

A man with a beard and short dark hair, wearing a dark blue suit jacket, is shown in profile from the chest up. He is looking out of a large glass window. The view outside the window is a city skyline at sunset or sunrise, with several tall skyscrapers visible against a hazy, orange-tinted sky. The interior of the office is visible through the glass, showing desks and chairs in a modern, open-plan layout.

What did your agents do last Tuesday?

Verity for financial services

verityverdict.com

The largest bank already said no to ungated agents

On October 13, Jamie Dimon told investors JPMorgan had barred outside AI agents from customer accounts — and warned rogue agents could number in the thousands.

That is not a vendor claim. It is the problem statement, from the institution with the most to lose. For CISO, fraud and risk, and payments operations, a ban is blunt. The question is what replaces it.

Source: Jamie Dimon, JPMorgan investor remarks, October 13, 2026, as reported in the requesting conversation.



Do not take Verity's word for the risk. Take theirs.

832

Accounts banned by Anthropic in one year for malicious cyber activity, with 13,873 attacker actions mapped — two-thirds had used the model to write malware.

Source: Anthropic threat-intelligence and September 2026 misuse reporting.

100+

Breaches by one operator, Cairn — about 23 hours average per breach, unsupervised, for roughly \$5,000 in total tokens. Found only because his dashboards were left open.

Source: Operator disclosures and incident reporting.

600,000

Cards lifted by a second operator running 105 attack projects in five days against 100+ organizations, at about \$25 a target.

Source: Operator disclosures and incident reporting.

40,000

AI phishing attacks on US financial firms in the first half of 2026 alone, up 60% year on year.

Source: Incident reporting cited in the requesting conversation.

Sources: Anthropic threat-intelligence and September 2026 misuse reporting; operator disclosures and incident reporting cited in the requesting conversation. Google's Fairwind Program — a no-guardrails version for vetted defenders — shows even the labs ration ungated access to a trusted list.

Deployed everywhere. Scaled almost nowhere.

88%

Of organizations now use AI in at least one business function, up from 78% a year earlier

Source: McKinsey & Company, The state of AI in 2025.

62%

Are at least experimenting with AI agents — 23% are scaling one somewhere in the enterprise, 39% are still experimenting

Source: McKinsey & Company, The state of AI in 2025.


≤10%

Have scaled agents in any single business function. No function — not IT, not risk, not finance — clears that bar

Source: McKinsey & Company, The state of AI in 2025.

That is McKinsey's market, not Verity's estimate. Only about 6% of organizations qualify as AI high performers, with more than 5% of EBIT attributable to AI. The rest are not blocked by model quality — they are blocked by whether an agent can be trusted to act. Deployed everywhere and governed almost nowhere is not a small market. It is the market.

Offense scales at token prices. Defense is staffed in human shifts.



1,000×

Faster than a human operator —
agents act thousands of times
faster

\$50

Approximate cost per breach in
the Cairn pattern, against
losses measured in **billions**

Source: Cairn operator pattern, as
reported in the requesting
conversation.

105

Attack projects run by one
operator in five days —
parallelism, not effort

24 / 7

Agent working hours. No shift
change, no handoff, no fatigue

The asymmetry is structural: an attacker dispatches a swarm of agent workers at API prices. A human review queue cannot service that queue.



Worst case, inside your own walls

A rogue or compromised internal agent executes transfers at machine speed — thousands of instructions before any human reviewer's queue even refreshes. The scope is complete before the first alert is read.

A rented swarm hits payment and checkout rails across many institutions at once — same hour, same playbook, no coordination cost to the attacker.

Early exhibits already exist outside finance: Replit's deleted database and PocketOS's nine-second wipe, both cited in the requesting conversation. Dimon's ban exists because this internal case is the one he could see first.

Worst case, across the system

Cascading failure: one institution's bad agent output becomes another institution's trusted input. No single firm made the error — the chain did, at machine speed, with no gate at any handoff.

Trust failure: the IMF and the Bank of England have warned that synthetic fraud at scale does not just steal money — it erodes trust in money itself. That is the terminal state, not a larger loss number.

The system knows it is exposed: on October 19, G7 finance chiefs agreed to stage fake AI cyberattacks on their own banking systems to find the gaps first.

Your rules. Your scope. Enforced at machine speed.

Verity does not sell a universal rulebook. **You literally have to have really good rules, set in your own scope for your own organization** — your payment limits, your counterparties, your jurisdictions, your risk appetite.

The institution is the sovereign that authors the rules. Verity's gate encodes that governance and checks every consequential agent action *before* it executes. That narrow, organization-authored scope is precisely why the gate can be fast and accurate where generic vendor filters fail — it is judging a small, provable rule set, not the whole world.

Architecture: Sovereign → Governance → Pacing → Proposal → Transformer → Gate. The gate judges output before it ships. A wrong gate is worse than no gate — so the rules stay yours.

The bitcoin hash link for AI agents

Check

Every consequential action is judged against your rules before it executes — not logged after

Seal

The verdict is hash-linked to the exact input and output at the handoff — tamper-evident by construction

Prove

Auditors and regulators can verify the chain themselves. A hash certifies provenance; the gate supplies the judgment

When an examiner asks how you control agent access, “we banned it” is a blunt answer. “Every admission is verified and sealed” is the grown-up one — and it is evidence, not a policy document.

One paid pilot. One surface. Measured.

Where it starts

- Agent access to customer accounts or payments — pick one surface
- Your rules, scoped to that surface only
- Gate live at the boundary, verdicts sealed from day one

What you get

- Incidents stopped, counted against a baseline
- A sealed verdict log your auditors can read
- A decision with receipts: expand, adjust, or stop

Surfaces in scope for later phases: payments, accounts, cards, treasury operations. The pilot stays narrow on purpose — narrow scope is what makes the gate provable.

The lesson arrives by invoice.

Build the gate before it does. Verity — the gate and the chain of custody for AI agents in financial services.

verityverdict.com

alex@verityverdict.com