions, one of which the bill's sponsors cite, because we could not verify the details at their original source.

Disclosure. Velotix is a former M37 Advisory client. Velotix did not review this article, which relies only on the company's published description of Atlas.

Predictions. Thresholds, named lists, counting windows and the deduplication rule were fixed before publication and will not be changed.

About the Author

Joe Batista advises boards, C-suite executives, and go-to-market teams on technology value creation, growth strategies, and the opportunities hidden beneath major technology shifts. He has spent more than 40 years in the technology industry, guiding companies on growth and technology decisions. His career spans the minicomputer era at Digital Equipment Corporation (DEC) through client-server, cloud, virtualization and now AI. He has held executive roles at Dell Technologies and Hewlett Packard Enterprise.

Through M37 Advisory, Joe helps leadership teams recognize recurring patterns, uncover second-order effects, find the opportunities others overlook, and identify the strategy worth betting on. His work is grounded in a simple belief: What if the biggest obstacle to your growth isn't the competition, but the way you think?

"Most executives have been bombarded with 'theories' about what might work. Based on our research, Joe's extensive playbook is rooted in creating alignment of interest and business outcomes that leaders should take notice of."

— Jim Dickie, Research Fellow & Sales Mastery Founder, CSO Insights

Email: Chief.Creatologist@M37-Advisory.com | Website: www.m37-advisory.com | Idea Vault: www.m37-advisory.com/idea-vault

NEXT IN THE SERIES

Article 9: "One Island." We move back down the map to Layer 3A, Semiconductor Manufacturing, where the chokepoint is leading-edge foundry capacity in Taiwan. Every agent in this article, and every record of what it was allowed to do, runs on chips, many of them made at the leading edge. Some risks can be diversified away. Others have to be planned for.

To subscribe to the series, complete the form.

Sources

PRIMARY DOCUMENTS

[1] American Express. "American Express Debuts Agentic Commerce Experiences (ACE)™ Developer Kit and Announces Industry-First Protection for Registered Agent Purchases." Business Wire, April 14, 2026. Link
[2] American Express. Intent Intelligence framework for agentic commerce, Amex newsroom. October 13, 2026. Link
[3] Google Cloud (Stavan Parikh and Rao Surapaneni). "Announcing Agent Payments Protocol (AP2)." September 16, 2025. Link
[4] Visa. "Visa Introduces Trusted Agent Protocol: An Ecosystem-Led Framework for AI Commerce." October 14, 2025. Link
[5] Mastercard (Pablo Fourez). "Verifiable Intent." March 5, 2026. Link. Open specification (draft v0.1) and reference implementation: verifiableintent.dev, which links to the GitHub repository. Link
[6] Microsoft. "Agent owners, sponsors, and managers," Microsoft Entra Agent ID documentation. Accessed October 2026. Link
[7] Microsoft. Microsoft 365 message center notice MC1395007, "Microsoft Entra: New service plans for Conditional Access and ID Protection for agents." June 17, 2026, as archived by merill.net. Link
[8] Immuta. "Introducing Agentic Data Access: The Next Era of Data Provisioning." March 26, 2026. Link
[9] Velotix. Description of Velotix Atlas on velotix.ai. October 14, 2026. Link
[10] NIST National Cybersecurity Center of Excellence (Booth, Fisher, Galluzzo, Roberts). Concept paper, "Accelerating the Adoption of Software and Artificial Intelligence Agent Identity and Authorization." February 5, 2026. Link
[11] NIST NCCoE. "Summary of Comments on the Concept Paper," Agentic AI Identity and Authorization Project resource hub. September 29, 2026. Link
[12] Office of Rep. Josh Gottheimer. "Gottheimer Introduces Bipartisan Bill to Stop Rogue AI Agents and Keep People in Control." September 9, 2026. Link
[13] European Union. Regulation (EU) 2024/1689 (Artificial Intelligence Act), as amended by Regulation (EU) 2026/1744; consolidated text as of July 27, 2026, Articles 12, 19, 26 and 113. Link; Article 19 via the European Commission AI Act Service Desk. Link
[14] U.S. Congress. H.R. 10362, Stop Rogue AI Act, 119th Congress. Introduced September 14, 2026. Link
[15] VeriSign. "VeriSign to Sell Authentication Services Business to Symantec for $1.28 Billion." May 19, 2010. Link
[16] Symantec Corporation. Form 8-K, Exhibit 99.01, first quarter fiscal 2018 results (DigiCert transaction). August 2, 2017. Link
[17] Google Security Blog (O'Brien, Sleevi, Whalley). "Chrome's Plan to Distrust Symantec Certificates." September 11, 2017. Link
[18] Google Security Blog (Emily Schechter). "A secure web is here to stay." February 8, 2018. Link
[19] Let's Encrypt. "Entering Public Beta." December 3, 2015. Link
[20] Visa. "Visa Announces Plans to Accelerate Chip Migration and Adoption of Mobile Payments." August 9, 2011. Link
[21] M37 Advisory. Article 1, "The Map Nobody Has Drawn," and Article 2, "The Accidental Kingmaker." 2026.

REPORTING AND COMMENTARY

[22] TechCrunch (Sarah Perez). "Amex introduces a framework to prevent AI agents from making unauthorized purchases." October 13, 2026. Link
[23] The Paypers. "American Express publishes framework for agentic commerce intent." October 14, 2026. Link
[24] Digital Commerce 360. American Express agentic commerce developer kit and purchase protection. April 14, 2026. Link
[25] The Next Web. "Google launches workplace AI agent that gets its own email address." October 8, 2026. Link
[26] Axios. "Exclusive: New bill cracks down on AI agents." September 3, 2026. Link
[27] Worldpay (Reshmi Suresh). "Agentic commerce liability is still being written." July 30, 2026. Link
[28] White & Case. "EU AI Omnibus enters into force, amending the AI Act." August 4, 2026. Link
[29] LexisNexis. "ICO seeks views on application of data protection law to agentic AI." October 8, 2026. Link