# Peptome.

*the operating system*

## for human biology.

### Intelligence layer · Execution layer · Closed loop.

An architectural whitepaper on agent-native scientific infrastructure, the consumer execution surface, and the revenue model that connects them. Settled machine-side in USDC on Solana via HTTP 402 + Solana Pay.

| | |
|---|---|
| VERSION | 3.0 · SOLANA EDITION |
| PUBLISHED | APRIL 19, 2026 |
| CLASSIFICATION | PUBLIC · ARCHITECTURAL · ECONOMIC |
| LAYERS | INTELLIGENCE (CORE) + EXECUTION (AGENT) |
| MACHINE RAIL | HTTP 402 · USDC · SOLANA PAY |
| HUMAN RAIL | FIAT SUBSCRIPTION · USDC-ON-SOLANA OPT-IN |
| TOKEN | $PEPT · SPL · SOLANA |
| AUTHORS | PEPTOME RESEARCH |

IMPRINT · COLOPHON

# On this document.

PURPOSE

This whitepaper describes Peptome's architectural model, evidence rubric, and payment primitive. It is not an implementation specification. It does not describe build order, infrastructure choices, or operational procedures. Those live in internal documents and are deliberately omitted here.

AUDIENCE

Institutional evaluators, agent-framework developers, research operators, biosecurity practitioners, and policy analysts assessing the feasibility and governance of agent-native scientific platforms.

SCOPE OF CLAIMS

Architectural and economic reasoning in this document reflects the authors' position as of April 2026. Specific pricing, chain support, and rubric versions are subject to change; the canonical values at any time are those published in the platform's live documentation.

TYPOGRAPHIC NOTES

Set in Inter Display, JetBrains Mono, and Instrument Serif Italic. Color palette aligned with the Peptome application surface. Rendered in the unit of the research instrument, not the marketing brochure.

CONTENTS

# Contents.

§00 Abstract 004

§01 The Problem 005

§02 Thesis 008

§03 System Architecture · Two Layers 010

§04 The Agent-Native Payment Layer 013

*Solana · USDC · HTTP 402 · $PEPT*

§05 Evidence Model + P-Tier 021

§06 The Execution Layer — Peptome Agent 024

§07 The Closed Loop 028

§08 Interoperability 030

§09 Safety & Governance 032

§10 Economics & Revenue Model 034

*Two rails · Token accrual · Unit economics*

§11 Roadmap 043

§12 Conclusion 044

*A small text for a specific reader.*

§ 00

# Abstract.

*a compressed statement.*

Peptide biology is among the most scientifically consequential and informationally fragmented domains in the life sciences. Over four hundred thousand peptide-related publications sit in PubMed. Thousands of trials are active on ClinicalTrials.gov. Tens of thousands of patents live in the USPTO's queryable surface, and a comparable volume of gray literature — preprints, vendor datasheets, clinician case reports, community protocols — exists outside any canonical index. No surface grades this corpus. No machine-readable consensus exists. And none of it currently connects to the human body that is supposed to act on it.

Peptome is a two-layer operating system that closes that gap, built as a Solana-native protocol. The intelligence layer indexes, grades, and grounds the peptide corpus behind a single queryable surface, exposed to autonomous agents through an HTTP 402 handshake settled in USDC on Solana via Solana Pay transaction requests. The execution layer takes that intelligence and applies it to individual humans: personalized daily protocols, adaptive adjustments based on biomarkers and outcomes, and a continuous feedback channel that converts real-world results back into evidence.

The architectural consequence is a closed loop between research and lived effect. Grades flow from the evidence engine to the protocol surface; outcomes flow from the user back to the evidence engine as a distinct P-tier of post-market, cohort-weighted data. No other peptide platform owns both layers, and no other platform closes the loop.

Coordination of the loop — fee routing, validator staking, evidence-rubric governance, and the Sovereign access tier — is carried by $PEPT, a single-chain Solana SPL token with narrow, functional utility. The token is not the product. The token is the coordination primitive that lets two rails share one evidence graph without a trusted intermediary.

KEYWORDS

TWO-LAYER OS · SOLANA · HTTP 402 + SOLANA PAY · P-TIER · $PEPT TOKEN · CLOSED LOOP

This [illegible] describes the [illegible] architecture, principles, the [illegible] payment primitive (HTTP 402 · USDC · Solana Pay), the evidence rubric (now six tiers including P), the execution layer's product scope, the dual-rail revenue model, the $PEPT token design and value-accrual mechanism, its unit economics, and the safety, governance, and interoperability frameworks that make the platform sustainable at the scale of an emerging agentic web — and, crucially, at the scale of real human use.

§ 01 · THE PROBLEM

# Fragmentation as the default condition.

Peptide therapeutics are the fastest-growing class of regulated pharmaceuticals, with FDA approvals accelerating through the early part of the decade and a pipeline of more than eight hundred candidates in clinical evaluation. The research community around them is correspondingly broad: academic laboratories, contract research organizations, clinical investigators, compounding pharmacies, longevity clinics, and a large informal research market. Each of these constituencies reasons over the same primary corpus — peptide literature, trial data, structural biology, patent claims, and vendor certifications — but accesses that corpus through different, incompatible surfaces.

The surface-level consequence is friction. A clinician searching for human evidence supporting a GLP-1 dual agonist for metabolic-associated steatohepatitis must cross-reference at minimum five data sources, each with a distinct query syntax and none with a consistent evidence-grade output. The researcher must then reconcile their findings across the surfaces themselves before reasoning can begin.

The deeper consequence is structural. Because the corpus is fragmented, no market-clearing mechanism exists for information quality. Vendor claims and peer-reviewed randomized-controlled-trial conclusions are weighted equivalently in a general web search. Preprints and meta-analyses co-mingle. Novel compound claims from compounding pharmacies are indistinguishable — to both a human researcher and an automated crawler — from validated clinical outputs. The evidence-weighting burden falls entirely on the reader.

| PUBMED PEPTIDE REFS | ACTIVE TRIALS | PEPTIDE PATENTS |
|---|---|---|
| 400k+ | 3.2k+ | 28k+ |

SOURCE NOTE · COUNTS ARE ORDER-OF-MAGNITUDE APPROXIMATIONS AS OF Q1 2026.

§ 01.2 · THE AGENTIC INVERSION

# The consumer class is changing.

Simultaneously, a new consumer class is emerging. Autonomous AI agents — operating on behalf of clinicians, researchers, and institutions — are increasingly the first interface to the scientific literature. These agents read, summarize, synthesize, and, increasingly, act.

They do not hold credit cards. They do not authenticate via email. They do not maintain OAuth refresh tokens across orchestrator restarts. They cannot meaningfully agree to a terms-of-service document as a legal matter. And they cannot, at the rates at which they now operate, tolerate the friction of per-query human authentication.

For agent access to scale, a payment-and-authentication primitive is needed that is machine-native by default. The cost of failing to provide one is not a missed opportunity. It is the complete exclusion of a consumer class whose query volume will, within a horizon of years, exceed the human readership of every major scientific database combined.

*Peptome's construction starts from this twin observation: that peptide biology requires a grading layer to be useful, and that its future primary consumer will be agentic, not human.*

— THESIS · § 02

§ 02 · THESIS

# Build the intelligence layer, *not another database.*

Peptome's thesis reduces to four architectural commitments. Each commitment constrains the surface area of the product and simplifies the reasoning about its economics.

01 **Ground every answer in source.**

No synthesized claim is published without explicit citation. Every card, every Copilot response, every Discovery Lab proposal carries a chain of reference back to primary literature or to the flagged absence thereof.

02 **Grade every source.**

No citation is treated equivalently to another. Evidence is ranked by an algorithmic rubric that separates human RCT meta-analyses from rodent case studies from unverified vendor claims. The rubric is public, versioned, and reproducible.

03 **Expose one surface twice.**

The human application and the agent API are two views on the same grounded knowledge graph. An agent paying USDC per call receives the same graded, cited answer a Clinician-tier subscriber receives through the web interface. One system of record. Two consumption modes.

04 **Price access in the unit of the consumer.**

Humans subscribe in dollars; agents pay per call in stablecoins. The payment primitive adapts to the caller. The underlying intelligence does not.

§ 02.2 · COROLLARIES

# What the commitments imply.

These four commitments produce a platform that is simultaneously a reference tool for humans and an infrastructure service for agents, without the tax of maintaining two distinct products.

The first corollary: the platform's data model is the product. No layer of presentation, personalization, or marketing workflow is permitted to mutate the underlying claim record. Grades are written once per rubric version and recomputed deterministically when the rubric evolves.

The second corollary: the agent surface and the human surface cannot diverge in semantics. If a card says a claim carries Grade C evidence, both the web application and the JSON response return Grade C. This is a non-negotiable architectural invariant.

The third corollary: the economic model is a consequence of the architecture, not the other way around. Per-call agent pricing is feasible because the underlying system produces per-call answers with per-call cost structure. It is not a business decision layered atop a subscription-shaped product.

§ 03 · OVERVIEW

# Intelligence above, execution below.

Peptome is built as two interoperable layers sharing one evidence graph. The upper layer — Peptome Core — is a grounded peptide-biology API priced in USDC and served to agents, labs, and AI systems. The lower layer — Peptome Agent — is the consumer surface where a single person stacks peptides, logs outcomes, and follows adapted protocols. Outcomes from below feed the evidence graph above. Intelligence above routes back into personalization below.


*The evidence graph is the seam. Below it, an individual. Above it, the field.*

§ 03 · SYSTEM ARCHITECTURE

# Four subsystems, one substrate.

Four subsystems share one substrate — the evidence graph. The graph is a continuously updated knowledge representation of peptides and their relationships to targets, indications, trials, patents, analogs, vendors, and the graded corpus of claims about them.

### EVIDENCE ENGINE

The retrieval and grounding layer. A hybrid index operates over the corpus — lexical over titles, abstracts, and full-text; dense vector over semantic embeddings of the same. Queries route through both indices in parallel; results are re-ranked by a cross-encoder trained on peptide-specific relevance signals. Every retrieved document carries a grade.

### CARD GENERATOR

Produces deterministic peptide profile documents. Each card specifies molecular weight, sequence, predicted half-life across administration routes, predicted bioavailability, predicted blood-brain-barrier permeability, predicted immunogenicity, and the strongest human clinical evidence available for each indication. Where a prediction cannot be made with high confidence, the field is flagged as absent rather than hallucinated.

§ 03.2 · REASONING & DISCOVERY

# The reasoning surface and the generative surface.

### PROTOCOL COPILOT

A conversational reasoning surface over the evidence graph. It differs from a general-purpose chat interface in three respects. First, every claim it makes carries an inline citation chip linked to the primary source. Second, it refuses to synthesize across unsupported claims, returning an explicit 'no evidence at this grade' when the graph does not support a requested answer. Third, it exposes a live protocol-building surface: conversation builds a structured, exportable protocol document that is itself graded and cited.

### DISCOVERY LAB

The generative subsystem. It proposes novel peptide analogs, de novo binders for specified targets, and cryptic peptide candidates mined from proteome data. Every proposed molecule is scored for predicted affinity, synthesizability, and immunogenicity. Every generation request is filtered through a biosecurity screening layer before execution — see § 07.

§ 03.3 · THE GRAPH AS INVARIANT

# The substrate is the product.

A design choice that is easy to miss: the substrate in Peptome's architecture is not a database schema. It is a knowledge graph with explicit grade, citation, and version metadata attached to every edge. The presentation layers — human web, agent API — are derivations of the graph, not sources of truth about it.

This matters because the platform's single most valuable asset is the grade. A graph where every claim is cited and graded is not a better search engine. It is the specification for the format that every downstream consumer must accept in order to reason correctly over the corpus.

In the long term, the graph is the commercial surface. The application and the API are interfaces onto it. A read-only license to derivative graph slices is conceivable; a future federated graph where academic consortia contribute and consume with signed provenance is an adjacent possibility; a neutral registry for peptide-claim disputes is a natural extension. All of these are functions of the substrate existing with integrity.

*The graph is the commercial surface. The application and the API are interfaces onto it.*

§ 04 · AGENT-NATIVE PAYMENT

# Why the default assumptions break.

Scientific information platforms have historically monetized through one of three rails: institutional subscription, single-user subscription, or pay-per-article. Each was designed for a human reader with a browser, an institutional affiliation, and a billing contact. None scales to the operational reality of the agentic web.

An autonomous research agent in 2026 — a clinician's triage assistant, a biotech company's compound-screening bot, a sovereign wealth fund's due-diligence system — issues queries at rates that make human-style authentication infeasible. The agent cannot reliably hold an OAuth refresh token across orchestrator restarts. It cannot meaningfully agree to terms of service as a legal matter. It cannot present a credit card for a $0.004 query. It cannot, at the workflow scale it now enables, tolerate the friction of a monthly-subscription price structure priced for a human.

Peptome's payment layer is designed around exactly this gap. It replaces account creation with cryptographic settlement, invoicing with on-chain proof-of-payment, and subscription gating with per-call clearing in USDC on Solana — the only chain the protocol settles against.

§ 04.1 · HTTP 402 + SOLANA PAY

# HTTP 402, revived.

*a long-dormant status code, now load-bearing.*

Peptome implements payment-required HTTP using the long-dormant 402 status code (RFC 7231) paired with Solana Pay transaction requests as the on-chain settlement primitive. The semantics are x402-style — request, quote, settle, verify, return — but the wire format is Solana-native end to end.

A compliant server responds to an unfunded request with a 402 response carrying a price quote in USDC, a Solana Pay transaction-request URL, a destination ATA (associated token account), and a challenge memo. The caller fetches the partially-built transaction from the provided URL, signs it with its Solana wallet, submits it to the cluster, and re-issues the original request with the transaction signature attached as a header. The server confirms the signature against the cluster, validates the memo, and returns the response.

For a human this is opaque. For an agent it is the most natural authentication primitive available: a verifiable, cryptographic, account-free signal that the caller has paid for the resource — confirmed in well under a second on Solana mainnet, at fractions of a cent. No session state. No secret management. No vendor-specific SDK.

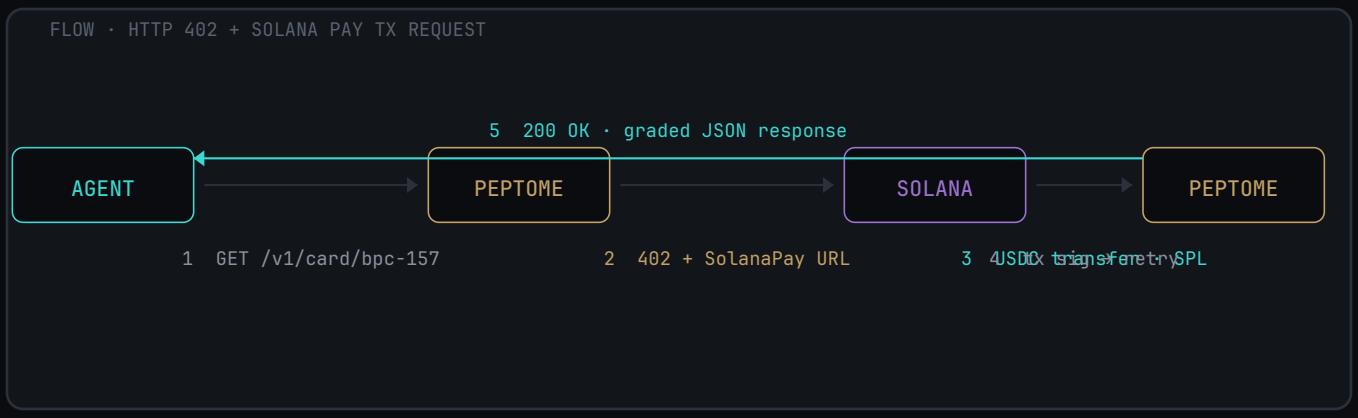

§ 04.2 · WHY USDC ON SOLANA

# Dollarized, fast, agent-compatible.

Peptome settles exclusively in SPL USDC on Solana. USDC is issued natively on Solana by Circle; no bridge, no wrapped asset, no cross-chain trust assumption sits between the caller and the unit of account. The choice is structural, not ideological.

### Dollarized pricing.

A $0.004 query denominated in SOL is not a $0.004 query twelve hours later. Institutional consumers cannot underwrite spot-price volatility in their operating costs. USDC settlement preserves the unit of account the consumer's finance function already uses.

### Sub-second finality, near-zero fees.

Solana confirms transactions in roughly 400ms at base costs measured in hundredths of a cent. For a protocol clearing per-call payments at $0.002–$4.000, the settlement layer must not impose a floor that exceeds the payment itself. Solana is the only major settlement rail where this is unambiguously true today.

### Native Circle issuance + Solana Pay standard.

USDC on Solana is an SPL token with direct Circle mint/redeem — treasury can sweep to Circle rails without a bridge hop. Solana Pay provides a signed transaction-request standard that gives the 402 response a clean, spec-conformant payload the caller's wallet already knows how to consume.

### Agent-compatibility via durable nonces.

Solana's durable-nonce primitive lets an agent pre-sign a payment that clears later, without racing a blockhash window. Combined with priority fees, this gives programmatic callers deterministic inclusion under load — the property an agent loop actually requires to build on a payment rail.

§ 04.3 · SOLANA-NATIVE PRIMITIVES

# One surface, one chain, one wallet.

A query arriving at Peptome's payment surface is priced in USDC. The 402 response carries a Solana Pay transaction-request URL; the caller's wallet fetches it, signs the embedded SPL-transfer instruction, and submits the signed transaction to any Solana RPC. The returned signature is replayed as a header on the original request.

Settlement is verified against a cluster-pinned RPC — Helius and Triton on hot paths, public RPC as fallback — checking that the signature lands, the amount matches the quoted USDC, and the memo contains the request-hash commitment. Confirmed payments unlock the response within one slot.

The architectural consequence is that Peptome is a Solana-native product. No bridge, no multi-chain router, no chain-selection UI. One chain, one asset, one wallet primitive. The simplification is not cosmetic — it removes a class of failure modes (bridge latency, wrapped-asset accounting, per-chain indexer drift) that the caller would otherwise have to tolerate.

SOLANA PRIMITIVES USED BY THE PROTOCOL

| | | | |
|---|---|---|---|
| SOLANA PAY | spec | tx-request | signed transfer payload in 402 |
| SPL USDC | asset | Circle-native | no bridge, no wrapper, no peg drift |
| DURABLE NONCE | primi. | pre-sign | race-free agent payment loop |
| PRIORITY FEE | primi. | inclusion | deterministic landing under load |
| HELIUS / TRITON | rpc | hot-path | low-latency settlement verify |

§ 04.4 · PRICING STRUCTURE

# Endpoint-level, cost-reflective.

Endpoint-level pricing reflects the underlying compute, storage, and retrieval cost structure. Indicative ranges at launch follow; pricing is subject to dynamic adjustment based on demand and compute costs. The 402 response always carries the current quote, and historical quotes are cryptographically logged for replay-resistance.

| ENDPOINT | TYPICAL PRICE | COST DRIVER |
|---|---|---|
| `/v1/search · lightweight` | $0.002 | dense vector lookup |
| `/v1/search · graded + citations` | $0.008 | hybrid index + cross-encoder |
| `/v1/card/{peptide}` | $0.020 | generation + freshness check |
| `/v1/copilot · per reasoning turn` | $0.050 | multi-retrieval + synthesis |
| `/v1/lab/structure` | $0.500 | GPU structure prediction |
| `/v1/lab/de-novo-binder` | $4.000 | diffusion + MPNN + filter pass |

The structural claim: per-call pricing in stablecoins enables Peptome's consumer base to scale with the agentic economy itself. A single research organization deploying a compound-screening agent at one query per second generates a revenue stream commensurate with a 150-user Pro cohort — without requiring any of the marketing, onboarding, or retention overhead of a human subscriber base.

The payment layer itself is not a token sale. It is not a speculative asset. It is a protocol for clearing per-call payments in the unit of account (USD) on a single settlement layer (Solana) with sub-second finality. $PEPT — described on the following pages — is a separate coordination layer, orthogonal to access.

§ 04.5 · $PEPT TOKEN DESIGN

# $PEPT — SPL token on Solana.

*Coordination, not access.*

$PEPT is an SPL token issued on Solana. It does not gate access — access is always priced in USDC and available to any agent that can sign a Solana Pay transaction. $PEPT exists to coordinate four functions the protocol cannot outsource: stake-based validation, rubric governance, cohort-privacy consent, and an optional subscription tier.

The design is deliberately narrow. A token that tries to be everything becomes a meme. A token bound to specific protocol work becomes infrastructure.

TOKEN FACTS

| SYMBOL | CHAIN | SUPPLY | DECIMALS |
|---|---|---|---|
| $PEPT | Solana (SPL) | 1,000,000,000 fixed | 6 |

**VALIDATOR STAKE**

Operators running Peptome evidence extractors, P-tier contributors, and third-party clients of the outcome-submission endpoint post $PEPT as skin-in-the-game. Stake is slashed for malformed submissions, adversarial outcomes, or biosecurity violations. This is the gate for machine-submitted evidence entering the graph.

**RUBRIC GOVERNANCE**

$PEPT holders vote on evolution of the evidence-grading rubric, P-tier admission thresholds, and the list of approved outcome schemas. The rubric is the protocol's most sensitive editorial surface; governance is the mechanism by which it survives a change in operators.

**STAKING DISCOUNT**

Agents and developers that stake $PEPT against their API key receive a calibrated discount on per-call fees. Discount scales with stake duration, not magnitude — rewarding long-horizon integrators over passing speculators.

**SOVEREIGN TIER**

The Sovereign subscription (see §10) is priced in $PEPT. This is the only tier where token holding is required; all other access remains USDC-priced and token-free.

§ 04.6 · VALUE ACCRUAL & LAUNCH

# Fees in, buybacks out.

*The flywheel is real revenue.*

Token value accrues from protocol revenue. The mechanism is mechanical, on-chain, and auditable. It does not depend on speculative demand, listing events, or narrative cycles.

ACCRUAL FLOW

| | | |
|---|---|---|
| 01 | PROTOCOL FEES | 15% of per-call USDC routed to the accrual program |
| 02 | PROGRAMMATIC BUYBACK | on-chain swap USDC → $PEPT via Jupiter aggregator |
| 03 | DISTRIBUTION | 60% validator staking rewards, 30% treasury, 10% burn |

INITIAL SUPPLY DISTRIBUTION

| | | |
|---|---|---|
| Public liquidity / fair launch | 30% | Meteora DLMM or Jupiter LBP |
| Protocol treasury | 25% | 4-year linear vesting |
| Ecosystem + validator rewards | 20% | streamed over 48 months |
| Team | 15% | 1-year cliff, 3-year vest |
| Strategic contributors | 10% | 6-month cliff, 2-year vest |

Launch venue is Meteora DLMM or Jupiter LBP — not a memecoin launchpad. The signal must match the product: $PEPT is infrastructure for a medical intelligence protocol, launched with disclosure, price discovery, and a lock-release schedule, not a fair-launch pump. Any listing on a memecoin venue would be inconsistent with the clinical posture Peptome requires.

§ 04.7 · LAUNCH SEQUENCE & RISK

# Ship the protocol first.

*Launch the token to pull demand through it.*

$PEPT does not launch into a vacuum. The token-generation event happens after the payment rail is live, paying customers exist, and the validator network is operating. The order matters — a token launched ahead of working infrastructure is a fundraising instrument; a token launched into working infrastructure is a coordination instrument.

LAUNCH PRECONDITIONS

| | | |
|---|---|---|
| T-90 | Payment rail live | USDC-on-Solana settlement in production for 90 days |
| T-60 | Paying revenue | $50K+ MRR across Rail 1 (per-call) + Rail 2 (subs) |
| T-30 | Validator beta | 20+ external operators staking testnet $PEPT |
| T-14 | Audit complete | two independent SPL + program audits published |
| T-0 | TGE | Meteora DLMM live, public-tier circulating, vesting on-chain |

STRUCTURAL RISKS — DISCLOSED

**Regulatory.**
Token utility is bound to operator-staking and rubric governance, not to revenue claims or profit expectations. Issuance entity, jurisdiction, and offering structure are determined with counsel; non-US persons only at TGE pending regulatory clarity in the operator's home jurisdiction.

**Concentration.**
Validator-stake concentration is the failure mode that converts a decentralized rubric into a captured one. Stake-weight caps, delegation-discount curves, and quadratic governance weighting are all active design constraints, not afterthoughts.

**Decoupling.**
Token price will at times decouple from protocol revenue. The accrual mechanism is mechanical; the secondary market is not. Protocol design must remain independent of token-price feedback to avoid degenerating into a ponzi-shaped incentive surface.

§ 05 · EVIDENCE MODEL

# Every answer cites. Every citation grades.

The grading rubric is algorithmic and versioned. It is the platform's most editorially consequential primitive, and the single surface on which Peptome's reputation is staked.

RUBRIC · v1.2 (CURRENT)

| | |
|---|---|
| A+ | Meta-analyses of human RCTs, n > 3, low heterogeneity |
| A | Single-arm human trials or peer-reviewed human RCTs |
| B | Open-label human data; primate peer-reviewed work |
| C | Rodent or in-vitro work with peer review |
| D | Anecdote, preprint, case report, or vendor claim unverified |

The rubric is versioned. Evidence graded under v1.2 is re-graded automatically when v1.3 ships. The governance of rubric versions is public: each version is specified in a signed artifact, and community input on rubric evolution is collected through structured channels.

§ 05.2 · ABSENCE · TEMPORAL HYGIENE

# Absence is a first-class signal.

Peptome's most distinctive evidence commitment is that the absence of high-grade evidence is itself a signal, surfaced prominently rather than hidden.

A user asking, 'what human RCT evidence supports BPC-157 for Achilles tendon repair?' receives an answer foregrounding the explicit fact that no such evidence exists, with the closest adjacent evidence — animal studies, case reports, compounding pharmacy registries — surfaced at their actual grades and clearly labeled as such.

The Copilot never synthesizes upward beyond the graph's support. If the strongest human evidence for a claim is Grade C, the Copilot returns a Grade-C-qualified answer. It does not assemble Grade C data into a Grade B conclusion. This is a behavioral invariant.

## Temporal hygiene.

The corpus is continuously refreshed. Publications discovered after a card's last update are automatically re-evaluated against the card's claims. Claims refuted by later evidence are explicitly flagged; claims reinforced are noted in the card's update history. Every card carries a freshness timestamp and a changelog. Staleness is a surfaced metric, not a silent failure mode.

*The Copilot never synthesizes upward beyond the graph's support.*

§ 05.3 · PLATFORM OUTCOME · P-TIER

# A sixth tier, earned on the platform itself.

The standard hierarchy — A1 systematic review, A2 RCT, B observational, C case series, D preclinical — describes evidence produced off-platform. When the execution layer accumulates de-identified outcomes from stacks that ran for weeks under known dose, route, and timing, a new class of evidence emerges: grounded in real use, structured enough to compare, noisy enough to require care. We admit it under a named, subordinate tier.

**P · PLATFORM OUTCOME**

RANK · BELOW CLINICAL · ABOVE ANECDOTE

| | |
|---|---|
| ADMISSION | Cohort size ≥ 50 completed protocols with the same active ingredient, dose range, and indication; outcome period ≥ 21 days; ≥ 2 outcome data points per user. |
| CONFIDENCE SCORE | Each P-tier claim carries a confidence in [0,1] derived from cohort size, variance, and drop-off rate. Confidence < 0.4 is rendered as 'signal only, not evidence'. |
| DISPLAY | In Copilot and Agent surfaces, P-tier claims render with a violet outline chip labelled 'P · platform outcome · n=N' linked to the de-identified cohort view. |
| OVERRIDE | P-tier is never ranked above any A-tier or B-tier claim. On contradiction, clinical evidence prevails; the P-tier claim is marked 'conflicts with [citation]' and demoted. |
| PRIVACY | Cohort membership is opt-in at stack creation. Raw outcomes never leave the user's record; only aggregate statistics feed the graph. |

P-tier is the closed loop made visible. It is also the platform's most defensible asset. Every additional cohort is evidence no competitor without the execution layer can generate.

§ 06

# The Execution Layer.

*the side that touches a person.*

Peptome Agent is the consumer surface of the platform. It is what an individual opens on a Tuesday morning to decide, with grounded help, what to inject, when, why, and what to watch for. It is also where outcomes are captured cleanly enough to feed back into the evidence graph.

The Agent is not a separate product. It is the same intelligence — Card Generator, Protocol Copilot, evidence rubric — wrapped in an interface that is short, clear, and reluctant to recommend anything it cannot cite. Where Core sells answers to other software, Agent sells decisions to a person.

The five-screen surface that follows was sketched by partners testing the thesis directly: a single mobile app where Home is a daily protocol, Agent is the conversation, Stack is the active regimen, Outcomes is the log, and Biology is the personal evidence layer that grows with every entry.

§ 06.1 · SURFACE MAP

# Five screens. One operating loop.

Each screen is a function in the closed loop. The order is intentional: decide, ask, commit, observe, learn. Nothing in the surface is decorative.

| | |
|---|---|
| HOME | Today's protocol — what to take, dose, route, timing windows. Adherence streak. Next outcome check. |
| AGENT | Conversational copilot grounded in the user's stack and the evidence graph. Every answer cites. Every refusal explains. |
| STACK | Active peptides and supporting compounds. Interactions checked against vendor verification and the evidence graph. |
| OUTCOMES | Structured logs: sleep, energy, libido, mood, biomarkers, side effects. Each entry timestamped and linked to dose. |
| BIOLOGY | Personal evidence layer — bloodwork, hormone trends, body composition, training load. The substrate the agent reasons over. |

§ 06.2 · ADAPTATION ENGINE

# The protocol moves with the person.

Static dosing schedules are the single largest source of failed peptide regimens. Body composition shifts, sleep degrades, side effects appear, biomarkers move. The Agent rewrites the protocol on a defined cadence, with explicit rules and a full audit trail.

**INPUTS**

Outcome logs (last 14 days) · adherence rate · biomarker deltas · self-report sliders · vendor lot changes · sleep & training load if connected.

**RULES**

Each protocol carries adaptation rules authored against the evidence graph. e.g.: 'if morning energy < 4/10 for 5 of 7 days, propose 20% dose reduction; cite [B-tier observational n=N].'

**OUTPUTS**

Proposed change · evidence chip · expected effect window · explicit decline path. The user accepts, modifies, or rejects. Nothing changes silently.

The adaptation engine is also the safety surface. A rule that proposes stopping a peptide entirely fires before any rule that proposes scaling. Stop-conditions are non-negotiable and authored at the protocol level, not the user level.

§ 06.3 · GUARDRAILS

# What the Agent will not do, and why.

NO PRESCRIBING

The Agent does not prescribe. It surfaces evidence, generates protocol drafts, and routes to clinical review where required. Any output presented as 'recommended' is gated behind a labelled clinical pathway.

NO UNGRADED CLAIMS

Every claim in the Agent surface carries a tier chip. If no graded source exists, the Agent says so and offers the closest analog evidence with a labelled inference distance.

NO SILENT CHANGES

Dose, route, timing, and stack composition never change without an explicit user accept. The diff is rendered. The reasoning is rendered. The decline path is one tap.

NO SCHEDULED-SUBSTANCE WORKFLOWS

Anything DEA-scheduled, GLP-1 included where state law restricts, or otherwise legally constrained in the user's jurisdiction is hidden behind a region-aware gate. Default is hidden.

NO VENDOR ENDORSEMENT WITHOUT VERIFICATION

Vendor links surface only when the vendor passes the verification regime in § 09: third-party CoA, batch traceability, sterility data on file. Unverified vendors are listed with a labelled warning.

NO COHORT WITHOUT CONSENT

Inclusion in the P-tier evidence cohort is opt-in per stack, revocable, and never bundled with subscription. Raw outcomes are never sold.

§ 07

# The Closed Loop.

*the moat is a circuit, not a wall.*

A research API on its own is a content product. A consumer agent on its own is a wellness app. Together — and only together — they form a circuit in which every individual outcome makes the underlying intelligence more specific, and every advance in the intelligence makes the next individual protocol better.

§ 07.2 · OUTCOME EVENT SCHEMA

# Structure beats anecdote.

Every outcome captured in the Agent surface is normalized into a single event schema before it is admitted to the cohort layer. The schema is narrow on purpose. Free-text notes are kept beside it, never inside it.

```
OUTCOME_EVENT · v1

"event_id"          : "oe_01HZX..."                    uuid · client-generated
"user_pseudonym"    : "u_a9f3..."                      rotated weekly · never reversible
"protocol_id"       : "prot_growth_v3"                 links to evidence-graphed protocol
"compound"          : "BPC-157 / 250mcg"               from verified vendor catalog
"route"             : "subq_lower_abdomen"             controlled vocabulary
"administered_at"   : "2026-04-19T07:14Z"              ISO 8601 · device clock
"day_of_protocol"   : 17                               integer · enforced ≥ 1
"outcome_metric"    : "morning_energy"                 from rubric · 24 metrics today
"value"             : 7                                0–10 scale · or numeric biomarker
"side_effect"       : null | "injection_site_redness"  controlled list
"adherence_window"  : "on_time"                        on_time · early · late · missed
"consent_cohort"    : true                             if false · event stored locally only
```

The schema is what makes P-tier admission possible. It is also what lets the Agent be honest about what it does not know — events that arrive outside the schema do not feed the cohort layer at all.

§ 08 · INTEROPERABILITY

# Three surfaces, one handshake.

The external surface is deliberately thin. Three endpoints compose the platform's entire public API; all three are HTTP-402-priced, settled in USDC on Solana, and return structured, graded JSON.

| | |
|---|---|
| `/v1/search` | Hybrid retrieval with grade filters and source typing. |
| `/v1/card/{name}` | Deterministic peptide profile; every field cited or flagged. |
| `/v1/copilot` | Streaming conversational reasoning with inline citations. |

Downstream integration is deliberately thin. An agent framework with tool-calling semantics adds Peptome via a single wallet-signed handshake. There is no SDK lock-in, no proprietary framework, no required middleware. The protocol is HTTP; the authentication is payment; the response is JSON.

Compatibility targets at launch include the Anthropic tool-use specification, the OpenAI Responses format, and LangChain/LangGraph tool protocols. Any framework that supports HTTP calls and can sign a Solana transaction — natively or via an embedded wallet provider — can call Peptome without a framework-specific adapter.

§ 08.2 · RESPONSE SCHEMA

# Grade as a required field.

Every JSON response carries grade metadata in a required top-level field, not as an afterthought embedded in prose. This makes the grade programmatically filterable: an agent can be instructed to accept only Grade-B-or-higher responses, and the filtering is enforceable without parsing natural language.

Similarly, every citation carries source metadata — type, date, author, DOI or equivalent identifier — as structured fields. No agent consuming Peptome needs to screen-scrape citations out of answer text.

The effect is that Peptome is not merely consumable by agents; it is composable with them. An agent can chain Peptome's /v1/search with its own reasoning, filtering on grade at each step, and produce a downstream answer whose provenance is fully machine-auditable back to the primary literature.

§ 09 · SAFETY & GOVERNANCE

# Three policies, non-negotiable.

### DUAL-USE SCREENING

Discovery Lab receives requests to generate peptides with specified biological properties. A subset of such requests may target proscribed protein classes — toxins, select agents, or sequences implicated in known biosecurity concerns. Peptome maintains a target-class filter derived from the Australia Group control list and the United States Centers for Disease Control and Prevention's Select Agents and Toxins list, with maintenance cadence tracking both authoritative sources. Submissions matching the filter are rejected. The submitting operator is notified. Repeated attempts by the same operator trigger escalation to the International Biosecurity and Biosafety Initiative for Science clearinghouse, with identifying metadata attached.

### MEDICAL DISCLAIMER

Peptome publishes research-grade intelligence. It does not provide medical advice. All Clinician-tier output is reviewed by a licensed medical advisor before release. The human-facing application surfaces a persistent disclaimer; the agent-facing API includes a disclaimer field in every response. Where a regulatory authority has specifically restricted a peptide — compounding status, DEA scheduling, FDA labeling — that status is surfaced in the card alongside the evidence, not buried.

§ 09.2 · VENDOR VERIFICATION

# Verification as signal, never as assumption.

Peptome's Certificate of Analysis interpretation surface is a signal, not a certification. Where verified batch data exists through BatchTrust — Peptome's sibling product focused on supply-chain provenance — a verification badge is surfaced on the relevant vendor and peptide pages. Where verification does not exist, the absence is surfaced explicitly; it is never silently assumed, and it never becomes a default positive.

Peptome and BatchTrust are separately branded, separately operated, and separately governed. A user of Peptome is not a user of BatchTrust by default. Cross-linking is contextual, opt-in, and reversible. The platforms share operators, aesthetic lineage, and an underlying commitment to grounded claims, but they do not share a brand surface.

The rationale is straightforward: the trust properties of a COA registry are distinct from the trust properties of an evidence-graded intelligence layer. Conflating them degrades both. Peptome grades evidence; BatchTrust attests to supply. The separation is an architectural choice, not a historical accident.

§ 10

# Economics & Revenue.

*two rails, one loop, compounding.*

The two-layer architecture collapses into a two-rail revenue model. Rail 1 is per-request USDC settlement on Solana via HTTP 402 + Solana Pay, charged to agents, labs, and AI systems calling Peptome Core. Rail 2 is a human subscription, charged to the individual using the Peptome Agent. Both rails are priced independently. Both rails compound each other.

Neither rail is speculative. Rail 1 is live today for any Claude-/OpenAI-shaped client calling an HTTP endpoint that returns a 402 and can sign a Solana transaction. Rail 2 is a conventional SaaS with conventional plumbing — Stripe for fiat, optional USDC-on-Solana for users who prefer it. The combination is unusual. The components are not.

This section works the model end to end: the two rails, the tier matrix, unit economics per surface, LTV/CAC/payback, and a deliberately labelled ARR projection with every assumption called out. The chain question is treated last — it is Solana, and the trade-off is explicit.

§ 10.1 · RAIL 01 · MACHINE

# Per-request USDC. One unit, one answer.

Rail 1 is the agent-native rail specified in § 04. Every Peptome Core endpoint returns 402 until paid. Prices are denominated in USDC and settled on Solana via Solana Pay transaction requests. Defaults below are assumption-labeled; they are the anchor, not a commitment.

| ENDPOINT | PRICE (USDC) | TIER | NOTES |
|---|---|---|---|
| /cards/:peptide | $0.08 | ENTRY | deterministic profile card · cached |
| /copilot/query | $0.35 | ANCHOR | grounded reasoning · multi-citation |
| /generate/analogs | $2.50 | DEEP | novel analog proposal · scored |
| /evidence/search | $0.12 | UTILITY | graded retrieval · sub-second |
| /protocols/draft | $1.20 | ANCHOR | complete protocol w/ adaptation rules |
| /vendor/verify/:lot | $0.05 | UTILITY | CoA + sterility verification |
| /cohort/p-tier/:query | $0.40 | PROPRIETARY | P-tier platform-outcome query |

ASSUMPTION · gross take-rate after Solana network fees and RPC/facilitator share: ~98% (Solana base fees + priority fees + Solana Pay relay are sub-cent at protocol's average request size).

ASSUMPTION · effective marginal compute cost per /copilot/query at steady state: $0.09 (Sonnet 4.6 default) · $0.21 (Opus 4.7 toggled).

ASSUMPTION · cache hit ratio on /cards and /evidence endpoints projected to 60–75% at scale; both are priced above marginal cost even at 0%.

§ 10.2 · RAIL 02 · HUMAN

# Subscription, priced honestly.

Rail 2 is Peptome Agent's consumer surface. A person pays a monthly or annual fee for the Agent, the Stack, the adaptation engine, and the vendor-verified supply graph. Fiat is the default because it converts. USDC-on-Solana is an opt-in for users who prefer self-custody.

### WHY SUBSCRIPTION, NOT TRANSACTIONAL

Consumer peptide users are not paying per question — they are paying to have a system. Per-request pricing at human scale kills retention. Subscription anchors the loop.

### WHY FIAT FIRST

Stripe default. 2.9% + $0.30 take. Apple/Google in-app purchase only if required by policy (with price raised 30% to net back). Crypto-only subscription excludes ~95% of the addressable market and is not worth that cost today.

### WHY USDC-ON-SOLANA OPT-IN

A meaningful segment of the target user base holds USDC. Accepting it on Solana costs hundredths of a cent to settle, aligns with Rail 1 chain choice, and creates a small additional retention signal — crypto payers churn less on SaaS than fiat payers in multiple industry studies.

### WHAT IS NEVER SUBSCRIPTION-GATED

Emergency safety rules. Stop-conditions. Poison-control-tier warnings. Vendor verification red flags. These are platform obligations, not product features.

§ 10.3 · TIER MATRIX

# Four tiers, two upgrade moments.

| | FREE $0 | STANDARD $29 / mo | PERFORMANCE $49 / mo | CLINICAL $199 / mo |
|---|---|---|---|---|
| PROTOCOL COPILOT | 3 queries/day | Unlimited | Unlimited | Unlimited + priority |
| STACK TRACKER | 1 stack | 3 stacks | Unlimited | Unlimited |
| OUTCOME LOGGING | 7-day window | Full history | Full + biomarkers | Full + lab uploads |
| ADAPTATION ENGINE | — | Weekly rewrite | Daily rewrite | Daily + clinician review |
| VENDOR VERIFICATION | Read-only | Included | Included + alerts | Included + direct escalation |
| P-TIER COHORT VIEW | — | Read-only | Read + contribute | Read + contribute + export |
| BLOODWORK INTEGRATION | — | Manual entry | Lab partner auto | Lab partner + interpretation |
| CLINICAL ROUTING | — | — | — | Tele-consult included 1x/mo |
| API CREDITS (RAIL 1) | — | $5/mo included | $15/mo included | $60/mo included |

| | |
|---|---|
| UPGRADE MOMENT 01 · FREE → STANDARD | First unanswered Copilot query · day 3–5 |
| UPGRADE MOMENT 02 · STANDARD → PERFORMANCE | First biomarker upload or second stack · day 25–40 |

§ 10.4 · UNIT ECONOMICS · PER SURFACE

# Margins are a model decision.

Gross margin is routed per surface. Every agent session runs on Sonnet 4.6 by default; deep-tier paths explicitly toggle Opus 4.7. Numbers below are steady-state estimates with caching in place.

| RAIL 01 · /copilot/query | |
|---|---|
| PRICE | $0.35 USDC |
| MODEL COST (SONNET) | $0.09 |
| RETRIEVAL + GRAPH | $0.02 |
| PAYMENT RAIL (SOLANA) | $0.001 |
| GROSS MARGIN | $0.239 · 68% |
| AT 1M CALLS/MO | $239k GM |

| RAIL 02 · STANDARD SUB | |
|---|---|
| PRICE | $29 / mo |
| STRIPE FEE | $1.14 (2.9% + 30¢) |
| MODEL + INFRA ALLOCATED | $4.20 |
| SUPPORT + OPS ALLOCATED | $1.80 |
| GROSS MARGIN | $21.86 · 75% |
| AT 10K SUBS | $218k GM / mo |

The deep-tier endpoints (/generate/analogs, Opus-toggled /copilot/query, clinical-tier subscription) carry lower percentage margin but higher absolute gross contribution per call. They also generate the highest-value P-tier evidence. The mix is the model, not any single line.

§ 10.5 · RETENTION DYNAMICS

# The loop is the retention mechanism.

Health-related subscriptions have punishing default retention. Peptome inverts that because the product gets measurably better per user over time. Their own outcome history — bloodwork, sleep, composition, mood — is locked inside the Agent. Leaving means leaving a personal evidence file that took months to build.

| BLENDED ARPU | GROSS MARGIN | LOGO CHURN |
|---|---|---|
| $38 / mo | 74% | 4.5% / mo |
| std + perf + clinical mix · assumption | weighted across all tiers | target · similar health-tracker comps are 6–9% |

| GROSS LTV | TARGET CAC | PAYBACK |
|---|---|---|
| $625 | $90 | 3.2 months |
| = ARPU × GM / churn | blended paid + organic · perf/clinical higher | = CAC / (ARPU × GM) |

ASSUMPTION · churn benchmarked against Whoop, Oura, and Levels (health subscriptions that accumulate personal data). The P-tier and outcome lock-in should push Peptome below the low end; model is deliberately conservative at 4.5%.

ASSUMPTION · CAC blend assumes heavy organic from the research surface (Copilot citations in AI chat outputs, X/Twitter practitioner community, direct referrals). Paid-only CAC at current prices would land closer to $140.

§ 10.6 · THREE-YEAR ARR MODEL

# A model, not a forecast.

The table below is a deterministic roll-up of the unit economics. Every input is a modelling assumption stated in § 10.5 and the pricing rails in § 10.1–10.2. Treat it as the anchor for planning, not a statement of expected results. Anything labelled ASSUMPTION below is a lever, not a fact.

| INPUT | YEAR 1 | YEAR 2 | YEAR 3 |
|---|---|---|---|
| Rail 2 paying subscribers | 2,500 | 18,000 | 65,000 |
| Blended ARPU (mo) | $32 | $38 | $42 |
| Rail 2 ARR | $960k | $8.2M | $32.8M |
| Rail 1 paying orgs / agents | 40 | 300 | 1,100 |
| Rail 1 average spend (mo) | $380 | $1,400 | $3,200 |
| Rail 1 ARR | $180k | $5.0M | $42.2M |
| **Total ARR** | $1.14M | $13.2M | $75.0M |
| **Blended gross margin** | 70% | 73% | 76% |
| **Contribution margin** | $800k | $9.6M | $57.0M |

ASSUMPTION · Y1 subscriber count assumes product-led growth only, no paid acquisition above experiment budget.

ASSUMPTION · Rail 1 Y2 growth is driven by a small number of AI-agent integrators (Claude Code, Anthropic-partnered clinical AIs, third-party peptide-protocol builders) rather than long-tail signups.

ASSUMPTION · no governance token, no TVL, no speculative revenue. All ARR is either fiat subscription or USDC-denominated API revenue.

ASSUMPTION · tax and compliance overhead (SOC 2 light, HIPAA adjacency, KYC on clinical tier) is absorbed in contribution margin starting Q3 Y1.

§ 10.7 · CHAIN · SETTLEMENT

# One chain. One asset. One wallet primitive.

Peptome settles on Solana. Only Solana. This is a deliberate simplification — the earlier draft of this paper proposed an EVM-multichain surface routed across Base, Arbitrum, Ethereum, Optimism, and Polygon. The multi-chain design was elegant on paper and operationally expensive in practice. Every chain is a parallel indexer, a parallel audit surface, and a fresh customer-education problem. Consolidation is the upgrade.

| SURFACE | CHAIN / RAIL | WHY |
|---|---|---|
| RAIL 01 · API | SOLANA | sub-second finality · sub-cent fees · native SPL USDC · Solana Pay standard · durable nonces for agent loops |
| RAIL 02 · SUB | STRIPE | default · fiat · 95%+ conversion · required for Apple/Google distribution |
| RAIL 02 · SUB | USDC · SOL | opt-in for users · one-click on-chain · retention benefit · shares chain with Rail 1 |
| TOKEN · $PEPT | SOLANA SPL | protocol coordination · validator stake · governance · Sovereign tier · same rail as payment |
| EXCLUDED | EVM CHAINS | no Base, no Arbitrum, no Ethereum · multi-chain fragmentation outweighs any one-chain's audience benefit |

The trade-off is explicit: clients whose treasuries sit exclusively on EVM chains must bridge to Solana before they can transact. This is an accepted cost. Bridging is a one-time friction; multi-chain settlement is a permanent operational tax. The protocol pays the friction once, at the caller's edge, and is freed from the tax forever.

§ 10.8 · WHY THE MODEL COMPOUNDS

# The moat is the loop between rails.

Each rail would be a respectable business alone. Together they are a category. Rail 2 subscribers contribute outcome events that enrich the P-tier layer. P-tier claims make Rail 1 API responses more specific than anything a pure-research competitor can produce. Rail 1 revenue funds the evidence engine that Rail 2 subscribers experience as a better Agent.

Competitors who ship only a research API cannot access P-tier evidence — they have no users generating outcomes. Competitors who ship only a consumer agent cannot fund a graded evidence graph at the depth Peptome requires — the API rail is what pays for the graph. The architecture is the moat.

The rails are also legally separable. If regulation requires Rail 2 to route through a licensed clinical partner in a given jurisdiction, it does so; Rail 1 is unaffected. If an agent-native client wants volume pricing under a B2B contract, Rail 1 can bill off-chain without touching the consumer surface. The architecture survives its own regulatory scenarios.

The $PEPT token is the coordination primitive that ties all of this together without a trusted intermediary. Fees in, buybacks out, validators staked, rubric governed on-chain. One chain — Solana — is the operational consequence of wanting the rails, the token, and the settlement layer to share one substrate.

*Two rails. One evidence graph. One chain. Each call strengthens the next answer.*

— § 10 · ECONOMICS

§ 11 · ROADMAP

# Directional, not committed.

Peptome's roadmap is expressed in directional quarters. Specific dates and features are subject to revision. Canonical values live in the platform's live documentation. The settlement chain is fixed: Solana.

**Q2 2026**

### Human interface v1.

Search, Card Generator, Protocol Copilot, Discovery Lab alpha. Closed cohort.

**Q3 2026**

### HTTP 402 + Solana Pay beta.

Institutional agent onboarding against the USDC-on-Solana rail. Medical advisor governance council seated.

**Q4 2026**

### Validator network + P-tier opens.

External operators post $PEPT testnet stake; outcome-submission endpoint opens; Discovery Lab admits verified researchers.

**Q1 2027**

### $PEPT TGE.

Meteora DLMM or Jupiter LBP launch, audits published, governance council handoff to on-chain rubric voting. Preconditions in §04.7 must be met first.

**BEYOND**

### Federated graph.

Peptome as a read surface inside third-party clinical research workflow tools, and federated graph integrations with selected academic consortia.

§ 12 · CONCLUSION

# A reference implementation *for the decade ahead.*

Three trends, on different timelines, are arriving at a convergence.

The first is the migration of scientific information consumption from human readers to autonomous agents. This trend is already well underway in software engineering, legal research, and intellectual property analysis. Peptide biology will follow on a timeline measured in years, not decades.

The second is the maturation of programmable stablecoin settlement on Solana. Sub-second finality. Fees measured in hundredths of a cent. Wallet primitives — including durable nonces, priority fees, and Solana Pay transaction requests — that machines can hold, sign with, and rotate without racing a blockhash window. This is the settlement layer the previous decade did not have.

The third is the emergence of evidence grading as infrastructure rather than editorial work. The corpus of scientific claims is too large for any human curation body to grade. Algorithmic, reproducible, versioned grading is not an affordance — it is a requirement of any system that claims to ground its answers.

Peptome is the reference implementation of the intersection. A graded, grounded peptide intelligence layer, exposed equivalently to humans and agents, settled in the unit of programmatic payment, governed by a rubric whose evolution is public.

The platform's ambition is not to own peptide biology. It is to be the surface that peptide biology, in its next phase, is reasoned over from.

# Peptome

*the operating system for human biology.*

INTELLIGENCE EXECUTION CLOSED LOOP $PEPT

*Two rails. One evidence graph.*
*One chain — Solana.*
*Grounded peptide intelligence,*
*served to humans and agents at*
*the same resolution.*

v3.0

SOLANA · $PEPT · DUAL-RAIL · AGENT-NATIVE