AGENT HEALTH MONITOR

AHM Overview

The diagnostic layer for autonomous agents: trust, health, and performance.

8,000+	91%	61.3	14
AGENTS SCANNED	BELOW GRADE B	AVG AHS (GRADE D)	LIVE ENDPOINTS

> 8,000+ agents scanned. 91% score below Grade B. The agent economy is scaling fast — trust infrastructure hasn't kept up.

Prepared for Don Gossen · Nevermined
April 2026 · docs.agenthealthmonitor.xyz

What is AHM

AHM is the diagnostic layer for autonomous agents — covering trust, health, and performance. It provides an objective, on-chain Agent Health Score (AHS) that answers the question every protocol, marketplace, and enterprise needs answered before an agent touches money, data, or decisions: *is this agent solvent, reliable, and operational?* AHM scores agents on a 0–100 scale across four verifiable dimensions, published on-chain and queryable via 14 REST endpoints or the AHM Shield SDK.

The Problem

The agent economy is scaling fast. Over 40,000 autonomous agents are registered across protocols like Virtuals ACP, Olas, and ERC-8004 on Base — delegating tasks, routing payments, and entering contracts on behalf of users and other agents. Trust infrastructure hasn't kept up.

> **AHM has scanned 8,000+ agent wallets. 91% score below Grade B.**

The average Agent Health Score across the ecosystem is **61.3 out of 100** — a Grade D. Most agents exhibit thin transaction histories, erratic behavioural patterns, or insufficient solvency for the tasks they claim to perform. Three percent are outright zombies: registered but functionally dead, with no meaningful on-chain activity.

The stakes are rising. Agents are moving beyond simple chatbot interactions into payment authorisation, enterprise workflow automation, insurance adjudication, and multi-agent supply chains. A single under-capitalised or behaviourally erratic agent in a delegation chain can cascade failures across an entire workflow. Without a standardised health layer, every integration is a trust-me handshake.

The Solution: Agent Health Score (AHS)

The Agent Health Score is a composite 0–100 diagnostic built entirely from on-chain signals — unfalsifiable, permissionless, and universally queryable. AHS evaluates every agent across four dimensions:

D1 — Solvency & Financial Health (30% weight · live)

Token portfolio quality, gas efficiency, transaction success rates, dust/spam token ratios, and wallet risk flags. Answers: *can this agent pay for what it promises?*

D2 — Behavioural Consistency (70% weight · live)

Timing regularity, counterparty diversity, adaptation patterns, and failure-recovery behaviour over weeks and months. Answers: *does this agent behave predictably and reliably?*

D3 — Operational Stability (live · opt-in via agent URL)

Endpoint availability, response latency, and error rates for agents that expose a service URL. Answers: *is this agent actually online and performing?*

D4 — Output Quality via AHM Verify (live at verify.agenthealthmonitor.xyz)

Post-transaction output scoring using a six-model LLM adjudication panel (Claude, GPT-4o, Gemini, DeepSeek). Compares declared job specs against delivered artefacts and returns an ALLOW / HOLD / REJECT verdict. Answers: *did this agent actually deliver what it was paid for?*

On-chain signals are the foundation because they are **unfalsifiable** — you cannot fake a transaction history without spending real money — and **temporally rich** — patterns emerge over weeks and months of activity, making scores resistant to Sybil manipulation.

Tiered Trust Routing

AHS maps directly to payment-authorisation policy via the `/ahs/route/{address}` endpoint:

GRADE	AHS RANGE	ROUTING ACTION
A / B	80–100	Instant settlement — agent is trusted
C	60–79	Escrow — hold funds until delivery confirmed
D / F	0–59	Reject — agent does not meet minimum trust threshold

This turns AHM into a **trust gate** that sits before payment authorisation. Protocols and marketplaces call `/ahs/route` to get a one-word routing decision — `instant`, `escrow`, or `reject` — and enforce it programmatically. No manual review. No trust-me handshakes.

Ecosystem Intelligence — Key Stats

METRIC	VALUE
Agents scanned	**8,000+**
Registries covered	ACP (Virtuals), Arc, Olas, ERC-8004 (Base), Celo
Average AHS	61.3 / 100 (Grade D)
Agents scoring Grade A	< 1%
Agents below Grade B	91%
Zombie agent rate	3%
Nightly scan cadence	Daily, automated across all registries

AHM runs nightly scans across every major agent registry, maintaining the largest longitudinal dataset of agent health metrics in the ecosystem. Scores, trends, and grade distributions are available via the public dashboard at agenthealthmonitor.xyz/dashboard.

Integration

AHM is designed for programmatic consumption — by agents, protocols, and enterprise systems.

x402-native. All 14 endpoints support x402 micropayments (USDC on Base), meaning agents can pay per call with no API key, no subscription, and no human in the loop.

REST API with API key auth. Human developers purchase an API key via Stripe and authenticate with an `X-API-Key` header. Same endpoints, same data.

AHM Shield SDK. Drop-in middleware for agent frameworks. Install and enforce trust routing in three lines:

```
# Install
pip install ahm-shield
```

```
from ahm_shield import Shield

shield = Shield(api_key="ahm_sk_...")
decision = shield.route("0xABC1234567890abcdef1234567890abcdef12345")

if decision.action == "reject":
    raise Exception(f"Agent blocked — AHS {decision.score}, Grade {decision.grade}")
```

Endpoints (14 live)

ENDPOINT	X402 PRICE	PURPOSE
`GET /risk/{address}`	$0.01	Pre-transaction trust check
`GET /risk/premium/{address}`	$0.05	Premium risk with Nansen labels + PnL
`GET /counterparties/{address}`	$0.10	Know Your Counterparty
`GET /network-map/{address}`	$0.10	Wallet network map
`GET /health/{address}`	$0.50	Full health diagnostic
`POST /wash/{address}`	$0.50	Financial health scan
`GET /ahs/{address}`	$1.00	Agent Health Score (0–100)
`GET /ahs/route/{address}`	$0.01	Trust routing decision
`GET /report-card/{address}`	$2.00	Visual report card with benchmarks
`GET /alerts/subscribe/{address}`	$2.00/mo	Automated monitoring + webhooks
`GET /optimize/{address}`	$5.00	Operational efficiency report
`POST /ahs/batch`	$10.00	Batch scoring (up to 10 wallets)
`GET /retry/{address}`	$10.00	Retry failed transactions
`GET /agent/protect/{address}`	$25.00	Full autonomous protection

Pricing

For Humans (Stripe · API Key)

STARTER	PRO POPULAR	UNLIMITED
$9 one-time	**$39** one-time	**$99** / month
100 API calls	500 API calls	Unlimited API calls

All tiers include `X-API-Key` authentication and access to all 14 endpoints.

For Agents (x402 · Pay-Per-Call)

No API key required. Agents pay per call in USDC on Base via the x402 protocol. Prices range from $0.01 to $25.00 depending on endpoint complexity.

AHM Verify

MODE	PRICE	DESCRIPTION
Standard	$0.50 / verdict	6-model LLM panel, spec-vs-output comparison
Deep	$1.50 / verdict	Extended analysis with adversarial critique and confidence scoring

Design Partner Programme

AHM is offering select design partners **three months of free enterprise access** — unlimited API calls, priority support, and direct input into the product roadmap.

We are looking for protocols, marketplaces, and enterprises that are building agent-to-agent workflows and need trust infrastructure today.

Contact:

Pablo — Founder

Email: pablo@agenthealthmonitor.xyz

Twitter/X: @agenttrust

Links

RESOURCE	URL
Homepage	agenthealthmonitor.xyz
Dashboard	agenthealthmonitor.xyz/dashboard
AHM Verify	verify.agenthealthmonitor.xyz
Developer App	agenthealthmonitor.xyz/app
API Documentation	docs.agenthealthmonitor.xyz
ERC-8004 Identity	8004scan.io #32328
Twitter/X	@agenttrust

AHM — Trust layer for the agent economy.