



WHITEPAPER · VERSION 1.3 · 17 SEPTEMBER 2026

Builder intelligence for Robinhood Chain

How HEY finds the projects that are still building on Robinhood Chain, measures what they have shipped from public evidence, and shows where builder momentum and market attention diverge.

SCOPE

Robinhood Chain, chain id 4663

SCORING VERSION

hbm-v3 (live on
heyresearch.xyz/methodology)

DATA AS OF

2026-09-03 UTC, production,
unless dated otherwise

This document describes a research product and its methodology. Nothing in it is investment advice, an offer, or a forecast. Activity status, Build Momentum, Discovery Gap and Still Building are observations about past, publicly verifiable activity and past market context. They are not predictions of price, and HEY never claims that continued building affects a token's value. The \$HEY token launched on 2026-09-10; section 11 states what is live, what is not, and how the launch differed from what the first edition of this paper said was intended.

Contents

1	Abstract	3
2	The problem: markets measure price, not building	3
3	Thesis: the Discovery Gap	4
4	The core model	4
5	Evidence and provenance	6
6	Methodology	7
7	Integrity by design	9
8	The product	11
9	For builders	14
10	Coverage today (2026-09-16)	14
11	\$HEY and the research economy	17
12	Architecture and operations	19
13	Roadmap	20
14	Risks and limitations	21
15	Glossary and references	22

Figures: 1 Core model · 2 Data-source tiers and polling · 3 From evidence to Discovery Gap · 4 Build Momentum weights · 5 The \$HEY research flywheel · Screenshots of the live product (captured 2026-09-03; the surfaces in \$8.4 and \$8.5 shipped after those were taken).

1. Abstract

Public markets for tokens report price, volume, liquidity, holders and social attention in real time. None of those figures answers the question that matters most to anyone evaluating an early ecosystem: *who is actually building, what have they shipped, and which of them has the market not noticed yet?* HEY Research Lab is a builder-discovery and project-showcase layer for Robinhood Chain (chain id 4663) built to answer that question from public evidence.

HEY indexes projects from every free public source that names the chain — eleven launchpad factories read directly from on-chain logs, launchpad and registry APIs, GitHub, project websites and changelogs, and market aggregators — and refuses to open a project page without evidence that a project exists behind a contract. Each concrete deliverable a project produces is recorded as a **ShipEvent** with its source and a verification state. From those events, deterministic and versioned rules derive an **activity status** (Shipping, Active, Quiet, Dormant, Resumed, Unknown), a 0–100 **Build Momentum** score, and — only where market context exists — a **Discovery Gap**: the difference between a project's build-momentum percentile and its market-attention percentile within the published cohort. **Still Building** marks projects whose verified building continued through a HEY-tracked market drawdown.

Three design commitments define the product. First, no wallet, holder, PnL or smart-money data exists anywhere in the system; a schema test fails the build if a column so much as mentions one. Second, no market figure and no payment can move a score: price, liquidity and volume are context and filter data only, and the planned \$HEY token is structurally separated from organic ranking by tests that fail the suite if a scoring module references the token ledger. Third, unknown stays unknown: a project without source coverage reads "Activity unknown", never "Dormant"; a token without a fresh market reading reads "Unavailable", never \$0.

As of 2026-09-03 the production catalogue holds 2,525 published projects screened from 856,702 token candidates across 11 launchpads, 2,204 of 2,214 published token cards carry a source-labelled market-cap reading, and 2,428 projects (96.2%) carry an evidence-based narrative. Ship-event coverage is the thin part — 490 events across 209 projects at the 2026-09-02 baseline — and this paper says so plainly. Sections 6 and 10 give the method and the numbers; section 14 the limits.

2. The problem: markets measure price, not building

A new chain produces two very different kinds of activity at the same time. One is easy to see: token launches, pools, price moves, trading volume, holder counts and social attention, all of it surfaced instantly by DEX aggregators, launchpads and terminals. The other is slow, quiet and mostly invisible: a team publishing releases, deploying and upgrading contracts, writing documentation, shipping integrations, or — for a meme community — delivering games, tools and creative work.

The instruments that exist today measure the first kind. A price chart, a volume leaderboard or a holder-distribution view will tell you that a token moved. It will not tell you whether anyone is still working on the project behind it, what they delivered last week, or whether the evidence for that delivery is public and checkable. Wallet-analytics products answer a different question again — *which addresses bought* — and answer it at the cost of building surveillance infrastructure that has nothing to do with whether a product exists.

A token is not a project. Anyone can deploy an ERC-20; market cap, liquidity, volume and a price move say nothing about whether a team is building.

Robinhood Chain makes the gap concrete. Between its first launcher (Hoodit, active from block 61,869) and the launchpads live in September 2026, HEY's factory scan has recorded on the order of 855 thousand token-launch candidates on the chain (2026-09-03). Roughly 2.5 thousand of those pass an evidence gate — an official website, documentation, a public repository, a registry listing, verified contract source, recent pushes or verified ships. The rest are launch records: real contracts, real launch transactions, and no observable project. A discovery product that starts from the token list will show the noise; one that starts from the price will show the attention. Neither shows the builders.

The consequence for the ecosystem is a discovery lag. Projects that keep shipping while their tokens receive little attention are typically discovered weeks after the work is visible on GitHub, in a changelog or on chain. HEY exists to shorten that lag by measuring building directly, from evidence, and displaying it before — and independently of — market attention.

3. Thesis: the Discovery Gap

HEY is built around one observable research hypothesis:

HYPOTHESIS

Market attention and builder activity do not always move together. Some projects remain small or ignored while their teams continue releasing code, contracts, integrations, documentation, product features, community tools or creative work. Discovery lags shipping, and the lag can be measured.

The measurement is a comparison of two independent facts, shown separately before any label is derived. **Build Momentum** is computed only from ShipEvents — the input type carries no market field, so the two are independent by construction. **Market attention** is a cheap, log-scaled percentile of market cap or FDV, liquidity and 24-hour volume within the published cohort. The **Discovery Gap** is the build-momentum percentile minus the market-attention percentile:

High Builder Momentum + low current market attention = Discovery Gap. It is a research signal about the past. It is not a valuation, not a prediction, and not a recommendation.

The framing matters and is enforced in copy as well as code. A high positive Discovery Gap means exactly one thing: *this project is shipping more actively than its current relative market footprint would suggest*. HEY never renders "undervalued", "ready to moon", "next 100x" or "guaranteed comeback", and no feature may claim that continued building affects a token's price. The flagship **Still Building** label is likewise factual and historical — market attention fell, verified building continued afterwards — and its copy is limited to statements such as "Down 66.9% from the HEY-tracked high, with 2 verified updates since."

This is why HEY is deliberately not a trading terminal or an analytics dashboard. The product's job is to make the builder story legible first — identity, what they ship, when, with what evidence — and to keep market context where it belongs: as context.

4. The core model

Everything HEY stores and shows follows one proprietary chain of objects. Each link is derived only from the link before it, so every score can be explained back to the evidence that produced it.

FIGURE 1 · THE CORE MODEL




Figure 1. The HEY core model. Market data is drawn from cached snapshots and enters the model at exactly one point, the Discovery Gap; it cannot influence activity status or Build Momentum, whose input types carry no market field.

Source: PRD V4 §§9-13; packages/scoring.

4.1 Project and identity

A project is a page HEY has decided to publish. Its identity is the pair (chain_id, contract_address) for a token project, and a name plus sources for a builder or protocol without a token — a ticker is how a reader finds a record, never what the record is. Two projects with the ticker ALPHA are two results, not one. A page is published only when the project shows presence on the chain in one of four ways (a token, a verified contract, a registry listing, or a repository fingerprinted for the chain) *and* carries enough evidence that a project exists behind it. Market cap, FDV, liquidity, volume or a price move never publish a page at any size, and tests assert this. One builder is one project: a team's contracts, sdk and app repositories become sources on one page named for the team.

4.2 Narrative

A narrative (AI Agents, DEX, StockFi, Infrastructure, Meme, ...) is a discovery label stored in the database, not a hard-coded enum, so new narratives can be added without a deploy. A project has one primary and up to four secondaries. A narrative never touches a score.

4.3 Sources and ShipEvents

Sources are the places HEY reads for a project: its website and documentation, feeds and changelogs, repositories, contracts, launch records and builder submissions. A **ShipEvent** is a structured record that the project did something concrete — a product launch, feature or app release, GitHub release, contract deployment or upgrade, integration, SDK or API release, documentation update, roadmap milestone, or, for memes, a game, tool, creative drop or community event. A raw commit is not a ship: code activity is aggregated into one event per project per UTC day ("Active development: 63 commits across 1 contributor") and its weight is capped low on purpose. Announcements are recorded but carry no build weight.

4.4 Self-reported and verified stay distinct

Every ShipEvent carries a verification state: SELF_REPORTED (a builder told HEY), SOURCE_LINKED (it came from a source registered to that project), PUBLICLY_VERIFIED or ADMIN_VERIFIED (HEY confirmed it against a public source such as a GitHub release or a contract), DISPUTED and RETRACTED (the record stays, its weight does not). Only source-linked or better events can move a status or a score, so a project cannot talk its way into looking active. The distinction is visible everywhere: verified labels render in the standard tone, self-reported ones in italic muted type; a project page states "Verified owner" only after a claim has been proved, and "Submitted by a builder; ownership has not been verified" otherwise.

5. Evidence and provenance

HEY is public-data-first by rule: every integrated source is a free public endpoint and the paid data-API cost is \$0 (source registry, verified 2026-09-01 and re-verified per source through 2026-09-05). Every third-party call goes through a source adapter in one package; adapters validate payloads at runtime, apply timeouts and budgets, send conditional requests, and ship with saved fixtures so continuous integration never depends on a live provider. Public pages never call a provider: they render from HEY's own database, and a test fails the build if the web app imports an adapter.

FIGURE 2 · DATA-SOURCE TIERS AND POLLING




Figure 2. Source tiers and the HOT/WARM/COOL/COLD polling cadence. Budgets and cadences are configuration values from docs/SOURCE_REGISTRY.md, packages/domain/src/refresh/tier.ts and the 2026-09-04 audit; all are free public endpoints. Blockscout answers HTTP 403 to programmatic requests and is deliberately not worked around.

5.1 Reading launchpads from the chain

Launchpads are the largest source of new projects on Robinhood Chain and the least reliable if read through their own listings. HEY therefore reads launch factories directly: each of the 11 indexed launchpads is registered as one or more (*factory address, launch-event topic*) pairs found in the launchpad's own app bundle or by tracing a known token's mint transaction, and the topic is the one read from a live log, never computed from a published ABI. Where the event carries the creator's name, ticker and metadata (Flap, Clanker, hood.fun, Robinlaunch, EasyA Kickstart) the identity is decoded from event data; where it does not (Pools, PAIR, Hoodit, Robinpad, Pons) the launch stays a candidate until name() and symbol() are read from the contract. An hourly scan spends at most 360 log requests across all factories, stops 64 blocks behind the head for reorg safety, and resumes from persisted sync state; a launch whose event carries neither name nor symbol never becomes an address-named page.

5.2 Provenance and conflicts

Every meaningful ShipEvent keeps an evidence row (490 of 490 at the 2026-09-02 baseline) and a stable external id for de-duplication. Every market snapshot records which provider produced it and when. Launch provenance — launchpad, factory and launch transaction — is written on the project row at promotion. Contradictory facts are never silently overwritten: when a repair demotes a launchpad page that had been recorded as a project's website, or a repository turns out to be a fork or a template, a source_conflicts record preserves the rejected value and the reason. Disputed and retracted events remain in the timeline with zero weight.


5.3 What is structurally absent

- **No wallet or holder analytics.** No holder lists, wallet histories, PnL, smart-money labels or clustering are collected or stored. Provider payloads that carry holder fields (Virtuals, pair.fund) are validated with schemas that do not declare them, and the fixtures prove the fields never come through. A schema test fails the build if a column or migration mentions one.
- **No LLM in any request path.** AI enrichment is optional, asynchronous, cached by content hash and disabled by default; narrative assignment covers 96.2% of the published catalogue with deterministic rules and no model.
- **No paid placement as evidence.** DEX Screener boost lists are read only to learn which tokens carry a team-written profile; the spend is dropped inside the adapter and never stored.
- **No full-chain history.** Chain RPC is used for verification and factory logs, not for indexing from genesis; old market snapshots are downsampled by hour after 7 days and by day after 90.

6. Methodology

Every rule in this section is deterministic — the same evidence always produces the same result — and every weight and threshold lives in one versioned configuration file (packages/scoring/src/config.ts). The scoring package cannot read the database and its input types carry no price, holder, wallet or volume field, so those quantities cannot reach a score even by accident. The version in force, hbm-v3, is stored with every score snapshot and published on heyresearch.xyz/methodology.

FIGURE 3 · FROM EVIDENCE TO STATUS, MOMENTUM AND GAP



Cohort for both percentiles: published projects only. Outputs describe the past; they are not valuations, predictions or recommendations.

Figure 3. The derivation chain. Activity status and Build Momentum read only meaningful ShipEvents; market context joins at the Discovery Gap and nowhere earlier. Source: packages/scoring/src/{significance,activity,hbm,discovery-gap,config}.ts, PRD V4 §§10–13.

6.1 What counts as building

An event is *meaningful* when it is a real build signal *and* is backed by something better than an unverified claim. Strong types are product, app, feature, game and GitHub releases, contract deployments and upgrades, integrations, API and SDK releases, roadmap milestones and community tools. Supporting types are documentation, demos, design and creative drops, community events, founder build updates and aggregated

code activity. Announcements and uncategorised items are never evidence of building. For meme projects, creative and community deliverables are real output; a meme is never required to pretend it ships software. Since hbm-v3 a project's code activity counts once per UTC day however many repositories produced it, so a builder page with four repositories no longer earns four events for four single commits.

ACTIVITY THRESHOLDS (DAYS)		BASE EVENT WEIGHTS (RECENCY AND SIGNIFICANCE INPUTS)	
RULE	VALUE	EVENT TYPE	WEIGHT
Shipping: last meaningful ship within	7	Product launch	10
Active: last ship within	30	App release · contract deploy	9
Active: or at least 2 ships within	45	Feature · integration · API · SDK release	8
Quiet: no ship for up to	60	GitHub release · contract upgrade · game release	7
Dormancy gap that makes a comeback "Resumed"	≥ 60	Roadmap milestone · community tool	6
Resumed: comeback ship within	14	Docs · demo · design · creative drop · code activity	3
		Community event · founder update	2
		Announcement · other	0

6.2 Build Momentum

Build Momentum (HBM) measures how consistently a project ships meaningful work. Five components, each 0–100, are combined with fixed weights (Figure 4). **Recency** sums the base weights of meaningful events in the last 30 days, each decayed exponentially by age with a 14-day decay constant, and saturates at a total of 24 so that a burst cannot grow without bound. **Consistency** is the share of the last six calendar weeks that contain a meaningful event. **Significance** credits breadth: events are grouped into six dimensions (product, protocol, integration, developer tooling, milestone, community); within a dimension the first event earns 1, the second ½, the third ¼ and further events nothing, community and tooling dimensions are weighted 0.8, and the full score needs four dimensions. **Verification** is the mean evidence weight of the events. **Source diversity** counts independent source classes (GitHub, website, on-chain, launchpad, builder), full at three so breadth is not demanded of every project. Every component is measured inside the six-week window; a project that registers sources but ships nothing scores 0, and years-old evidence cannot prop up a project that has stopped.




Figure 4. HBM component weights, PRD V4 §11.1, as implemented in hbm-v3. The user-facing view never shows an opaque number: each score is accompanied by active weeks, meaningful and verified event counts, dimensions covered, source classes and the top contributing events.

6.3 Discovery Gap and Under the Radar

Market attention is a percentile of three log-scaled figures within the published cohort — 50% market cap or FDV, 30% liquidity, 20% 24-hour volume — with missing figures renormalising the remaining weights rather than counting as zero, which would falsely depress a project that merely lacks coverage. The Discovery Gap is the project's HBM percentile minus that market percentile, and is undefined when no market reading exists. A project appears on the **Under the Radar** surface only when it is Shipping, Active or Resumed, its HBM is at least 45, it has at least two meaningful events in the last 30 days, its market data is fresh, and it is not hidden. Since hbm-v2 (2026-09-04) the percentile cohort is the published catalogue only; the earlier cohort included

thousands of hidden launch shells at HBM 0, which inflated every percentile and populated the surface dishonestly. On honest percentiles the surface is currently near -empty (section 10), and the threshold is a product decision that has not been taken.

6.4 Still Building

Still Building is the flagship label and directly expresses the thesis: verified building continued through a HEY-tracked market drawdown. A project qualifies when it is Shipping, Active or Resumed; has at least two meaningful source-linked events in the last 30 days; has a current market value and a tracked high from *the same* provider (a drawdown is never measured between two providers' arithmetic); its current value is at or below 50% of the highest value HEY itself observed in the last 90 days; and at least one meaningful ship occurred after that high. The tracked high is HEY's own observation, never a canonical all-time high, and the copy says so: "Down 66.9% from the HEY-tracked high, with 2 verified updates since."

6.5 Market bands and market readings

For filters, current market cap or FDV is banded as Micro (< \$100k), Small (< \$1M), Mid (< \$10M), Large, or Unknown. A card shows the freshest reading no older than seven days from the most direct source — DEX Screener, then GeckoTerminal, CoinGecko, Virtuals, then the Stock Token API — and labels it ("via Virtuals" marks the launchpad's curve valuation, not a market print). Anything older is history: the card reads "Unavailable", never \$0 and never a figure from another source. Ordering by market cap happens only when a reader explicitly asks for it; the default card order is by activity, with complete cards (contract, GitHub source and official X present) first — presence counts, values never do.

6.6 Narratives by evidence

Narratives are assigned by deterministic rules over evidence HEY already stores, applied in order of trust, each with a confidence and a plain-language reason: launch origin (Virtuals → AI Agents 0.9; PAIR → StockFi 0.9; memecoin launchers → Meme 0.7), registry categories from CoinGecko (0.8–0.95; tokenized-stock category → StockFi 0.95), repository topics and descriptions (0.7–0.8), and finally the project's own words (0.5–0.7). A primary is assigned only at confidence ≥ 0.6; names alone never assign; a meme is what a project says it is, never what it is called; and a project the rules cannot place stays unknown. Rows written by a project owner or a moderator are never overwritten by a rule. Rules version narratives-v1.

6.7 Versioning

Any change to a weight, threshold or formula bumps the scoring version, which is persisted with every snapshot so historical scores stay explainable. hbm-v2 restricted the percentile cohort to published projects and persisted Under-the-Radar eligibility with the snapshot; hbm-v3 collapsed same-day multi-repository code activity into one event. Nine builder pages in production were affected by the latter on 2026-09-04. Weights have not changed since PRD V4 §11.1.

7. Integrity by design

HEY's credibility rests on rules that are enforced in code and tests rather than in policy. The product rules below are non-negotiable in the repository's engineering contract (CLAUDE.md) and each has a mechanism behind it.

RULE	HOW IT IS ENFORCED
No wallet analytics, holder tracking, PnL, smart-money labels, clustering, wallet profiles or holder velocity.	A schema test that runs without a database fails the build if any column or migration mentions one; adapter schemas do not declare holder fields; fixtures prove they never pass through.

Activity status is based on developer and project activity, never on price.	The activity input type has no market field; the derivation reads ShipEvents only.
Market cap, price, liquidity and volume are context and filter data only.	Market data enters at the Discovery Gap and nowhere earlier; ordering by a market figure requires an explicit reader choice; presence of facts may order, values never do.
Organic ranking never depends on payment, \$HEY holdings, sponsored spend or price appreciation.	packages/scoring cannot read the database; a neutrality test fails the suite if any organic module references the research-request ledger, the priority column, the token config or a payer address.
Still Building is a factual observation, not a buy signal; no copy implies price appreciation.	Copy rules in PRD V4 §3 and §40; the badge text is generated from the observation ("Down X% from the HEY-tracked high, with N verified updates since").
Every meaningful ShipEvent keeps source provenance; conflicts are preserved.	Evidence rows and external ids on every event; source_conflicts records for every demotion; disputed and retracted events keep their record at zero weight.
Self-reported and verified states remain visibly distinct.	Distinct verification states, distinct typography, distinct ownership wording on the project page.
Token identity is (chain_id, contract_address), never a symbol.	Lower-case CHECK constraint on candidate addresses, case-insensitive token addresses, a daily identity audit with zero violations in production (2026-09-04).
Pages never call providers; no LLM in a request path.	A static test fails the build if the web app imports a source adapter; AI is disabled by default and asynchronous where enabled.

7.1 Recurring audits

Two audits run daily in production and store their results as a trend. The **data-integrity audit** checks 13 invariants — duplicate contracts, orphans, primary-token mismatches, demo rows, published pages without evidence or identity, stuck jobs, missing discovery sources, address-named shells, stale sources — and reported 11 OK and two warnings on 2026-09-04 (6,629 legacy hidden shells since retired to zero; one rate-limited source). The **data-quality audit** recomputes 74 coverage, provenance, ship-event, identity, narrative and submission counters per UTC day. A verified nightly backup, an hourly platform-health row, a worker heartbeat and a daily smoke run (29 checks) complete the picture.

7.2 Submission legitimacy


Anyone can submit a project, but a submission must name something HEY can check: a contract with bytecode on Robinhood Chain, a public non-empty GitHub repository, or a website that answers. Each anchor is evaluated as verified, absent or unknown — an unreachable RPC is not evidence that a builder's contract does not exist. Nothing checkable, or only definite negatives, is refused with a plain sentence and never enters the queue. Submitting is tied to a GitHub account in production; persistent per-day caps are kept in the database (3 accepted submissions per account, 10 per hashed client address, 3 attempts per contract; brand-new accounts hold one pending submission at a time; 5 claim starts per account and 20 per address), and a honeypot plus a signed minimum fill time stop naive automation. Anything odd is flagged, never dropped: a name that differs from the on-chain token name, a website host another published project already uses, a repository the submitter neither owns nor contributes to, a new account, evidence that could not be checked. A person decides; reject reasons enter the audit trail; a claimed-but-unapproved project is never public.

7.3 Claim proofs

A builder proves control of a profile in one of four ways, each nonce-bound and audited: GitHub ownership (admin or maintainer of the declared repository or organisation), a DNS TXT record on the official domain, a website challenge (`/ .well-known/hey-research.txt` or a `<meta name="hey-research-verification">` tag), or an on-chain signature — a 10-minute single-use challenge signed by a wallet that HEY’s own evidence ties to the contract (launch-transaction sender, owner(), or creation-transaction sender). Only the verification artefact is stored; no wallet profile is built from the address. A project with a verified owner refuses a second claimant and explains how to dispute.

8. The product

HEY’s interface follows a rule of extreme simplicity and progressive disclosure: a visitor should understand within five seconds that HEY finds active projects, builders and recent ships; a project’s builder story comes before any market figure; and every deeper layer — latest ship, timeline, evidence, market context, sources — is one step further in. Browsing never needs a wallet or an account, every core page works at 375 px without horizontal scrolling, and light and dark themes share one set of components.



Homepage, heyresearch.xyz, captured 2026-09-03 at 2×. Headline, search, the pulse strip (builders, ships in 7 days, active projects, last ingestion time — every value a real reading), activity tabs, fact and sort chips, and Card V7 project cards.


8.1 Surfaces

SURFACE	WHAT IT SHOWS
Home	Headline, search, pulse strip, 14-day ships heartbeat, activity tabs (All · Shipping · Active · Still Building · New), fact and sort chips, twelve cards, a short live feed.
Explore	The same browse surface with search, type, facts, sort, advanced filters behind a secondary control, and pagination (24 per page; the count is the page).
Ships	A card per evidence-backed ship event, newest first, with a “Latest ship” block (type, title, time, verified or self-reported, guarded source link); search, ship-type filter, facts and sort; 48 per page.

Pulse	The strip and heartbeat; Build vs Market (each bubble a project — position describes what has happened, not what will); Top Builders ranked by build activity, not price; most active narratives; activity cooling; public code ranked by observed development days, never raw commit count.
Radar	The eight discovery surfaces: Still Building, Under the Radar, Shipping Now, Most Active, New Builders, Back From Dormancy, Utility, Memes. Honest empty states.
Narratives	A directory of 23 narratives, each page listing primaries with the same controls.
Search	Published projects and, kept visibly apart, launch records — hidden candidates matched by name or ticker, labelled “Launch record — not researched by HEY”, with contract, launchpad and stored market cap only. An exact contract address reaches the whole discovered universe.
Project	Identity, narrative and type, verified-builder mark, status, last ship, market cap, Still Building; then latest ship, builder activity, build timeline, drawdown chart, research notes, sources; Build Momentum with its components in the margin; market context last and labelled.
Badge	An embeddable per-project badge (/badge/<slug>) linking to the evidence-backed profile.

8.2 The card (Card V7)

Every browse card and Explore row shows the same eight facts and nothing else: name, ticker, logo, market cap (freshest source-labelled reading ≤ 7 days, otherwise “Unavailable”), contract address (shortened, copyable, linked to the explorer), activity status, launch origin (“via Pons”, “via Virtuals”, or “Unknown” / “Independent”), and official X account (omitted when HEY cannot vouch for one). Tokenless builders show their last ship and their own site in the two token slots. There is no momentum or gap score on a card, no commit strip, no holders, volume or liquidity: activity status is visually the most important element, and everything deeper lives on the project page. A logo is shown only when the project supplied one; HEY never generates artwork.



Project page, captured 2026-09-03 at 2x. The builder story first — identity, narrative, verified-builder mark, status, last ship and the Still Building observation — then the latest ship with its evidence link. Build Momentum and its five components sit in the margin; market context comes last and is labelled with its source and age.

8.3 Listing controls

Every browse surface carries search, tabs, fact chips, sort and pagination as plain links and GET forms, so a view is a URL, nothing needs JavaScript, and every control keeps every other parameter. One vocabulary — q, tab, has, sort, offset — means the same thing on every page; unknown values fall back to “no filter”, never to

an error. All predicates are applied in SQL on the same joins the card reads, so the count is the page and the empty state is honest.

8.4 Asking about one address: the builder scan

Everything above answers the building question for a project HEY already tracks. Since **2026-09-15** it answers the same question *on demand, for a contract address it has never seen, at /scan*.

The framing is the whole of it. Every other scanner on this chain answers “**will this rug?**” — holders, wallet clusters, honeypot simulation, liquidity locks. None answers “**is anyone building this?**”, which is the only question HEY asks. The two do not overlap at any point, which is what makes the word *complementary* true rather than a courtesy: HEY reads no holders and no wallets here at all.

The gap it closes is measurable. The median time between a launch happening and HEY’s scheduled pipeline seeing it is **122 minutes**; a scan answers now.

What it reads. The chain for bytecode, name, symbol and supply. The decoded chain for the creating call — which carries two accounts a block explorer reports as one: what executed the creation (a factory, for a launchpad token) and the account that sent the transaction, which is what a reader means by “who built it”. The explorer for whether the source is published. The token’s own declared site, checked for whether it *names the contract*, because a site that does not claim the token is not evidence the token is that project’s. And any declared repository, for its last push and whether it is shared or forked.

What the contract answers. For the many addresses that declare no website and no repository, the decoded chain still carries the one fact that bears on the question: which methods the contract answers. A plain token answers `balanceOf`, `transfer`, `approve` and little else; a protocol answers a surface somebody designed and maintains. Neither is judged — a plain token is an ordinary and honest thing to be, and the team may be building at another address entirely — and a caller count is a count of distinct senders, never resolved to an address and never stored.

What it refuses to be. There is no score, no count of passed checks and no colour. A scan box is the surface a reader would most readily mistake for a safety rating, and one number is all that takes. Absences are named as absences: where a check could not run, the line says so, and absent means unknown — never a pass, and never a failure. Two independent ceilings hold it: ten scans an hour per client, and a daily provider budget so a reader pasting addresses can never spend what the scheduled pipeline needs.

A scan is also an intake path. It writes one candidate row through the same identity and qualification gates as any launch the factory sweep finds, so asking about an address never publishes it — and the second person to ask gets the stored answer.

8.5 The daily index, and what it is not allowed to become

Since **2026-09-13** HEY keeps its own daily index of the chain rather than reading a provider in the request path: price, liquidity, trades and volume day by day for every token project (`/project/{slug}/market`), and Robinhood Chain itself day by day on Pulse. **HEY Signal** turns measured changes into a feed with before- and-after figures and the source behind each; the **Builder Radar** ranks builders by verified development, on-chain use and research standing, never by price; weekly reports are archived.

Since **2026-09-14**, and by an explicit decision of the founder rather than by drift, a market page also carries **token distribution**: a bubble map of a single token’s largest balances drawn to scale, with pools, launchpad lockers and burned supply named and kept out of the concentration figure. This is the one and only place HEY holds balance data, and the boundary around it is narrow on purpose. An address there is a point on a chart of one token’s supply. It is never a person, never scored, never ranked across tokens, never followed, and never an input to activity status, Build Momentum, the Discovery Gap or the Builder Radar. A schema test fails the build if a table or column outside three named tables so much as contains the word *holder*, and fails it outright for wallet, PnL, smart-money or copy-trade anywhere.

The concentration figure is stated carefully for a reason. The Nakamoto coefficient HEY publishes counts the addresses holding half the supply *that is not infrastructure* — pools, lockers and burns are excluded from both the count and the supply it is measured against. An earlier edition compared that count against total supply, which read every token as more concentrated than it is.

9. For builders

HEY is designed so that a builder's best strategy is simply to build in public. The path is short and every step is verifiable.

1. **Submit with evidence.** Name a contract on Robinhood Chain, a public repository, or a website that answers. Nothing else is required, and nothing unverifiable helps. A project already on HEY — published or pending — is not duplicated; the form points to the profile and offers the claim.
2. **Claim the profile.** Prove control through GitHub, DNS, a website file or meta tag, or an on-chain signature. Ownership then reads "Verified owner" on the page; until then, "Submitted by a builder; ownership has not been verified."
3. **Let the ships speak.** Register the project's own site, docs, feed and repositories as sources. Releases, changelog entries, contract deployments and aggregated code activity become ShipEvents with provenance. Self-reported updates are welcome but stay visibly self-reported and carry 0.35 of a verified event's weight until a source corroborates them.

What visibility means on HEY. Visibility is earned only by evidence. A project rises on the activity tabs by shipping meaningful, source-backed work; it appears on Still Building when verified building continued through a drawdown; it appears Under the Radar when its build-momentum percentile exceeds its market percentile and the eligibility gate holds. Nothing else moves placement: not the token's price, not holders, not a launchpad boost, not a payment, and not \$HEY. Conversely, a project with no public sources is "Activity unknown", not "Dormant": absence of evidence is not evidence of abandonment, and HEY will never label a project dead, rugged or abandoned from inactivity alone. A memecoin community is welcome without fake utility; a game, tool, art drop or event is real output and is scored as such.

10. Coverage today (2026-09-17)

All figures are production readings taken from heyresearch.xyz on 17 September 2026. Where a figure has moved sharply since an earlier edition the earlier reading is given beside it, because the direction of travel is part of the answer — and on 17 September the direction reversed. The published count fell by 60% in one day. It fell because a defect in the publication gate was found and closed, and §10.5 says exactly what it was; a paper that quietly swapped one number for another would be doing the thing it accuses everyone else of.

4,674

published projects (11,141 on 2026-09-16 before the correction in §10.5; 2,525 on 2026-09-03)

1,480

verified builders (1,443 on 2026-09-16; 215 on 2026-09-03). Unaffected by either correction: every one had evidence

1.35M

token candidates screened (856,702 on 2026-09-03)

15

launchpads read from on-chain factory logs

8,809

ship events on the public feed, every one carrying its source (11,939 earlier on 2026-09-17 before the correction in §10.6; 490 on 2026-09-02)

1,596

ship events in the last seven days (2,541 before §10.6)

567

projects shipping right now

8,660

sources on published projects, re-read on a schedule (27,081 on 2026-09-16, most of them on pages that are no longer published)

2,119

published projects carrying a source-labelled market cap (of 3,102 with a token)

13,650

signals raised in the last seven days

244

projects with publicly verified code activity (261 on 2026-09-16)

23

narratives

The catalogue grew 4.4x in thirteen days, and most of that growth was a mistake. The first edition of this section said that discovery scales with the chain while research does not, and read the gap between them as the honest shape of the figures. Half of that was right. The other half was a publication gate admitting pages it should have refused, and when it was closed on 17 September the published count fell from 11,141 to 4,673. What remains is the catalogue the 3 September edition described, grown at the pace research actually allows. §10.4 says where the thin parts are; §10.5 says what was wrong.

10.1 Launchpad coverage

LAUNCHPAD	HOW IT IS READ	LAUNCHES MEASURED (2026-09-03)
Pons (V1 active, V1 published, V1 legacy, V2, V2 early)	five factory addresses, on-chain logs	~455k candidates
Flap (flap.sh)	factory logs; identity in event data	~90,000 from block 30M, plus 67,579 found so far in the 4.22M–14.72M gap scan
Pools (pools.trade, Uniswap Labs)	factory logs; name/symbol read from contract	~50,700
Virtuals	API, chain filter applied in the request	25,284 candidates (25,253 agents)
Hoodit	factory logs; dormant since block 6.72M	21,443
Bankr	factory logs; symbol in event data	≥ 17,629
Clanker v4	factory logs; full metadata in event data	14,898
PAIR (pair.fund)	factory logs plus public API for identity	1,474
hood.fun (main and community)	factory logs; inline JSON metadata	352 (+ 2 community)
EasyA Kickstart	factory logs	227
Robinpad	nine factories, five creation topics	141
Robinlaunch	fifteen factories (bonding, direct, boost)	~75
PonsPad · hood.dev	public API · Goldsky subgraph	7 · 45

Rejected after checking: launchpad.meme (no public data), NOXA Fun (gone), Openfair and ramenpad (parked pages), TrustSwap (not on 4663), Mint Club (no pools yet). Clanker's API was verified and not integrated because the factory event already carries everything it adds (14,898 launches on chain against 5,954 in the API).

10.2 Market-cap context

On 2026-09-03, 2,204 of 2,214 published token cards carry a market-cap reading, labelled by source: DEX Screener or GeckoTerminal where a token has a DEX pair (256 published tokens at the pre-run baseline), CoinGecko for the 44 published tokens among the 701 coins it catalogues with a Robinhood Chain address, and Virtuals for 2,158 agents priced from the launchpad's own curve or pool valuation, converted with a

same-run VIRTUAL/USD rate. The Virtuals figure is the launchpad's arithmetic rather than a market print; the label says so and it ranks last in precedence. Before the audit only 62 published projects had a snapshot at all (2026-09-02).

10.3 Narratives

Primary narratives after the 2026-09-03 backfill: AI Agents 2,206 · DEX 78 · Launchpad 34 · Infrastructure 16 · Gaming 14 · StockFi 13 · Dev Tools 12 · AI 8, with smaller counts for Lending, Perps, RWA, Payments, Analytics, Meme, Agentic Finance and Data; 97 published projects stay unknown, 45 of them with no classification evidence at all. The caveat is stated in the assignment record: 85% of the published catalogue is Virtuals launches filed under AI Agents by launch origin; product narratives those agents describe are kept as secondaries.

10.4 What is thin, stated plainly

- **Narrative coverage is 49.0%, and the rise from 20.4% is arithmetic, not progress.** 2,289 of 4,673 published projects carry an evidence-based primary narrative, against 96.2% on 2026-09-03. The numerator barely moved between the two editions (2,272 to 2,289); the denominator lost the 7,663 pages §10.5 describes. Assignment still reads a project's own sources one at a time, and a project without a primary narrative is shown without one rather than guessed into a category.
- **Building evidence is concentrated, and thinly.** 244 published projects carry publicly verified code activity, out of 4,673. Much of the catalogue is still launchpad tokens that publish a site and social channels but no repository, and for those "Activity unknown" is the honest state: 67.5% of published projects, down from 87% before the correction, and that too is the denominator moving rather than the evidence.
- **And of the projects that do carry code activity, 41% produced all of it inside a single week.** That figure is not flattering to the index and it is the clearest statement of why the filtering is the product: building something that looks finished has never been cheaper, and still being there in two months has not got any cheaper at all.
- **Market data for curve-stage tokens.** 2,119 of the 3,102 published projects with a token carry a source-labelled market cap. Tokens still on a launchpad curve with no DEX pair have no liquidity or volume print, so their Discovery Gap rests on a renormalised market-cap percentile alone.
- **Claims are rare.** Two projects have a verified owner claim. The mechanism works — four methods, proofs kept, an audit trail — but almost nobody has used it, and a page that says "unclaimed" is saying something true about attention rather than about the project.
- **Identity backlog.** 1.35 million screened candidates sit behind 4,673 published projects. The gap is not waste: a candidate without a readable name or ticker is held rather than published as a shell, and a launch carrying any off-chain identity at all is **466×** more likely to reach a published page (16.1% against 0.035%, measured 2026-09-15). The queue is ordered by that evidence, not by arrival.

10.5 The correction of 17 September

On 17 September an audit of the publication gate found that 7,663 of the 11,507 projects then published carried `https://charms.ai/<slug>` as their website — the mint generator's own token page, not a site anyone built — and that 6,830 of them had qualified on exactly that URL, a verified contract and a launchpad name. Two rules in the gate disagreed about one field. The identity score had always refused a launchpad-generated page; the separate test for whether a page had *anything to show* asked only whether some URL was present. A generator page therefore scored zero for identity and still counted as something to show, and the looser rule decided publication. That is how the catalogue went from 3,664 to 11,492 in four days.

The gate now applies one rule, and the daily sweep re-evaluated every published row under it. 7,957 pages returned to indexed status; 4,673 kept theirs. 786 of the charms-hosted projects stayed published because they carry a repository, a recorded ship or documentation of their own — the URL was never the reason they were there. Verified builders were untouched: 1,479, because every one of them rested on evidence the gate had checked properly.

Two things follow. The 4.4× growth the previous edition reported was mostly not growth, and the coverage shares it reported as falling were falling against a denominator that should not have existed. And the first edition's own claim — that discovery scales with the chain while evidence does not — was true in a way its author did not intend: a discovery pipeline scaled straight past a gate that was supposed to stop it. No demoted page was deleted; each is a launch record HEY holds and will publish the day it can show something, which is the standard the gate was meant to hold all along.

10.6 The second correction of 17 September: ships

The same day, a pass over the ship-event tables found that the “active development” summary — the row that says *N commits in the last 90 days* for a repository HEY reads — had been keyed to the day of the last commit until 15 September, and to its ISO week after. The generator changed; the rows already written did not. A repository that committed on eight different days had eight rows all restating one ninety-day measurement, and every weekly re-read since the change had added a week-keyed twin beside a day-keyed row it was meant to replace. 3,243 of 6,011 such rows were restatements, 3,152 of them on published pages.

They were collapsed on 17 September: one row per repository per week, the latest measurement kept, and no project's last-ship date moved (this was checked before the change: zero). The public ship count fell from 11,939 to 8,809 and the seven-day count from 2,541 to 1,596. Nothing that was a release, an announcement, a documentation change or a feature was touched; every remaining ship still carries its source. The lesson is the one §10.5 already drew, applied to a different table: a change to how evidence is recorded has to be applied to the evidence already recorded, or the old and the new count twice.

11. \$HEY and the research economy

STATUS ON 2026-09-16

\$HEY is **live**. It launched on 2026-09-10 at 0xB33eb16782776b4D738c0Fd643577cb0284Db610 on Robinhood Chain, through **Pons V2**. Research bounties, bonds, Request Research and holder tiers are running; Scout staking and evidence challenges remain designs behind refused flags. The first edition of this paper described all of it as prepared and not live, and §11.3 records where the launch differed from what was then intended.

\$HEY is research working capital for HEY Research Lab: it exists so that useful research work is easier to fund, verify, reward and repeat. It is not the product, and price is not the loop. The governing principle is four words long — *fund research, reward evidence, recycle value, never sell rankings* — and the last clause is a hard invariant enforced in code (section 7).

FIGURE 5 · THE RESEARCH FLYWHEEL




Figure 5. The research flywheel and its treasury loop, from docs/HEY_TOKEN_FLYWHEEL.md and docs/HEY_REQUEST_RESEARCH_V1.md (2026-09-03). Once the token is live the loop is independent of the launch mechanics: the same pool continues, with no bonding curve, migration or second phase. Home, Explore, Search, Ships, Pulse, project pages, Methodology and organic rankings are never token-gated.

11.1 Day-one utility: Request Research

The first and, at launch, only utility is **Request Research**. A reader pays the configured amount of HEY (application default 1,000 HEY, adjustable) to the HEY Treasury on Robinhood Chain; the worker verifies the transfer on chain — chain id 4663, transaction success, sender equals the recorded payer, transfer emitted by the official token contract from canonical configuration, recipient equals the treasury the request was priced against, amount at least the price at request time, one request per transaction hash — and the project moves up the research backlog. The normal research pass then runs and its result is public whatever it says.

Payment buys research priority. It does not buy a positive result, Verified Builder, Build Momentum, Still Building, Under the Radar, Top Builder or homepage placement; the payment writes one column, `research_priority`, which only the research backlog reads. The ledger records where the HEY came from, why, where it went and what work it funded — a record of product actions, not wallet analytics. The free research-request form remains. Neutrality is measured, not asserted: the flywheel metrics report paid requests with negative or inconclusive outcomes and paid projects that did not become Verified Builder.

11.2 What is live

The first edition of this paper described bounties, bonds and holder tiers as designed and flag-gated. They shipped. As of 16 September 2026:

- **Research bounties.** The treasury, or any reader in \$HEY, funds a question; a scout claims it, hands in evidence, a moderator reviews. Half is paid at approval and half after ninety days *if the evidence still stands* — the delay is the point, because evidence that does not survive scrutiny should not have been paid for in full. Rewards are set in dollars by tier (Small \$10, Standard \$25, Deep \$50) and quoted in \$HEY once, from HEY's own newest price reading, when the bounty is posted. Six bounties are open.
- **Bonds.** A claim is backed by a bond of about \$2, quoted at deposit and returned when the work is reviewed. It is kept only by a moderator's explicit decision, with a reason, and then counts as research money. No automatic slashing exists, because a dispute process must come before one.
- **Request Research.** Roughly \$10 to move a project up the queue. It buys a place in the queue and nothing else: not a result, not a status, not a score. Free requests remain for GitHub accounts within a quota.
- **Holder tiers.** Holder (10,000 \$HEY) and Holder+ (100,000) give discounts on research and placements, higher API quotas, research notes 48 hours early, the first 24 hours on every new bounty, and a monthly advisory vote weighted by balance. None of them touches a ranking.

What has *not* shipped, and is still only designed: Scout staking and evidence challenges. Their flags exist and configuration validation refuses `true` for each.

11.3 Launch mechanics, as they happened

\$HEY launched on **2026-09-10** at `0xB33eb16782776b4D738c0Fd643577cb0284Db610` on Robinhood Chain. The first edition said the founder intended to launch on **Pons V1**; it launched on **Pons V2**, and this paper records the difference rather than quietly reconciling it.

A fixed supply of one billion, a bonding curve paired with ETH that graduated at **4.2 ETH** raised, and the full-range Uniswap v4 position **locked permanently**. A 2% creator tax goes to the treasury and is spent in thirds — research bounties, infrastructure and data, platform operations — and **never on a buyback**. Every movement is recorded line by line in a public ledger (`docs/HEY_TREASURY_LEDGER.md`), which is also served at `/docs/treasury-ledger`.

HEY's own token is held to the same rules as every other token in the catalogue. The site names the contract in its footer, so the token reads as *Token verified*; the project page was claimed with the deployer's on-chain signature; and its activity is read from the **public** repository only, with the page saying in as many words that the private repository is not read. Reading the private one was available and would have flattered the figures.

Neutrality, enforced rather than promised. No payment, balance, stake or treasury holding reaches Build Momentum, Verified Builder, activity status, Still Building, the Discovery Gap, Under the Radar or Top Builders. This is asserted by a test in the codebase (`packages/domain/src/hey/neutrality.test.ts`), so a change that broke it would fail the build rather than the trust.

Nothing in this section is an offer, and none of it is investment advice. \$HEY never influences any organic ranking, score or signal. Token markets are volatile and a token may lose all value; readers should make their own assessment and consult their own advisers. Statements about primitives that have not shipped describe designs that may change or not ship.

12. Architecture and operations

HEY is a strict-TypeScript pnpm monorepo: `apps/web` (Next.js App Router, public site and route handlers), `apps/worker` (scheduler, background jobs, data CLI), and `packages` for configuration, database (Drizzle schema and migrations), domain services, source adapters, deterministic scoring and shared UI. **PostgreSQL is both the database and the job-coordination layer**: the worker claims jobs with `FOR UPDATE SKIP LOCKED`, retries with exponential backoff, honours provider cool-offs, and reaps jobs abandoned by a dead worker. There is no Redis, Kafka, Elasticsearch, graph database or self-hosted node, and none may be added without a measured justification. Every external payload is validated at runtime; every adapter has saved fixtures and contract tests; jobs are idempotent and retry-safe; every schema change is a migration (27 applied, 0000–0026, 2026-09-03).

The whole pipeline is scheduled. Every ten minutes: market discovery, due market refreshes, due source reads, rescoring with evidence-first and published-first ordering, the quality gate, and the next site-to-source resolutions (paid research requests first). Hourly: launchpad intake across the factory registry and launchpad APIs, website checks, narrative assignment. Every fifteen minutes: candidate promotion through the identity gate. Daily: ecosystem registries and repository search, a full quality-gate sweep, the builder collapse, logo backfill, the integrity and data-quality audits, retention. Nothing depends on a person typing a command.

Production is a single host (1 vCPU, 2 GB RAM) running Docker Compose behind Caddy with TLS. Deployments go through one script that builds the migrate, worker and web images from a single commit with the build SHA stamped in, runs migrations and stops on failure, restarts worker and web with no side effects on Postgres or Caddy, waits for health checks, runs the 29-check smoke suite, and rolls back automatically to the previous images if anything fails; `--rollback previous` is a seconds-long operation without a build.

Backups are nightly verified custom-format dumps with 7-day retention and a tested restore (30,957 projects across 27 tables restored into a scratch database in 35 seconds during the 2026-09-04 drill); a pre-deploy dump precedes every migration. Observability: `/api/health` reports the running version and `/api/health/deep` the database latency; the worker writes a heartbeat every 60 seconds and carries a container health check; an hourly `platform_health` row summarises queue, sources and resources; container memory limits and log rotation are set; request ids, a nonce-based Content-Security-Policy, HSTS and per-route frame policy are applied at the application layer. Test suite at the 2026-09-03 merge: 908 unit tests in 102 files, 423 integration tests in 32 files on isolated databases, and a Playwright suite of 351 passing journeys on desktop and a 375 px viewport.

13. Roadmap

This roadmap separates what is decided from what is not. Dates are not promised.

13.1 Decided and sequenced

- **\$HEY go-live**, when the founder launches: verify the contract on chain, register its identity through the ordinary candidate-to-project path with no special treatment, switch the token status to live, then enable Request Research after a person has inspected production. A failing verification leaves prelaunch untouched.
- **Research bounties** after at least one review cycle of real Request Research usage; then Scout staking, evidence challenges and API / MCP credits, each behind its own flag and validated before the next opens.
- **Identity backfill** of the launch backlog continues on the scheduled cadence (Flap gap scan below block 15M; Pools and Pons contract reads at 150 per hour).
- **Per-repository code activity** on multi-repository builder pages (already scored once per day under `hbm-v3`) and periodic fork and template re-checks of mapped repositories.

13.2 Open decisions

- **The Under-the-Radar floor.** On honest percentiles no published project reaches HBM 45. Either the threshold is lowered with a documented rationale, or Radar stays empty until builders earn it. The audit deliberately did not change it.
- **Tokenized equities.** Robinhood's 194 stock tokens are real contracts on the chain with no team building behind each. An adapter for the issuer's price API is implemented and off by flag; none of the 153 candidates is published. Whether they get one issuer page, a filter, or nothing is undecided; until then the quality gate must not publish them as projects.
- **Collapsing same-site relaunches.** 1,135 published projects share a website host. A canonical page per site with the other launches as relationships is the candidate shape; it touches the card and identity rules and has not been decided.
- **Infrastructure.** Building images off the production host (or adding swap monitoring) before the next large deploy; per-endpoint rate limits that survive a restart; connection pinning against DNS rebinding for URL fetches. Whether Blockscout can be used at all depends on a public API path being published.
- **Narrative administration** — an admin interface for overriding a rule-assigned narrative (today a SQL change with an audit entry) — and whether the card should visually distinguish a launchpad curve valuation from a market print beyond the current label.

14. Risks and limitations

14.1 Data and method

- **Coverage bias.** HEY can only observe public sources. A team that builds privately, or whose repositories the project's own site does not link, reads "Activity unknown" — not measurable, never a zero. Conversely, the catalogue over-represents launchpads with rich public data (Virtuals is 85% of published projects) and under-represents builders without a token.
- **Thin evidence.** Ship events are concentrated in a few hundred projects (section 10.4). Build Momentum, Discovery Gap and Still Building are only as good as that evidence, and percentiles within a small scored cohort are coarse.
- **Curve valuations.** For tokens still on a launchpad curve, the market-cap context is the launchpad's own arithmetic, labelled as such; liquidity and volume are absent and the market percentile renormalises accordingly.
- **The tracked high is HEY's own.** Still Building compares against the highest value HEY observed in 90 days from one provider, never a canonical all-time high; a token that peaked before HEY tracked it will show a smaller drawdown than its history.
- **Gaming.** Meaningful events require source-linked evidence, commit volume is capped, repeated updates in one dimension diminish, and a self-reported claim carries 0.35 of a verified event's weight — but a determined actor can still generate real-looking releases. Moderation flags, source conflicts and disputes are the response; scores are not truth, they are explained observations.
- **Rule changes.** Thresholds are configuration and will change; versioning keeps history explainable but a change will move individual projects, as hbm-v3 did for nine pages.

14.2 Sources and dependencies

- Every source is free and public, so every source may rate-limit, change shape or disappear without notice. Adapters degrade to the last known value and mark the source stale rather than fail, but coverage can fall. GitHub's unauthenticated limit is 60 requests an hour; a token raises it to 5,000. CoinGecko's free tier rate-limits aggressively. GeckoTerminal lists at most 200 pools. Blockscout returns 403 to programmatic requests and HEY will not impersonate a browser.
- The shared public RPC is undocumented and treated as scarce; factory scans are serialised and budgeted, which bounds how quickly the launch backlog is named.
- Websites supplied by builders are fetched under a strict SSRF guard with DNS resolution on every hop; DNS rebinding between lookup and connect remains possible until connection pinning is added.

14.3 Operations

- A single 2 GB host is over-committed by container limits and still builds images on the box; an out-of-memory event during a build is the remaining infrastructure risk. Rate limits are per web process and reset on restart. Sessions last 30 days and are revocable server-side.
- Recovery heuristics classify provider errors by message text; a new wording could be retried wrongly, which affects retry timing, never data.

14.4 Token

- \$HEY launched on 2026-09-10 through **Pons V2**, not the Pons V1 route the first edition of this paper described as intended. The full-range liquidity position is locked permanently, but a launchpad factory is an upgradeable proxy controlled by a single owner account that also receives launch fees; HEY verifies behaviour, not bytecode, and cannot guarantee a factory's future behaviour. The opening valuation was small by construction; a token may lose all value.

- The economic designs beyond Request Research are unbuilt. Usage may be low, Scout rewards may attract spam, and the treasury policy is an operating policy rather than a contract; each is a stated failure mode the design tries to avoid, not a guarantee.
- Regulatory treatment of tokens differs by jurisdiction and may change; nothing here is legal advice.

15. Glossary and references

Activity status

One of Shipping, Active, Quiet, Dormant, Resumed, Unknown, derived only from meaningful ShipEvents (§6.1).

Build Momentum (HBM)

0–100 measure of how consistently a project ships meaningful work; five weighted components, six-week window, versioned (§6.2).

Discovery Gap (DGS)

HBM percentile minus market-attention percentile within the published cohort. An observation, not a valuation (§6.3).

Under the Radar

The Radar surface for active projects with $HBM \geq 45$, ≥ 2 meaningful events in 30 days and fresh market data.

Still Building

Verified building continued after a HEY-tracked drawdown of at least 50% from the 90-day high observed by HEY (§6.4).

ShipEvent

A typed, dated, sourced record that a project delivered something concrete (§4.3).

Meaningful event

A strong or supporting ShipEvent that is source-linked or better and not disputed, retracted or unapproved.

Verification state

Self-reported, source-linked, publicly verified, admin verified, disputed, retracted (§4.4).

Verified Builder

A published project with verified ships on record; the strongest catalogue state.

Launch record

A token launch HEY has indexed from a factory log or launchpad API but not researched; shown in search only, labelled as such.

Hidden candidate

A discovered project that has not passed the evidence gate; never listed, reachable only by exact contract.

Narrative

A discovery label (one primary, up to four secondaries) assigned from evidence with a confidence; never a score input (§6.6).

Market context

Cached market cap or FDV, liquidity and 24-hour volume with source and age; shown as context, never used for status or momentum.

HEY-tracked high

The highest market value HEY itself observed for a token from one provider in a window; not an all-time high.

Polling tier

HOT, WARM, COOL or COLD; how often a project's market and sources are refreshed, recomputed from its activity (§5).

Request Research

The day-one \$HEY utility: verified payment buys a place in the research queue, never a result (§11.1).

Scout

A researcher who completes verification tasks; staking and rewards are later, flag-gated phases.

Chain 4663

Robinhood Chain's chain id; the only chain HEY reads.

References

1. HEY Research V4 — Final Builder Discovery & Project Showcase PRD (docs/HEY_RESEARCH_FULL_PRD_V4.md), the product and engineering source of truth; §§9–13 (ship events, activity status, Build Momentum, market context, Discovery Gap and Still Building), §15 (narratives), §18 (claims), §22 (cost rules).
2. Engineering and product contract (CLAUDE.md): non-negotiable product rules, UI/UX rules, architecture and public-data-first rules.

3. Scoring configuration and rules, version hbm-v3
(`packages/scoring/src/{config,version,significance,activity,hbm,discovery-gap,math}.ts`).
4. Public methodology page, heyresearch.xyz/methodology, and the token page, heyresearch.xyz/hey (read 2026-09-03).
5. Full platform audit index, 2026-09-04, with its area reports: ingestion pipeline; domain, scoring and database; web API and security; UI/UX and mobile; operations and deployment; data quality and coverage; market coverage sources; submission legitimacy; launchpad coverage; deferred follow-ups (`docs/audit-2026-09-04/`).
6. Public Source Registry (`docs/SOURCE_REGISTRY.md`, verified 2026-09-01 with per-source re-verification through 2026-09-05) and Data sources (`docs/DATA_SOURCES.md`).
7. Narrative assignment, rules version `narratives-v1` (`docs/NARRATIVES.md`); Project cards, Card V7 (`docs/PROJECT_CARD_V7.md`); Search (`docs/SEARCH.md`); Listing controls (`docs/LISTING_CONTROLS.md`).
8. \$HEY token flywheel (`docs/HEY_TOKEN_FLYWHEEL.md`); Request Research V1 (`docs/HEY_REQUEST_RESEARCH_V1.md`); M13 token-economy implementation report, 2026-09-03 (`docs/HEY_M13_TOKEN_ECONOMY_IMPLEMENTATION.md`); Pons V1 pre-launch audit, read live from chain 4663 on 2026-09-02/03 (`docs/HEY_PONS_V1_PRELAUNCH_AUDIT.md`).
9. Public sources read by HEY: Robinhood Chain public RPC (rpc.mainnet.chain.robinhood.com); launch factories of Pons, Pools (pools.trade), Flap, Clanker, Bankr, PAIR (pair.fund), hood.fun, Hoodit, Robinpad, Robinlaunch and EasyA Kickstart on chain 4663; Virtuals, PonsPad, hood.dev (Goldskey) and pair.fund APIs; DEX Screener, GeckoTerminal and CoinGecko public APIs; GitHub public API; Sourcify; Robinhood Stock Token API (`docs.robinhood.com/chain/stock-token-apis`, adapter off by flag); ponsfamily.com and github.com/ponsdotdev/ponsfamily for the Pons V1 contracts.

HEY Research Lab · heyresearch.xyz · Whitepaper v1.3, 17 September 2026. Figures are dated at the point of use, taken from live production readings on 17 September 2026, and will change daily. This document is not investment advice.