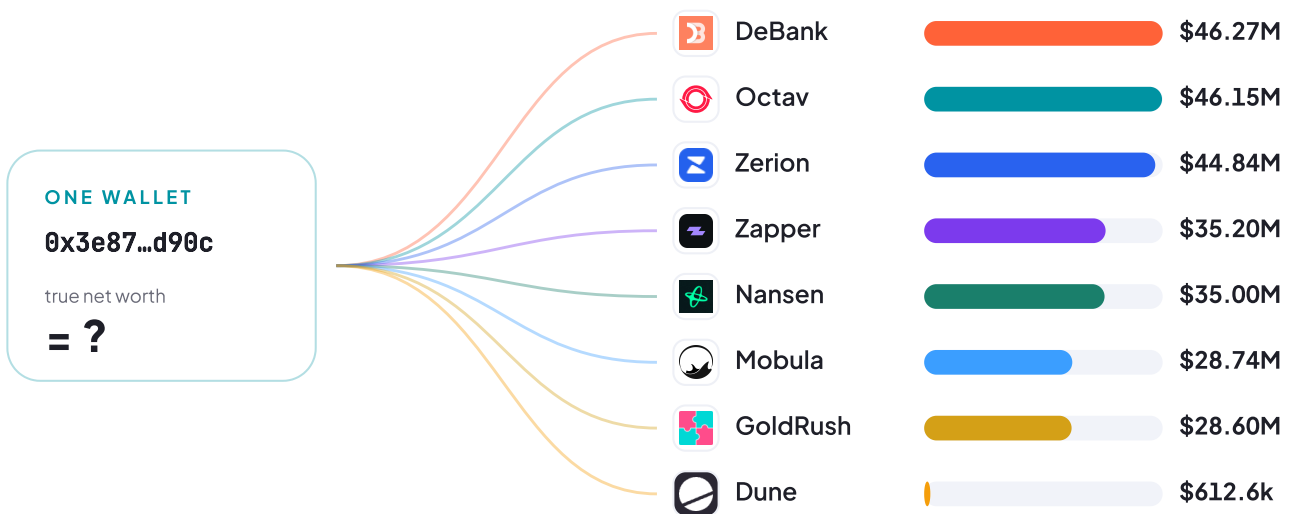


# The Crypto Portfolio API Benchmark

One wallet address, eight portfolio APIs, nine real-world wallet archetypes. A side-by-side measurement of coverage, DeFi decoding, accuracy, performance, cost and developer experience — every number reproducible from an open harness.

**8**PROVIDERS  
BENCHMARKED**9**

WALLET ARCHETYPES

**1**SHARED ENTRY  
POINT**0**HAND-TYPED  
NUMBERS

The same wallet, valued by eight APIs — from **\$0.6M** to **\$46M**. Which one is right, and why they differ, is what this benchmark measures.

## PROVIDERS BENCHMARKED



# What the data says

We run nine real-world wallets through eight portfolio APIs — one address at a time — and diff the results. The single biggest differentiator is **DeFi decoding**: whether an API turns a wallet's on-chain positions — lending, LP, staking, perps, options — into a portfolio, or just lists raw token balances. That one capability moves reported net worth by **10x to 600x** on active wallets.

## 10–600x

net-worth gap between DeFi-decoding and token-only APIs on active wallets






## 1 of 8

APIs that decode on-chain options (Derive) — the rest see only the collateral

## 0.2%

spread on a clean single-protocol wallet — proof the deltas are coverage, not noise

### Five findings


1. **DeFi decoding is the dividing line.** On the multi-sector whale, token-only APIs report \$612.6k–\$28.60M against \$46.27M from the full-portfolio APIs.
2. **Options and perps are barely covered.** Only  **Octav** decodes the Derive options book (\$1.33M vs ~\$401.6k elsewhere).
3. **Solana DeFi is uneven.**  **Octav** and  **Nansen** decode Kamino; most others see only the token side;  **DeBank** and  **Dune** don't cover Solana at all.
4. **Even strong APIs disagree.** On the yield wallet, coverage ranges from \$1.43M to \$3.48M — the difference is Pendle LP valuation.
5. **Price is not the same as coverage.** A clean Aave wallet reconciles across every EVM API to within 0.2%; the gaps elsewhere come from *what* each API sees, not how it prices.

### How to read this report


It is organised **by capability** (coverage, DeFi decoding, options, Solana, performance, cost, developer experience) rather than provider-by-provider, so you can weigh the dimensions that matter for your use case. Every figure is drawn from a single time-aligned snapshot captured 2026-07-02; the harness that produced it is open source. Octav publishes this report and is included in it — see *Methodology & disclosure*. We flag every place a number is estimated or a provider's own data is anomalous.

# The field, ranked by use case


There is no single winner — the right API depends on what you're building. The weighted leaderboard (full rubric on the scoring page) is shown alongside, but the use-case picks below matter more.

 **Octav** Deepest DeFi + options + Solana


Only API that decodes options, perps and Solana DeFi in one call — with a per-asset price source.

 **DeBank** EVM DeFi gold standard

Ties for the richest EVM decoding (health rates, perps, vesting); EVM-only, no Solana.









 **Zerion** Fastest & cheapest full portfolio

One call, ~320 ms, ~\$0.0006 — strong all-round coverage, NFTs and PnL.

 **Dune** Cheapest at scale (token data)



Free Sim tier for token balances across 95 chains; no DeFi decoding.

**Weighted overall score** (measured; rubric disclosed)

	DeBank	<div style="width: 86%; height: 10px; background: linear-gradient(to right, #ff7043, #ccc);"></div>	86
	Octav	<div style="width: 83%; height: 10px; background: linear-gradient(to right, #009688, #ccc);"></div>	83
	Zerion	<div style="width: 79%; height: 10px; background: linear-gradient(to right, #3954ab, #ccc);"></div>	79
	Zapper	<div style="width: 69%; height: 10px; background: linear-gradient(to right, #9c27b0, #ccc);"></div>	69
	Mobula	<div style="width: 66%; height: 10px; background: linear-gradient(to right, #2196f3, #ccc);"></div>	66
	Nansen	<div style="width: 53%; height: 10px; background: linear-gradient(to right, #43a047, #ccc);"></div>	53
	Dune	<div style="width: 51%; height: 10px; background: linear-gradient(to right, #ffc107, #ccc);"></div>	51
	GoldRush	<div style="width: 48%; height: 10px; background: linear-gradient(to right, #ffc107, #ccc);"></div>	48

Dimensions: coverage 25 · DeFi 20 · chains 15 · cost 15 · performance 10 · DevX 8 · tooling 5 · features 2.

**Reading the numbers honestly**

Two providers returned anomalous values we *kept* rather than hide:  **Mobula** mispriced a dead token (LUNC) into an astronomical total on one wallet, and  **Nansen** double-counts a receipt token when its two portfolio endpoints are merged. Both are real, reproducible behaviours, flagged wherever they appear. See *Data-quality gotchas*.

# Why we built this

Portfolio APIs all promise the same thing — "give us an address, we'll give you the portfolio." In practice they disagree wildly. We wanted to know exactly how, and by how much, for two reasons.

## REASON ONE

### To make our own data better

Octav builds portfolio infrastructure. The fastest way to find our own blind spots is to line our output up against everyone else's on the same wallets. Building this benchmark already surfaced — and let us fix — a real parsing bug and a timeout behaviour on our side. Measuring competitors is how we pressure-test ourselves: where do we miss a protocol, mis-price a token, or under-cover a chain?

## REASON TWO

### To publish a transparent benchmark

Anyone choosing a portfolio API today has to trust marketing pages. There is no neutral, reproducible comparison of what these APIs actually return for real wallets. So we built one — with an open app you can run on any address, an open capture harness, and a commitment to report the numbers as they fall, including where Octav loses.

## The question that matters

A "portfolio" is more than a list of tokens. A serious wallet has money working inside protocols — supplied to Aave, LP'd in a Pendle pool, staked in EigenLayer, held as a Hyperliquid perp or a Derive option. The core question this report answers is simple: **which APIs actually see that money, and which only see the tokens sitting idle in the wallet?**

To answer it we assembled nine wallets that each stress a different capability — from a single clean Aave deposit (where everyone should agree) to a 54-protocol, 19-chain whale (where they don't) to Solana-native DeFi (where most EVM-first APIs go blind). We then called all eight APIs on each, at approximately the same moment, and measured everything.

### What "we" means here

This report is published by **Octav**, which is one of the eight APIs measured. We've tried hard to earn your trust rather than assume it: the harness is open, the rubric is disclosed, Octav's numbers get no special treatment, and we lead with use-case fit rather than a single trophy. Where Octav is slower or more expensive than a rival, the report says so plainly.

# How we measured

One backend, one entry point, one snapshot. Every provider is called through the same normaliser so results diff apples-to-apples, and every raw payload is cached so any number can be re-derived without spending an API credit.

## Single entry point

Each API sits behind one adapter that returns a normalised portfolio (tokens + decoded protocol positions). The frontend calls one endpoint; differences are the providers', not our plumbing.

## Time-aligned snapshot

All eight APIs were called on all nine wallets in one pass on 2026-07-02, so market moves don't distort the comparison. Latency is the provider's measured response time.

## Reconciled & reproducible

Where an API exposes its own net-worth total we cross-check our parse against it (Octav, Zapper and Mobula reconcile to 0.0%). Raw payloads are cached; the harness is open source.

## The nine wallet archetypes







ARCHETYPE	CHAIN	WHAT IT STRESSES
Multi-sector whale	EVM	54 protocols across 19 chains — lending, LP, staking, perps, options, prediction markets.
Options trader	EVM	Large on-chain options book on Derive (Lyra v2) plus vesting and staking.
Lending whale	EVM	Single, clean Aave V3 USDC supply — the control case where everyone should agree.
Yield / restaking	EVM	Pendle LP + Ethena + Swell/Renzo restaking — yield-token valuation stress test.
Curve / Convex farmer	EVM	Deep Curve/Convex/Frax LP + gauge positions across many pools.
Retail power user	EVM	Realistic retail wallet: DeFi + perps + options + Polymarket at small scale.
Solana DeFi whale	SOL	Multi-million Kamino lending book plus LSTs and SPL tokens.
Solana all-rounder	SOL	Jupiter Lend + Kamino + Drift + Jito across many SPL tokens.
Solana leveraged loop	SOL	Leveraged Kamino position (health rate 0.09) — value locked inside the protocol.

## What we deliberately did not do

We did not tune any adapter to flatter a provider, and we did not "correct" a provider's own data. When an API returns an anomalous value, we keep it and flag it. Cost figures are normalised to a single-address fetch from each provider's published pricing (verified against live docs where possible) and marked *estimated* where a rate is not public. Some capabilities — NFT valuation depth, historical accuracy — are out of scope for this edition and noted in *Limitations*.






# Eight portfolio APIs

They fall into three families: **full-portfolio** APIs that decode DeFi positions, an **analytics** platform with a portfolio endpoint, and **token-data** APIs that return balances only. That family line predicts most of what follows.


PROVIDER	CATEGORY	EVM	SOL	DEFI	CHAINS	CALLS	PRICING MODEL
 <b>Octav</b> 1-call full portfolio + price source	Full portfolio	✓	✓	✓	74+	1	Credits (never expire)
 <b>DeBank</b> DeFi-position gold standard (EVM)	Full portfolio	✓	-	✓	84+	3	Prepaid Units
 <b>Zerion</b> Fast 1-call, cheap, NFTs + PnL	Full portfolio	✓	✓	✓	64+	1	Subscription + quota
 <b>Zapper</b> GraphQL portfolioV2, 60+ chains	Full portfolio	✓	✓	✓	56+	1-4	Credits (3/query)
 <b>Mobula</b> Cheap multichain, built-in PnL	Full portfolio	✓	✓	✓	88+	2	Credits (1/chain)
 <b>Nansen</b> Analytics-grade, labels, SOL DeFi	Analytics + portfolio	✓	✓	✓	37+	2	Credits + subs
 <b>Dune</b> Sim API — token balances (free tier)	Wallet data API	✓	-	-	95+	1+	Sim free tier
 <b>GoldRush</b> Token balances, EVM + Solana	RPC · token API	✓	✓	-	100+	4 multichain · 1 SOL	Credits

✓ supported - not supported **Calls** = requests to assemble a full portfolio for one address



## FULL PORTFOLIO

 **Octav** ·  **DeBank** ·  **Zerion** ·  **Zapper** ·  **Mobula** — decode wallet tokens *and* protocol positions. The question is how deep and across which chains.

## ANALYTICS + PORTFOLIO

 **Nansen** — analytics-grade with labels and a DeFi-holdings endpoint; priciest per call, no token logos, strong on Solana.

## TOKEN DATA

 **Dune** ·  **GoldRush** — accurate token balances and prices across many chains, but **no DeFi decoding**. Excellent baselines; not full portfolios.

# Who sees the whole wallet

Coverage is the headline. On our multi-sector whale — 54 protocols across 19 chains — full-portfolio APIs land near \$46.27M, while token-only APIs see a rounding error of it.




Fig. 1 — Reported net worth for the multi-sector whale (0x3e87...d90c). Token-only APIs (Dune, and GoldRush's EVM tokens) miss the DeFi positions that hold most of the value.

## Coverage across every wallet

Each cell = the share of the reference net worth (highest credible value among the DeFi-grade APIs) that a provider captured. Darker = more complete. "!" = anomalous payload; "-" = not supported.









WALLET	OCTAV	DEBANK	ZERION	ZAPPER	NANSEN	MOBULA	GOLDRUSH	DUNE
Multi-whale	100%	100%	97%	76%	76%	62%	62%	1%
Options	100%	34%	34%	30%	33%	30%	30%	30%
Lending	100%	100%	100%	102%	205%	102%	102%	102%
Yield	100%	100%	100%	53%	54%	41%	42%	42%
Curve	100%	100%	82%	83%	48%	!	0%	0%
Retail	99%	100%	48%	50%	49%	31%	2%	37%
SOL whale	100%	-	37%	31%	70%	37%	37%	-
SOL all-round	100%	-	6%	6%	10%	6%	6%	-
SOL loop	100%	-	0%	0%	98%	0%	0%	-

### The control case

On the **lending whale** — a single clean Aave V3 USDC supply — every EVM API agrees to within 0.2% (\$83.64M ±). That is the point: when the wallet is simple, everyone is accurate — proving the large gaps elsewhere are about *coverage of complex positions*, not price noise. (🚩 Nansen's 205% is the receipt-token double-count, flagged later.)





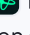
# The dividing line

Decoding a position means turning a receipt token (aUSDC, a Pendle LP, a Kamino obligation) back into what it represents — a supply, a debt, an LP with underlying assets, a reward. It is the single capability that separates a portfolio API from a token API.


PROVIDER	DECODES DEFI	POSITIONS SEEN*	HEALTH RATE	PRICE SOURCE	DEPTH
 Octav	✓	118	✓	✓	Lending, LP, staking, perps, options, prediction mkts
 DeBank	✓	103	✓	–	Lending, LP, staking, perps, options, vesting (EVM)
 Zerion	✓	94	–	–	Lending, LP, staking, deposits (+NFTs, PnL)
 Zapper	✓	81	–	–	Apps + positions; weak on Pendle-style yield tokens
 Mobula	✓	0	–	–	Advertised; inconsistent run-to-run this benchmark
 Nansen	✓	36	–	–	DeFi-holdings endpoint; strong on Solana, misses some LP
 Dune	–	0	–	–	None — token balances only
 GoldRush	–	0	–	–	None — token balances only (DeFi endpoints deprecated)

\* Total decoded protocol positions across the nine test wallets.  **Mobula** advertises DeFi but returned no positions on several wallets this run — see below.




## What good decoding looks like

On the yield wallet,  **Octav**,  **DeBank** and  **Zerion** value the Pendle LP correctly (~\$3.48M total).  **Zapper** and  **Nansen** show the Pendle app but price the LP token at a few thousand dollars — a known weak spot for yield-token valuation — landing near \$1.85M.

## Inconsistency is real

 **Mobula** returned a full DeFi breakdown on one snapshot and an empty one on the next for the same wallet (\$3.86M → \$1.43M). If you depend on DeFi positions, run-to-run stability matters as much as peak coverage.

## Beyond value: what the position tells you

Two APIs go further than a dollar figure.  **Octav** and  **DeBank** return a **health rate** for leveraged positions (how close to liquidation), and  **Octav** attaches a per-asset **price source** so you can audit where a valuation came from — the only API in the set that does.

# Options & perps: the blind spot

On-chain derivatives are where coverage collapses. Our **options trader** holds a large book on Derive (Lyra v2). Only one API turns that book into portfolio value; the rest see the collateral sitting in the wallet and stop there.




Fig. 2 — Options trader (0x43a9...ed7). **Octav** decodes the Derive options position (\$1.33M); every other API reports roughly the wallet-token value (~\$401.6k), a 3.3x difference.

### Options

Only **Octav** decodes the Derive book. **DeBank** and **Zerion** catch part of the collateral (~\$449.2k); the token-only APIs see just the wallet tokens. For anyone tracking structured-product or options desks, this is the difference between a real number and a wrong one.

### Perps

Perpetuals fare a little better. On the retail wallet, **Octav** and **DeBank** decode the Hyperliquid and Lighter positions others miss, which is why they report \$56.0k against ~\$27.4k from the DeFi-lite APIs. Perp equity vs. open-notional accounting also varies between providers — worth verifying per API.

### Why this matters disproportionately

Derivatives concentrate a lot of value in a few positions. An API that misses them doesn't just lose a line item — it can under-report a serious trader's net worth by an order of magnitude, as Fig. 2 shows.

# Solana splits the field

Solana is where EVM-first APIs go blind. Two of eight don't support it at all; of the rest, only some decode Solana DeFi (Kamino, Drift, Jupiter Lend) rather than just SPL token balances.




Fig. 3 — Solana DeFi whale (4yhX...Z3do). Octav (\$26.04M) and Nansen (\$18.21M) decode the Kamino lending book; token-focused reads land near \$9.67M; DeBank and Dune are EVM-only.

### Decode Solana DeFi

Octav and Nansen turn Kamino/Drift positions into value. On the leveraged loop wallet they report \$254.7k / \$249.7k while everyone else sees ~\$0.

### Token side only

Zerion, Zapper, Mobula and GoldRush return Solana SPL balances accurately (all ~\$9.67M on the whale's token side) but miss the value locked in protocols.

### No Solana

DeBank (EVM DeFi specialist) and Dune (Sim) don't cover Solana at all — a hard gap for any multi-chain product that includes SOL users.

### The takeaway

If Solana matters to you, the field shrinks fast. Full Solana *DeFi* coverage in the same call as EVM is rare — in this set, Octav and Nansen — and only Octav does it while also decoding EVM options and perps.

# Breadth of tokens and chains

Coverage isn't only about DeFi — it's also how many chains an API reaches and how well it filters spam. We don't score raw token counts (a wallet with 1,100 memecoins isn't "richer"); we report supported chains and what each API actually returned across the test set.





PROVIDER	CHAINS SUPPORTED	EVM	SOLANA	TOKENS RETURNED*	MAX CHAINS / WALLET
 Octav	74+	✓	✓	202	19
 DeBank	84+	✓	-	162	22
 Zerion	64+	✓	✓	201	14
 Zapper	56+	✓	✓	238	15
 Mobula	88+	✓	✓	166	7
 Nansen	37+	✓	✓	172	10
 Dune	95+	✓	-	152	10
 GoldRush	100+	✓	✓	133	9

\* Summed across the nine wallets after each provider's own spam filtering. Higher isn't automatically better — it reflects both reach and how aggressively an API filters dust.

## Chain reach vs. chain depth

Several APIs advertise 80–100+ chains, but reach is not depth:  **GoldRush** and  **Dune** cover many chains for *tokens* yet none for DeFi. The full-portfolio APIs cover fewer chains but decode positions on them — the more useful kind of breadth.

## One call vs. many

 **Octav** and  **Zerion** return every chain in a single call.  **GoldRush** batches its chains through a multichain endpoint (10 per call), and  **Dune** paginates — which shows up directly in the latency and cost pages that follow.

# Speed, calls & payload size

A portfolio call sits in a user's critical path, so latency and request count matter. The spread is wide: from a single ~320 ms round-trip to multi-second, multi-call assemblies.

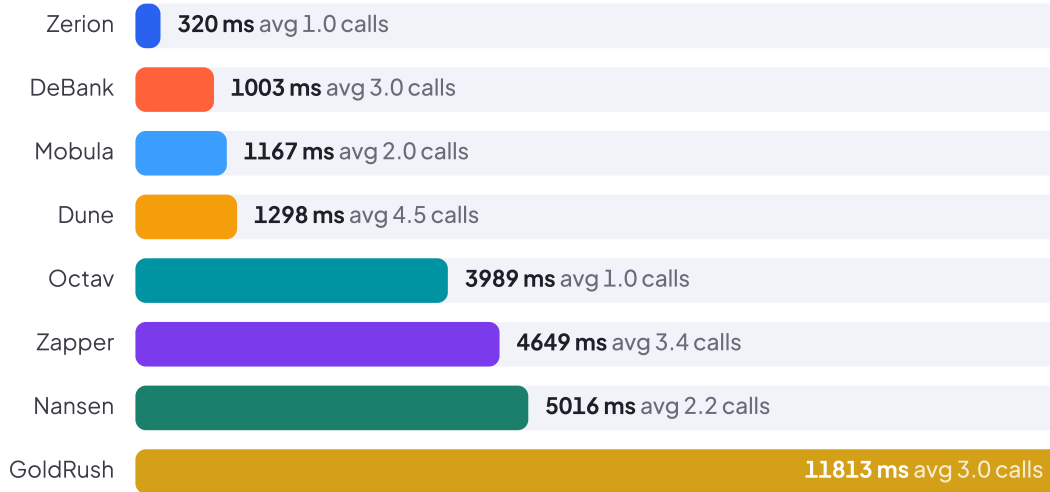


Fig. 4 — Median response time per provider across the test wallets, annotated with average calls per portfolio. One-call APIs cluster at the top; fan-out and heavy-sync APIs trail.

*Latency is the wall-clock time to assemble the full portfolio. For multi-call APIs, independent requests are issued in parallel, so the figure reflects concurrent execution — not the sum of each call. Call counts are averaged across the nine wallets, so they read as fractional: larger wallets trigger extra pagination pages and per-chain providers fan out, so the same API can make a different number of calls per wallet.*

**Fastest**

**Zerion**

~320 ms in a single call — the performance leader, even on large wallets.

**Heaviest payloads**

**Zerion**

Speed has a cost: Zerion returns the largest payloads (~750 KB avg), including one ~62 MB Solana response.

**Depth over raw speed**

**Octav**









Octav returns everything in one call but decodes deeply, so it can be slower on very large wallets.


**Depth and breadth both cost latency**

Latency tracks how much each API does per call — in two ways. The APIs that *decode positions* do more work, so the deepest can be slower on the largest wallets. And the APIs that query *many chains* pay for breadth: **GoldRush** is token-only yet the slowest here because it sweeps all ~39 EVM chains through the multichain endpoint. Both are trade-offs to weigh against how complete you need the portfolio to be.



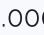

# What it costs at scale

Pricing models don't compare directly — credits, prepaid units, compute units, subscriptions. We normalise to **\$ per portfolio fetch**, counting how many *calls* each API needs and how many *credits* those calls cost (they're rarely 1:1), at each provider's first paid package.





PROVIDER	PRICING MODEL	CALLS / FETCH	\$/FETCH	1K / MO	100K / MO	1M / MO	OUR TEST COST
 Dune	Free CU tier <sup>a</sup>	1-5	\$0	\$0	\$0	\$0	<b>\$0</b>
 Zerion	Subscription (Builder)	1	\$0.0006	\$0.6	\$60	\$600	<b>\$0.005</b>
 Zapper	Credits (3/query) <sup>b</sup>	1-4	\$0.0031	\$3	\$307	\$3,070	<b>\$0.03</b>
 Mobula	Credits (1/chain) <sup>b</sup>	2	\$0.0051	\$5	\$511	\$5,110	<b>\$0.05</b>
 GoldRush	Credits <sup>b</sup>	4.1	\$0.0070	\$7	\$700	\$7,000	<b>\$0.06</b>
 DeBank	Prepaid Units <sup>c</sup>	3	\$0.0116	\$12	\$1,160	\$11,600	<b>\$0.1</b>
 Octav	Credits (never expire)	1	\$0.0250	\$25	\$2,500	\$20,000	<b>\$0.23</b>
 Nansen	Credits + subs (Pro) <sup>d</sup>	2	\$0.0544	\$54	\$5,440	\$54,400	<b>\$0.49</b>

Verified against live pricing 2026-07-02. First paid package, no free tier (except  Dune); **1 call ≠ 1 credit** — calls are mapped to each provider's billing unit. "Our test cost" = the real credits/units our 9-wallet benchmark consumed × that rate. Octav = ground truth. <sup>a</sup>Dune Sim is being sunset 1 Aug 2026. <sup>b</sup>Credit cost is endpoint-dependent (per-chain / per-query). <sup>c</sup>DeBank -20% on ≥\$100k prepay. <sup>d</sup>Nansen publishes no \$/credit; we use its base Pro plan (\$49 / 2,000 credits = \$0.0245/credit). Enterprise/volume tiers are sales-gated.

## Cheap at scale






 Dune is free (token data, but sunsetting);  Zerion is the cheapest paid full portfolio (~\$0.0006/fetch);  Zapper and  Mobula stay well under a cent.

## Premium




 Nansen is the most expensive per fetch (~\$0.054 — you pay for analytics and labels);  Octav at \$0.025 sits mid-high, reflecting the decoding depth per call. Our 9-wallet run costs most on  Nansen (\$0.49) and  Octav (\$0.23) — no free tier — but pennies overall.

# How hard is it to integrate




Beyond the data, what is it like to build against? We weigh auth friction, whether a full portfolio is one call or a fan-out, SDK/docs quality, and Solana support — from actually integrating all eight.

PROVIDER	AUTH	CALLS	DEVX SCORE	INTEGRATION NOTES
 Octav	Bearer key	1	90	1-call, price source, MCP; slow on huge wallets
 Zerion	Basic (key)	1	88	1-call, fast, cheap; large payloads
 Mobula	Header key (flexible)	2	80	Cheap, simple; DeFi inconsistent run-to-run
 DeBank	AccessKey header	3	78	Rich EVM decoding; 3 calls; EVM-only
 Dune	X-Sim-API-Key	1+	74	Simple, free; token-only; DeFi API sunset
 Zapper	x-api-key (GraphQL)	1-4	72	GraphQL flexibility; verbose; Pendle mispriced
 Nansen	apikey header	2	66	Two endpoints; merge carefully (double-count)
 GoldRush	Bearer key	4 multichain · 1 SOL	60	Multichain balances endpoint; DeFi endpoints deprecated

## Smoothest to integrate

 Octav and  Zerion : one call, clean auth, a coherent shape that maps straight onto a normalised portfolio.  Mobula is close and cheap. These got us to a correct portfolio fastest.

## Sharper edges

 GoldRush fans out per chain and its DeFi endpoints are deprecated;  Zapper 's GraphQL is powerful but verbose and its Pendle valuation is off;  Nansen splits a portfolio across two endpoints that double-count if you merge them naively. All integrable — just budget more time.

## A note from building it

Every integration gap we hit is documented in the open harness — including bugs we fixed on our *own* Zerion adapter (a receipt-token double-count) and Octav's timeout behaviour. The DevX scores reflect real integration effort, not marketing claims.

# The tooling & agent stack

In 2026 an API is only as good as what you can build against it — and the fastest-moving layer is AI. We surveyed each provider's official **SDKs**, **MCP servers**, **x402 pay-per-call**, **AI-agent kits**, **CLI** and streaming (community-only tooling doesn't count).

PROVIDER	SDK	MCP	X402	AI AGENT	CLI	STREAM	SCORE	HIGHLIGHTS
GoldRush	✓	✓	✓	✓	✓	✓	90	TS + Python SDKs, React kit, MCP, CLI, x402, GraphQL/WS streaming
Zerion	-	✓	✓	✓	✓	✓	86	Hosted MCP, agent skill, CLI, x402, Kafka streams + webhooks
Octav	~	✓	✓	✓	✓	-	82	Community SDK, official MCP (14 tools), Rust CLI, x402, AI-dev docs + skill
Zapper	-	✓	✓	✓	-	~	78	Official MCP, x402, agents.txt; GraphQL (no subscriptions)
Nansen	-	✓	✓	✓	✓	-	74	Official MCP, CLI, agent hub, x402; no SDK / streaming
Mobula	✓	-	✓	~	-	✓	64	TS SDK, GraphQL + WebSocket + webhooks, x402; no MCP/CLI
DeBank	~	~	-	-	-	-	40	REST only; SDK & MCP are community, not official
Dune	-	-	-	-	-	~	38	Sim: webhooks only; sunseting Aug 2026

✓ official ~ partial / community - none **MCP** = Model Context Protocol server for AI assistants

### The MCP + x402 wave

Six of eight ship an **official MCP server** and five support **x402** pay-per-call (Base/Solana) — the new default for agent access. Octav , Zerion , Zapper , Nansen and GoldRush lead the agent-native shift.

### Richest stack

GoldRush is the most complete — official TS + Python SDKs, a React kit, MCP, CLI, x402 and GraphQL-over-WebSocket streaming. Zerion is close and the most AI-forward (only real Kafka streaming + webhooks here).

### Thinnest

DeBank is REST-only with *no* official SDK/MCP/CLI/x402 (all community). Dune 's Sim API offers only webhooks and is **sunseting Aug 2026**. Both are usable, but you build the tooling yourself.


### Why this is its own dimension

An MCP server or x402 endpoint can turn a week of integration into an afternoon — and increasingly your *users* AI agents call these directly. We weight tooling lightly (5%; data still comes first), but it's the fastest-growing reason teams switch providers.


## Where the data lies (and we kept it)

A benchmark you can trust has to show the ugly parts. Two providers returned values we could have quietly "fixed" — we didn't. Here is exactly what happened and how to defend against it in production.


### Mobula — a mispriced dead token

On the Curve farmer,  Mobula returned a unit price of  $\sim 3.2 \times 10^{18}$  for LUNC (a defunct token), producing a portfolio total of  $\sim 10^{23}$  dollars. The wallet's other values are reasonable; a single bad price feed on one long-tail token blew up the sum. **We keep the raw number and exclude it from charts** — it is a real Mobula data issue, not ours. *Defence*: sanity-bound unit prices and cross-check long-tail tokens against a second source.

### Nansen — a double-counted receipt token


 Nansen splits a portfolio across two endpoints: token balances and DeFi holdings. On the Aave wallet the aUSDC receipt token appears in *both* — once as a wallet token and again as the decoded Aave position — so a naive merge reports  $\sim 205\%$  of the true value (\$171M vs  $\sim$ \$84M). It only surfaces when a receipt token lands in the balance list. **We keep it and flag it** as an integration gotcha. *Defence*: dedupe receipt tokens that also appear as protocol positions.

### What this says about the field

Neither issue is exotic — a bad long-tail price and an endpoint-merge double-count are exactly the failure modes you inherit in production. The lesson isn't "these two are bad"; it's that **every** portfolio API needs sanity checks around it, and an API that decodes positions *and* exposes its price source (only  Octav here) makes those checks easier.

# A transparent scorecard




We resisted reducing this to one number, but a disclosed rubric is useful if you read it as a starting point, not a verdict. Each dimension is scored 0–100 from the measured data; the overall is a weighted sum. Re-weight it for your use case and the order changes.



PROVIDER	COVERAGE	DEFI	CHAIN	COST	PERF	DEVELOPER	TOOLING	FEATURES	OVERALL
 DeBank	100	100	80	79	74	78	40	63	<b>86</b>
 Octav	100	100	78	54	52	90	82	75	<b>83</b>
 Zerion	48	100	69	99	96	88	86	50	<b>79</b>
 Zapper	50	100	61	94	27	72	78	38	<b>69</b>
 Mobula	34	55	92	91	79	80	64	63	<b>66</b>
 Nansen	54	100	43	0	31	66	74	25	<b>53</b>
 Dune	34	0	90	100	59	74	38	25	<b>51</b>
 GoldRush	30	0	100	87	26	60	90	25	<b>48</b>

Weights: coverage 25 · DeFi 20 · chains 15 · cost 15 · performance 10 · DevX 8 · tooling 5 · features 2. Coverage uses the neutral reference (highest credible value among DeFi-grade APIs per wallet), so no provider is defined as ground truth.

## Read it as a spread, not a podium

The top three (DeBank 86 · Octav 83 · Zerion 79) are within a few points and win on different axes:  **DeBank** and  **Octav** tie on coverage and DeFi decoding, with DeBank edging ahead on chain breadth and cost and Octav leading on developer experience, Solana and its unique price source;  **Zerion** wins on raw speed and cost. Weight the dimensions you actually ship on and pick accordingly — which is what the next page does.

# Provider scorecards

Strengths, watch-outs and the use case each API is genuinely best for.


**DeBank**

Full portfolio

**86/100**

DeFi-position gold standard (EVM) · 84+ chains · EVM · 3 call(s) · \$0.0116/fetch

**STRENGTHS**

- Richest EVM DeFi decoding (perps, options, vesting)
- Health rates; strong reconciliation
- Fast (~1s), 3 clean calls

**WATCH-OUTS**

- EVM-only — no Solana
- Prepaid-units billing
- No price source / PnL

**Best for:** EVM-first products that want gold-standard DeFi decoding and don't need Solana.


**Octav**

Full portfolio

**83/100**

1-call full portfolio + price source · 74+ chains · EVM + Solana · 1 call(s) · \$0.025/fetch

**STRENGTHS**

- Only API to decode options, perps & Solana DeFi in one call
- Per-asset price source + health rates
- Reconciles to its own totals at 0.0%

**WATCH-OUTS**

- Mid-high cost (~\$0.025)
- Can be slower on very large wallets
- No NFTs

**Best for:** Products needing the deepest, broadest single-call coverage — derivatives, Solana DeFi, multi-chain net worth.


**Zerion**

Full portfolio

**79/100**

Fast 1-call, cheap, NFTs + PnL · 64+ chains · EVM + Solana · 1 call(s) · \$0.0006/fetch

**STRENGTHS**

- Fastest (~320 ms), one call
- Cheapest paid full portfolio (~\$0.0006)
- NFTs + PnL + 24h change; EVM & Solana tokens

**WATCH-OUTS**

- Very large payloads (up to ~62 MB)
- Misses options/perps & most Solana DeFi
- Subscription quota limits

**Best for:** Consumer apps that need a fast, cheap, broad portfolio and can live without derivative/Solana-DeFi depth.


**Zapper**

Full portfolio

**69/100**

GraphQL portfolioV2, 60+ chains · 56+ chains · EVM + Solana · 1-4 call(s) · \$0.0031/fetch

**STRENGTHS**

- GraphQL flexibility, 60+ chains incl. Solana
- Decodes most EVM DeFi + NFTs
- Field-selectable payloads

**WATCH-OUTS**

- Under-values Pendle-style yield tokens
- 3-4 calls, slower
- Verbose to integrate

**Best for:** Apps already on GraphQL that want broad app coverage and can tolerate yield-token gaps.

**Mobula**

Full portfolio

**66/100**

Cheap multichain, built-in PnL · 88+ chains · EVM + Solana · 2 call(s) · \$0.0051/fetch

**STRENGTHS**

- Cheap multichain incl. Solana
- Built-in PnL & allocations
- Simple, flexible auth

**WATCH-OUTS**

- DeFi coverage inconsistent run-to-run
- Long-tail price outliers (LUNC)
- Weaker decoding depth

**Best for:** Cost-sensitive multichain apps that need tokens + light DeFi and add their own sanity checks.**Nansen**

Analytics + portfolio

**53/100**

Analytics-grade, labels, SOL DeFi · 37+ chains · EVM + Solana · 2 call(s) · \$0.054/fetch

**STRENGTHS**

- Decodes Solana DeFi (Kamino, Drift)
- Analytics, labels, smart-money context
- Broad chain set

**WATCH-OUTS**

- Most expensive (~\$0.054)
- Two endpoints double-count if merged naively
- No token logos

**Best for:** Analytics and intelligence products where labels and Solana DeFi matter more than price.**Dune**

Wallet data API

**51/100**

Sim API — token balances (free tier) · 95+ chains · EVM · 1+ call(s) · \$0 (free tier)

**STRENGTHS**

- Free at benchmark scale
- Token balances across 95 chains
- Simple auth, reliable

**WATCH-OUTS**

- No DeFi decoding at all
- EVM-only
- DeFi Positions API sunset

**Best for:** Dashboards and internal tools that need cheap, wide token balances and no protocol positions.**GoldRush**

RPC · token API

**48/100**

Token balances, EVM + Solana · 100+ chains · EVM + Solana · 4 multichain · 1 SOL call(s) · \$0.007/fetch

**STRENGTHS**

- Token balances across EVM + Solana
- Multichain endpoint — all EVM chains in a few calls
- Broad chain coverage

**WATCH-OUTS**

- No DeFi decoding (endpoints deprecated)
- Slow — full-chain multichain query ~13 s
- Trial keys suspend under bursts

**Best for:** Token-balance baselines and multi-chain wallet views where DeFi positions aren't required.

# What we couldn't measure

Honesty about scope is part of a credible benchmark. This edition has real boundaries.

- **Nine wallets, one moment.** A time-aligned snapshot is great for fairness but is a single point in time. Recurring/streaming accuracy and cold-start behaviour over days are not captured here.
- **Cost is normalised, not billed.** We convert each provider's published pricing to \$/fetch; real invoices depend on plan tier, volume discounts and caching. Figures marked estimated are not official.
- **NFTs are out of scope.** Only some APIs value NFTs and methods vary wildly, so we excluded them to keep the net-worth comparison fair. A future edition may add an NFT track.
- **Reference, not truth.** "Coverage %" is measured against the highest credible value among the DeFi-grade APIs, not an absolute ground truth — no oracle exists for a live wallet's true net worth.
- **Perp accounting varies.** Equity vs. open-notional conventions differ between APIs; we report totals as returned and flag the ambiguity rather than force a convention.
- **Provider plans differ.** Some capabilities (labels, extra chains, streaming) sit behind higher tiers than the keys we tested; where a product exists but our key couldn't reach it, we say so.

## Everything here is reproducible

The capture harness, the normaliser, the scoring rubric and this generator are open. Run the app on your own address, or re-run the benchmark, and you'll get these numbers — or newer ones.

# Pick by what you're building

There is no universal best portfolio API — there is a best one for your use case. The data points to clear picks.

## Octav

### You need the deepest, broadest coverage

Options, perps, Solana DeFi and EVM DeFi in one call, with a price source to audit valuations. You pay for it in cost and, on giant wallets, latency.

## DeBank

### You're EVM-only and want gold-standard DeFi

The richest EVM decoding in the set. Skip it only because you need Solana or a price source.

## Zerion

### You want fast, cheap and broad

One ~320 ms call, ~\$0.0006, NFTs and PnL. Accept that options, perps and Solana DeFi won't be decoded.

## Dune

### You just need token balances cheaply

Free Sim tier, 95 chains, dead simple — as long as you never need a decoded DeFi position.

## The one thing to take away

Decide first whether you need **positions** or just **balances**. That single choice removes half the field. If you need positions — and any serious wallet has them — the depth of DeFi, options, perps and Solana decoding is what separates a portfolio that's right from one that's off by an order of magnitude. Everything else — cost, speed, DevX — is a tie-breaker among the APIs that clear that bar.

### Why we're comfortable publishing this

Octav places well here on coverage and decoding, and worse on cost and raw speed — and we've shown both. If the numbers ever stop favouring an honest read, the harness is open for anyone to check. That's the point: a benchmark is only worth something if you can reproduce it.

# Run it yourself

## Try the live app

Paste any EVM or Solana address at **benchmark.octav.fi** and see all eight APIs diffed in real time.

## Read the harness

The capture harness, normaliser and this report generator are open at **github.com/Octav-Labs/portfolio-api-comparator**.

## Reproduce the numbers

Every figure is derived from a cached snapshot — re-run `capture.mjs` on the same wallets to regenerate them.

---

Published by Octav. Octav is one of the eight APIs measured; the benchmark, rubric and code are open so the comparison can be checked and challenged. Provider names and logos are trademarks of their respective owners. Figures reflect a single snapshot captured 2026-07-02 and each provider's published pricing at that date; treat estimated costs as directional. This document is informational, not financial or procurement advice.