

What machine commerce claims — and what actually happens

MCRI measures what publicly listed machine-payment infrastructure claims, what it actually does, and where the two diverge.

Of **955** endpoints **402index** labelled degraded or down, **754** — **79%** — returned a valid payment challenge on every independent probe we made in late August. **52** labelled healthy never returned one.

FINDING 01

The observability layer is stale and wrong

Machine commerce already has an observability layer: health labels published by observability providers (e.g., **402index**, **PulseFeed**, **x402-trust**, **vet402**), which agents use to route purchases. Two distinct failure modes show up in the observability layer: stale measurements and incorrect classifications.

STALE

The observation may be old, or the provider’s stated cadence may not be supported by its visible freshness evidence. The strongest cadence claims in the field sit beside counters that do not support them — or beside no counter at all (each provider’s own published words beside its own published counter, all read 2026-09-01):

PROVIDER	THEIR CADENCE LANGUAGE	THEIR OWN FRESHNESS COUNTER, SAME SURFACE
vet402	every catalog-listed x402 endpoint, probed daily	2,750 of 14,662 active endpoints measured in the last 7 days (18.8%)
PulseFeed	probe rotates through the registry over weeks (PulseFeed’s own note)	median age of a currently-served live verdict: 214 hours; oldest 17.9 days
x402-trust	probes every listed endpoint continuously ... many times a day	no freshness or coverage counter published anywhere on the surface
probe402	per-route cadence stated on each grade page (daily; a small weekly class)	8,528 of 9,440 covered routes observed in the last 7 days (90.3%); 907 never yet observed

An old observation is not necessarily a wrong one — a label can be weeks old and still true.

WRONG

A fresh label can still contradict independent observation.


Across 21–31 August 2026 we recorded 29,022 observations and compared 7,225 independently graded endpoints with **402index**’s health labels. On 1 September the `last_checked` stamps on the disagreeing labels were hours old, so staleness does not explain what follows.

The disagreement ran in both directions — endpoints labelled bad that presented a valid payment challenge containing a payable quote (the machine-readable response that shows a price was offered, not that anything was paid), and endpoints labelled healthy that did not:

degraded / down: 754 of 955 endpoints (79.0%) returned a valid payment challenge on every probe; 489 of them did so on three or more probes across at least two days.

healthy: 52 of 5,459 endpoints (1.0%) never returned a valid payment challenge; 15 more were unreachable on every probe.

FIG. 1 • BOTH DIRECTIONS OF THE DISAGREEMENT



Each bar is its own population of 402index-labelled endpoints, compared with independent observation, 21–31 August 2026; the disagreeing labels' last_checked stamps were hours old on 1 September.

402index's **healthy** label means the endpoint answered HTTP 402 within its timeout; MCRI tests whether the response contained a payable quote. That definitional difference explains part of the gap, but it does not eliminate the underlying observation: the labels and the independently observed payment-wall behaviour can diverge. (Coverage widened during the window, so most endpoints carry days of record, not weeks; one of the 754 was graded on all eleven days.)

WHY IT MATTERS

An agent, or a buyer of agent infrastructure, routing on a health label may be relying on a measurement that is days or weeks old — or simply wrong when it was checked.

WHAT IT CHANGES

Health labels need to expose when they were last tested and what the test actually established. Without that context, buyers cannot judge whether a label is safe to route on.

FINDING 02

Listed does not mean machine-executable

A service (e.g., an x402 API endpoint) can be publicly listed, carry a price, and advertise machine payment — and still not publish enough of a request contract for an agent to construct a real call.

Across the x402 and MPP route catalogues (Bazaar and mpp.dev), 16,197 routes were listed on 1 September, combined and de-duplicated. **1,521 of the 8,056 POST routes (18.9%) publish no usable request example.** Of those, 975 publish no example at all, 481 publish an empty {}, and 65 substitute a schema, an unsupported encoding or an unusable type. On the MPP rail the gap is structural: **794 of 794 catalogue-discovered POST routes publish nothing sendable**, because that catalogue's format has no request-example field at all — a statement about what the catalogue makes available, not about what each operator may publish elsewhere.

FIG. 2 • POST ROUTES WITHOUT A USABLE REQUEST EXAMPLE



Of those: 975 publish no example at all · 481 publish an empty {} · 65 substitute a schema, an unsupported encoding or an unusable type.

These are not dead listings. Knocked at the payment wall — a minimal call sending nothing, claiming liveness and the quote and nothing more — 598 of the 831 contract-less routes graded (72.0%, 31 August–1 September) returned a valid payment challenge containing a payable quote: price, asset, pay-to address. The endpoint exists. The payment wall exists. What is missing is enough public request information to safely construct the call. (The share describes the routes selected for missing contracts, never the sector; “not machine-executable from the published contract” is the claim — *not buyable* is never the claim.)

794 / 794

catalogue-discovered POST routes on the MPP rail publish nothing sendable — the catalogue’s format has no request-example field at all.

WHY IT MATTERS

The catalogues list these routes as available, and this failure is invisible at the catalogue layer unless someone attempts to construct the actual request.

598 / 831

contract-less routes graded returned a valid payment challenge containing a payable quote (72.0%) — the payment wall exists; the published contract is not enough to construct the call.

WHAT IT CHANGES

Catalogue standards should treat executable request information as part of discoverability. A listing is not sufficient for machine integration unless an agent can construct the request from the published information.

FINDING 03


Infrastructure is more concentrated than catalogue diversity suggests

The x402 and MPP route catalogues (Bazaar and mpp.dev) list 16,197 routes across 1,768 hosts as of 1 September. That apparent breadth of listings compresses into a much smaller underlying infrastructure stack.

SETTLEMENT

Most settlement routing cannot be identified from published payment challenges: 1,229 of 1,677 catalogue hosts (73.3%) publish no information identifying the party that would settle the payment — the facilitator. Among the hosts where routing is visible, Coinbase dominates: **405 of the 438 attributable hosts (92.5%) route to Coinbase** (next: PayAI, 38). Two things this figure is not: hosts bound to a facilitator are not settlement volume, and the invisible 73.3% are unknown, not Coinbase.

FIG. 3 · SETTLEMENT ATTRIBUTION AMONG CATALOGUE HOSTS



1 September. Two things this figure is not: hosts bound to a facilitator are not settlement volume, and the invisible 73.3% are unknown, not Coinbase. Next: PayAI, 38.

GATEWAYS

On the MPP rail, apparent brand diversity is mostly catalogue integration structure. 109 of 142 listed services (76.8%) are third-party gateway mounts by the catalogue’s own integration field, and **978 of 1,456 endpoint rows (67.2%) are mounted through the gateway domains of Orthogonal, Locus and Tempo** — a statement about how the catalogue is integrated, not about economic control. The marquee names — Anthropic, OpenAI, Google Gemini, fal.ai, Firecrawl — appear as mounts on Tempo’s gateway, not as first-party listings; first-party services number 33, unchanged across our reads since 20 August.

OPERATORS Three operators — Orthogonal, delx, m2mcent — hold **2,979 of 16,197 listed routes (18.4%)** on four hostnames between them.

WHY IT MATTERS

Headline provider counts suggest a broad market; the settlement, gateway and operator layers under it are each narrow.

WHAT IT CHANGES

Catalogue diversity can therefore overstate infrastructure diversity: a failure, policy change, or access decision at a small number of underlying parties could affect a much larger apparent market.

OTHER SIGNALS

Payment terms move under static listings. Of 6,835 endpoints quoted 21–31 August (22,318 observations), 33 changed their payment address between observations (281 changes; one endpoint changed on nearly every observation) and 73 changed price. A buyer relying on a cached quote — price or address — could act on obsolete terms.

The denominator will not hold still. The listed universe moved 16,170 → 16,197 routes between two complete reads 93 minutes apart on 1 September. Because the catalogue denominator moves, market-size claims need a timestamp.

Paid delivery is the open question. Current ecosystem observability measures payment acceptance more readily than successful paid delivery; probe402's paid-delivery audit is not yet complete.

WHAT THE SECTOR SHOULD FIX

- 01 Health labels should expose freshness** — when the underlying observation was made, and what was actually tested.
- 02 Catalogue standards should require an executable request example or schema** — discovery should give an agent enough to construct the request.
- 03 Payment protocols should expose facilitator identity, or an equivalent settlement-routing identifier, where appropriate** — today's visible settlement attribution is too opaque.

METHOD & LIMITATIONS

MCRI rests on probe402's own observations of publicly listed machine-payment endpoints, made from a single vantage on a stated per-route cadence. Routes with a published request contract are executed as documented; contract-less routes are tested only at the payment wall. Observations are dated and retained; third-party figures are the providers' own published numbers, quoted with their denominators and our read date. Full methodology, corrections and dispute route: probe402.com/method.

MCRI #002 (September) adds provider movers.

Free sample: 100 observations at probe402.com/publications/pub_mcric_001/r1/sample.csv.

We welcome corrections, questions, and suggestions for MCRI #002.

GET THE UNDERLYING MCRI DATASET**Research dataset**

endpoint-level observations and historical data for independent analysis.

API access

programmatic access to current and historical observations.

Intelligence subscription

MCRI monthly, historical trends, provider benchmarking and alerts.

Custom research

a specific question answered using the probe402 dataset.

Contact hello@probe402.com for dataset access, pricing, feedback, or suggestions for the next issue.