HACKING BIOLOGY

BIOHACK.IT

An open laboratory for longevity

WHITE PAPER

Turn self-experimentation into evidence.

Presented at **The 7th TimePie Longevity Forum** · Biohacker track
September 12–13, 2026 · Shanghai

VERSION 0.2 · 2026-09-06
OPEN SOURCE · AGPL-3.0 · NON-PROFIT
BIOHACK.IT

ABSTRACT

Longevity intervention has run far ahead of longevity measurement. People already take off-label drugs, peptides, hormones and undergo therapies in pursuit of a longer healthspan — but the record of what they do lives in spreadsheets, forum posts and chat prompts: not comparable, not followable, and not readable by a physician. A thousand people running the same protocol produce a thousand anecdotes instead of one dataset.

BIOHACK.IT is open infrastructure for structured human self-experimentation. It lets a person document a protocol in a standardised, machine-readable form, measure its efficacy and safety through standards-coded biomarkers, and share both the protocol and its outcomes in the open — so that many individual experiments become one body of comparable human data. Safety is engineered into the act of copying a protocol, which is where beginners are most exposed. The software is free, open source under AGPL-3.0, and stewarded by the non-profit Hacking Biology Foundation.

This paper states the problem, the design, the honest limits of the evidence such a system can produce, and the path from a minimum viable product — turning the lab reports people already have into a legible clinical history — to distributed, self-organised trials. The full functional specification is public.

01 · The problem

Intervention has outrun measurement.

The longevity field can propose more than it can observe. The imbalance is the opportunity.

Over the last decade, the practice of longevity has moved faster than its evidence base. Rapamycin, metformin, GLP-1 agonists, a growing catalogue of peptides, hormone protocols, plasmapheresis, hyperbaric oxygen, exercise and fasting regimes — all are adopted today by motivated individuals, often years ahead of the trials that would tell us whether, and in whom, they work. This is not going to stop, and telling people to wait is neither realistic nor, for many, acceptable.

The waste is not that people experiment. The waste is that the experiments evaporate. A protocol is written as prose in a forum post. A result is a screenshot of a lab report. A regime lives in a personal spreadsheet nobody else can open. None of it is comparable across people, none of it accumulates, and almost none of it is in a form a doctor — or a researcher — could actually use.

> One person with a lab and a research team produces data. Everyone else produces anecdotes. The difference is not effort or sincerity — it is structure.

The premise of this project is narrow and, we think, correct: the single most valuable thing to build for longevity right now is not another compound or another wearable, but the *shared representation layer* that turns what people are already doing into data. That layer does not exist. Building it is the work.

02 · The ethical core

The beginner who copies is the person most at risk.

Harm reduction is not a feature bolted onto the platform. It is the reason the platform is defensible.

Three people stand around any shared protocol, and they are not equally safe.

PRIMARY USER

The experienced biohacker

Designs and self-administers, iterates — but today the results are anecdotes: claimed, not measured, impossible to compare.

HIGHEST RISK

The beginner who copies

Reads a forum thread, starts taking a compound with no baseline, no safety markers, no idea what to monitor.

RECIPIENT, NOT USER

The treating physician

Receives, at best, a handwritten sheet — and deserves a legible, complete view of what is being taken and why.

The person who copies an off-label compound off a forum thread — without knowing to monitor kidney and liver function, without a baseline to compare against — is exactly the case where structured infrastructure protects health rather than merely recording it. Biohackers already take the risk today. They take it *without* a baseline, without inherited safety markers, without an overdue-test warning. The platform does not add risk to the world; it adds a floor of safety under a behaviour that is already happening.

This is the module that justifies the project. If BIOHACK.IT does nothing else, making the first copied protocol safer than a screenshot is worth building.

03 · The core idea

The same protocol, written three ways, is one entity.

The real asset is not an app. It is the standardisation of something nobody has standardised.

Ask three biohackers to describe the same regimen and you get three strings of text:

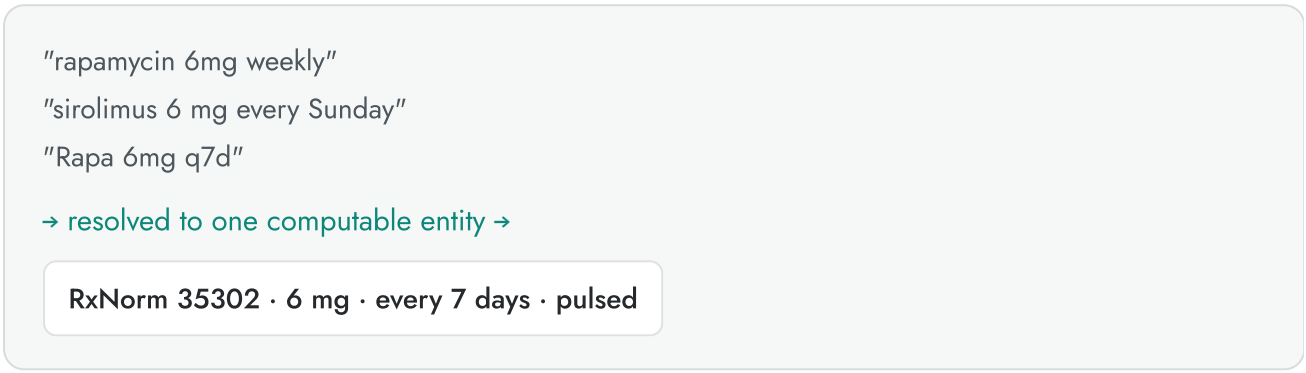

FIGURE 1 — THREE HUMAN DESCRIPTIONS, ONE MACHINE-READABLE INTERVENTION

To a human these are obviously the same thing. To software they are three unrelated strings — and that is the whole problem, because you cannot compare, aggregate or reason over prose. BIOHACK.IT resolves each of them to a single canonical entity: a substance keyed to standard vocabularies (RxNorm, ATC, PubChem, UNII), with a dose in defined units and a real cycle schema. Do this ten thousand times and a forum's worth of prose becomes a dataset.

The same discipline runs through every part of the system. Biomarkers are coded to LOINC and expressed in UCUM units, with the laboratory and assay method carried alongside each value, so that a result from one lab can be compared to a result from another. Nothing the platform stores is free text where a code will do. This is unglamorous, and it is the entire value.

04 · From act to evidence

From one intervention to shared evidence.

A single pipeline connects what one person does to what a community can learn.

FIGURE 2 — THE PIPELINE FROM PERSONAL ACT TO OPEN EVIDENCE

- **Intervention** — what you take and do: dose, timing, real cycle schema, versioned with every change.
- **Measurement** — standards-coded biomarkers (LOINC/UCUM) with lab, method and provenance attached.
- **Individual effect** — your own before/after, bound to the exact protocol version that was active at the time.
- **Cohort** — everyone running the same protocol, weighted by adherence and data completeness.
- **Evidence** — an open dataset the community, and researchers, can actually interrogate.

Each stage is only possible because the one before it is structured. You cannot compute an individual effect over free-text doses; you cannot form a cohort over incomparable biomarkers; you cannot produce evidence over a cohort whose adherence you never recorded. The pipeline is the argument for the standardisation.

05 · What it does

Document your protocol. Measure it. Share it.

Three verbs describe the whole product.

01 — DOCUMENT

What you take & do

Substances, therapies, exercise, nutrition — with dose, timing and real cycle schemas. Create from scratch, or copy any public protocol with one click.

02 — MEASURE

Efficacy & safety

Biomarkers with reference, optimal and safety ranges, plus biological clocks (PhenoAge, epigenetic, glycation). Upload past lab reports, get your history in graphs.

03 — SHARE

Open by default

Publish your protocol and its outcomes. Quality comes from three layers: deterministic validation, community review, expert oversight.

The social mechanic is deliberately familiar: follow and copy the people who get results, seeing their public biomarkers rather than their claims — copy-trading, for health. But unlike a trading feed, what you copy is a safety-instrumented, versioned object, and what you see is measured, not asserted.

The platform is built on the community that already exists, not a forum rebuilt from scratch: community features attach to Rapamycin News, where longevity biohackers already gather.

06 · Architecture

GitHub for biological self-experimentation.

If you have ever used version control, you already know how this works.

A protocol is not a static document; it is a living object that changes over time, and every change matters. Borrowing the mental model of version control makes the design legible at a glance:

repository	a protocol
fork	copying someone's protocol
commit	a dose or intervention change, with its stated reason
release	a protocol version
CI	safety monitoring — baseline, inherited markers, overdue checks
issues	a contested claim, discussed in public
telemetry	your biomarkers over time
dataset	the open research output

FIGURE 3 — THE VERSION-CONTROL METAPHOR, MADE LITERAL

Every change to a protocol is versioned and carries a mandatory reason, so that when an effect appears in the data it can be traced to the exact state of the protocol that produced it.

Three questions, deliberately kept apart

Much of the confusion in existing tools comes from collapsing distinct questions into one record. The domain model keeps them separate:

TestingProtocol
What I measure, and when. A versioned set of measures and timepoints.

TreatmentPlan
What I take and do. Versioned interventions — self-managed or physician-assigned.

Goal
What I want to change. Mapped onto the Hallmarks of Aging.

Substance → SafetyRule → Biomarker → Measurement → Dashboard

FIGURE 4 — "WHAT YOU TAKE" AND "HOW YOU ARE" ARE THE SAME DATA STRUCTURE, READ FROM TWO SIDES

Technically the system speaks in established vocabularies rather than inventing its own: HL7 FHIR internally, OMOP for research export, LOINC and UCUM for analytes, RxNorm/ATC/PubChem/UNII for substances. Evidence content structure follows the Forever Healthy Foundation's Evipedia; AI-

generated review text is produced through the AI4L audit-based-prompting framework. Where AI is and is not used is published as an explicit map — and, as a firm rule, the model writes words, never numbers.

07 · Harm reduction

Follow a protocol — inherit its safety.

Safety is never opt-in, because the person who most needs it is the least likely to choose it.

When a user adopts a protocol, they do not merely copy its interventions. They inherit its safety apparatus automatically:

- **Mandatory baseline** — before you can follow a protocol at all, plus an explicit acknowledgement of the risk you are accepting.
- **Safety markers inherited automatically** — from every active compound. If a substance warrants a liver panel, that panel comes with it; safety monitoring is never something you have to remember to add.
- **Dose sanity check** — a plausibility validation on the numbers, catching the extra zero. This is not clinical advice; it is a guardrail against the obvious mistake.
- **Overdue is a safety gap** — "rapamycin active for 94 days, liver panel overdue by 34" is surfaced as a defect in the protocol, not buried.

> The platform describes; it never prescribes. It tells you what to measure — never what to take, or how much.

Interventions that are medical- or research-grade, and doses that derive from animal rather than human studies, are marked as such and gated behind an informed-decision panel: adoption and outcomes, the evidence corpus and its grading, the medical or research level, dose provenance, potential impacts and links to the relevant forum discussion. Everything remains one-click copyable — nothing is documentable-but-hidden — but nothing is copyable without the context needed to judge it.

08 · The minimum viable product

Turn the lab reports you already have into a clinical history.

The first thing we ship is useful to a single person, with no community and no risk narrative attached.

Ambition kills products by making the first version un-shippable. The MVP is deliberately narrow: **the Blood Layer**. Upload the lab reports you already have — PDFs, images, the mess your labs actually give you — and get your biomarker history in graphs, free, coded to LOINC/UCUM with reference, optimal and safety ranges per analyte.

A HARD BOUNDARY

The model writes words, never numbers

Extraction, mapping and thresholds are deterministic and reproducible. The model only produces the plain-language narration — generated once and cached.

BUILT FOR AGGREGATION

LOINC + UCUM + method, every value

The only architecture compatible with releasing open data. Three ranges per analyte — laboratory reference, longevity-optimal, safety — each with its source.

THE POINT

Useful before the network

A sentence you can say without pronouncing the word "biohacking", with no regulatory exposure — and it works for one user, alone.

The Blood Layer earns trust with a task everyone recognises and no one enjoys — making sense of a stack of lab PDFs — while quietly building the standards-coded substrate that everything else depends on. It is the wedge, not the whole.

09 · Intellectual honesty

We are not a clinical trial, and we say so first.

Credibility comes from declared limits, not from claimed rigour.

This is self-managed, self-declared, self-selected experimentation. There is no randomisation, no blinding, no controlled conditions. People change their protocol halfway through, measure irregularly, and drop out. A platform that hid these facts would deserve the criticism it received; BIOHACK.IT states them on the front page.

> The honest comparison is not BIOHACK.IT versus a clinical trial — we lose that one, and we are not trying to win it. It is BIOHACK.IT versus what exists today: a dose written in prose, a screenshot of a lab report, a spreadsheet nobody else can open.

Against that real baseline, structure is a large improvement. But structure is not the same as rigour, and the platform enforces the difference in how it reports:

- **Every number carries its confidence** — value, uncertainty, provenance and evidence level, or it is not shown at all.
- **Observed, never caused** — findings are phrased as "among qualifying users exposed to X, the observed change was Y." The system does not state or imply causation, on any surface, including AI-generated text.
- **Negative results are welcome** — "I tried X. Nothing happened." is a first-class contribution. The absence of an effect is exactly the data biohacking usually loses.

The specification carries a full statistical-methodology capability that names the threats such data faces — confounding, regression to the mean, healthy-user and survivor bias, measurement error, multiple testing — and requires that under-powered comparisons report "not determinable" rather than a precise-looking number. And it publishes the bar a study must clear before the words *distributed trial* are used: a pre-registered endpoint, a shared protocol version, declared timepoints, an adherence floor, a minimum cohort. Nothing starts as a trial. A study earns the term.

10 · What is unique

What the literature predicts, beside what practitioners actually show.

The originality is not a chart. It is the ability to put two charts on the same axis.

Anyone can draw the line that the literature predicts — it is in the papers. What almost no one can draw is the second line: what a living community of practitioners actually observes, on standards-coded biomarkers, at real doses, over real time. BIOHACK.IT exists to put those two lines side by side on the same axis — outcome against biomarker, prediction against observation — and to let the gap between them be the finding.

This is only possible on a platform where the observed line is built from comparable, provenance-carrying measurements rather than testimonials. It is the payoff of every unglamorous standardisation decision described above, and it is the thing no spreadsheet, forum or wellness app is structurally able to produce.

11 · Openness

Open data first, linked data second, API third.

The research output is not a feature to add later. It is what the whole system is for.

Everything public — profiles, protocols, treatments, measurements, and, behind a heightened and explicit consent gate, genomic data and original lab reports — is published as a clonable open-data snapshot: enough to stand up an independent instance and seed it. This is the first-order output, generated by construction, not an export bolted on.

- **Open & linked data, self-generated** — everything public, coded and linked, sufficient to seed an independent instance.
- **A research API, second** — a documented endpoint returning de-identified, methodologically annotated cohort data: ask for an intervention and a biomarker, get an answer with its caveats attached.
- **Impossible to enclose, by construction** — AGPL-3.0 means anyone running it as a service must open the source; a capable user can self-host and keep the community's data in the community's hands.

Data is published raw rather than pre-obfuscated — a maximal-openness stance that will be put to a community vote on Rapamycin News when the public beta opens. Standards, not silos: HL7 FHIR internally, OMOP for research export, LOINC/UCUM for analytes — so what the platform produces is research data, not another wellness silo. We do not, in any form, aim to publish papers; the output is the open dataset itself, and what others build on it.

12 · Specification & roadmap

The whole specification is public before the first line of product.

Breadth in thinking, discipline in shipping.

The functional specification — 26 capabilities and over 150 requirements, each with concrete scenarios, spanning protocols, biomarkers, safety, community, evidence, statistical methodology and open data — is written down, in the open, and browsable at biohack.it/specs, growing with community review. Specifying broadly costs little; building broadly costs everything. So what we build first is deliberately narrow.

NOW	**Foundations** — domain model, biomarker registry, coded catalogue, safety rules.
MVP	**Useful to one** — the Blood Layer, safety, the doctor sheet. Useful to a single person, alone.
NEXT	**Adherence** — the daily log; thirty seconds a day.
THEN	**Useful to many** — public protocols, follow and copy, pre-registered N-of-1.
THEN	**Open data** — cohorts, open and linked data, then a research API.
DESTINATION	**Distributed trials** — shared protocols with endpoints declared in advance.

FIGURE 5 — MVP-FIRST ROADMAP: EACH STAGE IS USEFUL BEFORE THE NEXT EXISTS

Deferred without regret, so the first release stays shippable: procurement, advanced nutrition, therapeutics scheduling, genomics beyond simple import, and agent access.

13 · Governance & sustainability

The software and the foundation.

An instrument this sensitive should not be owned by a company that could be pressured to close it.

BIOHACK.IT — the software, the brand and the copyright — is owned by the **Hacking Biology Foundation**, a non-profit incorporated in the Próspera ZEDE on Roatán, Honduras. The organisation is Hacking Biology; the software and the initiative are BIOHACK.IT. Keeping the two distinct is deliberate: the software can be forked, self-hosted and outlive any single steward, precisely because the licence and the ownership are structured to prevent enclosure.

The Foundation's board is drawn from a basket of practising biohackers and software hackers — the two communities whose competence the project depends on. Sustainability is pursued through public grants and, where startups build on the open data and standards, aligned equity participation — never by making the platform itself commercial, and never at the cost of the openness of the data.

The project began in 2022, when its founder — a hacker who turned to biology after his mother died of lung cancer — went looking for the instrument he wished for himself in hacking his own biology, and could not find it. BIOHACK.IT is an attempt to build that instrument in the open, for everyone who comes to the same search after him.

14 · Participate

Build this with us.

This is open infrastructure, and it is early. Every one of these roles is genuinely open.

BIOHACKERS

Bring your experience into the specification, and your protocols and measurements into the software.

SOFTWARE HACKERS

Build it. AGPL, public specification, open issues — written down before it is written in code.

RESEARCHERS & SCIENTISTS

Keep us honest on validity, uncertainty and evidence — while keeping it usable by a beginner.

LONGEVITY ENTHUSIASTS

Volunteer to spread it: community, translation, events. A tool nobody knows about produces no data.

FUNDRAISERS & PHILANTHROPY

Build the funding pipeline — grants and philanthropy — so it never has to become commercial.

Contacts & references

EMAIL
research@hackingbiology.com

WEBSITE
biohack.it

SPECIFICATION
biohack.it/specs

SOURCE
github.com/hackingbiology/biohackit

LINKEDIN
company/biohackit · in/secret

X · TWITTER
@fpitrosanti

TELEGRAM
t.me/+0qg9-HC4Nx45OTI8

RAPAMYCIN NEWS
rapamycin.news/u/hackingbiology.com

BIOHACK.IT — An open laboratory for longevity. Turn self-experimentation into evidence.

Open source under AGPL-3.0. A non-profit project of the Hacking Biology Foundation. · This white paper accompanies the public functional specification at biohack.it/specs; where the two differ, the specification is authoritative.