HACKING BIOLOGY

BIOHACK.IT

Functional Specification

REFERENCE DOCUMENT · GENERATED FROM OPENSPEC

26	157	178
CAPABILITIES	REQUIREMENTS	SCENARIOS

Generated from the OpenSpec tree. **Do not edit this document** — edit the specs and regenerate. The narrative companion is the white paper; where the two differ, this specification is authoritative.

VERSION 0.1 · 2026-08-20
OPEN SOURCE · AGPL-3.0 · NON-PROFIT
BIOHACK.IT/SPECS

CONTENTS

§ 0

Principles & product posture

> Settled directions that shape every capability. Where one supersedes an earlier spec decision it says so; these are the frame within which the requirements below are read.

- **Server-hosted, not local-first.** The individual layer runs on the server. There is no browser-only mode.
- _Supersedes the "local-first mode" open question (spec v0.3 §11 Q28) — resolved: no local-first (confirmed by Fabio 2026-08-04)._
- **Self-hosting is a different, supported thing:** the software is AGPL, so a technically capable user may stand up their own full instance and seed it from the public OpenData snapshot.
- **Fully clonable — open by code and by data.** Everything public — public profiles and their public data, defined protocols, treatments/interventions, and measurements — is **OpenData**, published as a clonable snapshot sufficient to run an independent instance. Withheld per-item data is never included; public genomic data (including the gVCF) and original lab report files are included under the same public-by-default rule. See `specs/analytics-and-open-data`.
- **Public by default.** A profile, its protocol and its biomarker outcomes are meant to be public and shared. The UX actively invites sharing rather than defaulting to private.
- _Supersedes spec v0.3 Service Principle #2 ("public by choice, never by default"). Kept from #2: per-marker granularity still exists for the few things a user withholds._
- **Raw by default, validated by the community.** Public data is published **raw** (raw lab reports and gVCF included) — no forced obfuscation. This maximal-openness stance will be put to a community poll on Rapamycin News when the public BETA opens. See `specs/analytics-and-open-data`.
- **Genomics is public too.** Reversing spec v0.3 Principle #9: genomic data — including the gVCF — is public by default like other data, behind a heightened, explicit consent gate (it is maximally identifying and effectively irreversible once public). See `specs/genomics`.
- **The platform describes, never prescribes** (Principle #3). It suggests *what to measure*, never *what to take or how much*.
- **Deterministic where it counts.** Extraction, mapping, threshold and index computation are deterministic and reproducible; the LLM only produces narration, generated once and cached. *The LLM writes words, never numbers* (Principle, M2/M8).
- **Built on Forever Healthy's work.** Evidence content structure follows **Evipedia** (Forever Healthy, CC BY 4.0), integrated via its MCP; AI-generated reviews and evidence queries use the **AI4L** audit-based-prompting framework (MIT). Where AI is and isn't used is published as an AI surfaces map. See `specs/ai-uses-and-attribution`.
- **Fails loudly** (Principle #6) and **stops instead of guessing** (Principle #8).
- **No chicken-and-egg.** At To the team seeds **content-curated protocols** (starting from Fabio's own) and invites a first cohort of **alpha-tester biohackers** with real profiles. Community features attach to Rapamycin News, which already exists.
- **Protocols are living objects.** They are created from scratch, copied/forked from others, or imported from an already-running regimen; they **change over time and every change is versioned and tracked**.
- **Everything is copyable, informed.** Every intervention — peptides, plasmapheresis, IV included — sits in the public **one-click-copyable** catalog; nothing is documentable-but-not-copyable. The mitigation is an

informed-decision panel (adoption + outcomes, evidence corpus & grading, medical/research level, dose provenance, potential impacts, forum links) plus the safety gate. Resolves spec v0.3 §11 Q7.

- **No economic figures** appear in specs or artifacts (per Fabio, 2026-08-04). Cost *tracking as a user feature* may exist; investment/opex numbers do not.

§1 · DOMAIN

Domain Model (cross-cutting)

> Defines the entities every other capability builds on, and the one relation that carries the whole system. This capability owns no screens; it constrains all the others. Derived from `docs/hackingbiology-project-spec.md` §5, extended with the protocol-lifecycle clarifications of 2026-08-04.
>
> The load-bearing relation:
>
> ``` Substance → SafetyRule → Biomarker → Measurement → Dashboard ```
>
> "What you take" and "how you are" are the same structure read from two sides.

§1.1

Separation of the three protocol facets

The system SHALL model what-I-measure, what-I-take, and what-I-want as three distinct, independently versioned entities — `TestingProtocol`, `TreatmentPlan`, and `Goal` — rather than a single fused "protocol" record.

Scenario 1 · A treatment changes without disturbing the measurement plan

- **WHEN** a user edits a dose inside their `TreatmentPlan`
- **THEN** the change creates a new `TreatmentPlan` version
- **AND** the associated `TestingProtocol` and `Goal` are untouched and keep their own version history

Scenario 2 · Goals map to hallmarks

- **WHEN** a user declares a `Goal`
- **THEN** the system SHALL let them map it onto one or more Hallmarks of Aging (Schmauck-Medina 2022 framework)
- **AND** the user's top three goals are flagged as weighted for later analysis

§1.2

Protocol as public composite

The system SHALL expose `Protocol` as a read-facing composite of a `TestingProtocol` + `TreatmentPlan` + `Goal`, carrying an origin of `curated` or `community`, that is the unit which is published, followed, forked, and copied.

Scenario 1 · Publishing composes the facets

- **WHEN** a user publishes a `Protocol`
- **THEN** the composite references specific versions of its three facets
- **AND** it records origin, author, license posture, and any forked-from lineage

§1.3

Intervention carries a cycle schema

The system SHALL represent every `Intervention` with an explicit cycle schema — `pattern` (continuous | pulsed | titration | on-off), `on_days` / `off_days`, cycle length, cycles per year, titration steps — not merely a frequency.

Scenario 1 · Pulsed senolytic is representable

- **WHEN** a user records "dasatinib 100 mg, days 1–3 of each month"
- **THEN** the intervention stores pattern=pulsed with on_days=[1,2,3] over a monthly cycle
- **AND** the pattern is never flattened to "Monthly" with the cycle relegated to free text

§1.4

Intervention subtypes

The system SHALL support intervention subtypes with their own fields: Substance, Exercise, Device/Therapy, Procedure, Hormonal, and Nutrition/Fasting.

Scenario 1 · A procedure differs from a pill

- **WHEN** the intervention is a Procedure (e.g. IV infusion, plasmapheresis)
- **THEN** it carries operator, site, consent and a distinct risk profile
- **AND** the system does not treat it as interchangeable with an oral Substance

§1.5

Biomarker with role and three ranges

The system SHALL model `Biomarker` with a role (`efficacy` | `safety` | `baseline`), a collection modality, and three distinct ranges per analyte — laboratory reference, longevity-optimal, and safety threshold.

Scenario 1 · The three ranges stay distinct

- **WHEN** a biomarker value is displayed
- **THEN** the reference range, the optimal range, and the safety threshold are shown as three different things
- **AND** they are never collapsed into a single band

§1.6

Measurement provenance

The system SHALL store each `Measurement` with date, source, laboratory, method/assay, unit (UCUM), and validation state.

Scenario 1 · Method and lab are mandatory metadata

- **WHEN** a measurement is persisted
- **THEN** it retains its method/assay and originating laboratory
- **AND** a measurement missing these is flagged as not aggregation-eligible

§1.7

Safety, adherence and wellness as first-class entities

The system SHALL model `SafetyRule`, `AdherenceLog` (four states: taken-as-planned | partial | intentionally-skipped | forgotten), and `WellnessCheck`, and treat them as core, not optional.

Scenario 1 · Adherence qualifies an N-of-1 result

- **WHEN** a cohort result is computed
- **THEN** each contribution carries its adherence percentage
- **AND** a low-adherence contribution can be down-weighted or excluded (see `analytics-and-open-data`)

§1.8

Cohort and Study

The system SHALL model `Cohort` (the set of people practising the same protocol) and `Study` (a self-experiment container with pre-declared endpoints and timepoints).

Scenario 1 · Cohort forms around a protocol

- **WHEN** two or more people practise the same published `Protocol`
- **THEN** they belong to its `Cohort`
- **AND** the cohort is the unit of comparison and analytics

§2 · AI USES

AI Uses, Provenance & Previous Work (cross-cutting)

> Consolidates three things the scientific and biohacker audience asks first: **where AI is used and — crucially — where it is not**, **which frameworks and external knowledgebases the platform builds on**, and **the attribution and licensing of prior work**. It adopts the published "AI surfaces map" pattern and names the shoulders biohack.it stands on — above all **Forever Healthy's Evipedia and AI4L**, plus `get-based` and the open vocabularies. It also makes biohack.it a **connector** that evaluates, credits and notifies the upstream projects whose data effort it binds together. This capability governs the others; it owns the trust artifact, not a product screen.

§2.1

Published AI surfaces map — "what is AI, what is NOT AI"

The system SHALL maintain and publish a canonical map of every surface where AI is used and every surface where it is deliberately not, so a reader can trust the numbers.

Scenario 1 · The numbers are not AI

- **WHEN** a user views a computed value (safety threshold, index, sex/age percentile, trend, schedule)
- **THEN** the AI surfaces map documents these as deterministic and reproducible, not AI-generated
- **AND** only extraction, narration, and evidence-review generation are marked as AI surfaces

§2.2

AI4L as the framework for AI-generated reviews and evidence queries

The system SHALL use Forever Healthy's **AI4L** audit-based-prompting framework for any AI-generated evidence review or evidence data query, rather than ad-hoc prompting.

Scenario 1 · Audited generation, not single-shot

- **WHEN** the platform generates or refreshes an evidence review, or answers an evidence query with AI
- **THEN** it runs the AI4L create → audit → correct cycle against the QA checklist, restricted to trusted scientific sources
- **AND** outputs are cached and reproducible, and the LLM never writes numbers into constrained fields (see `interventions-and-catalog`, `biomarkers-and-labs`)

§2.3

AI4L output conforms to the project's strict normalization

The system SHALL ensure AI4L-generated reviews and evidence queries conform to biohack.it's strict data normalization — analytes coded to LOINC/UCUM, substances resolved to RxNorm/ATC/PubChem/UNII, outcomes mapped to biomarkers — rather than free-text entities.

> OPEN
>
> OPEN: technical feasibility to decide — either (a) **extend AI4L to use these databases natively** (LOINC/UCUM/RxNorm/ATC/PubChem/UNII), or (b) **extend AI4L's prompting to use biohack.it's databases**. Evaluate both.

Scenario 1 · Generated review is normalized

- **WHEN** AI4L produces or refreshes a review, or answers an evidence query
- **THEN** its interventions, analytes and outcomes are resolved to the project's coded vocabularies (see `interventions-and-catalog`, `biomarkers-and-labs`)
- **AND** non-resolvable entities are marked ambiguous for review, never left as silent free text

§2.4

Integrate the Evipedia knowledgebase, attributed

The system SHALL integrate Forever Healthy's **Evipedia** as an external evidence knowledgebase via its MCP server (`mcp.evipedia.ai`) and public endpoints (`/reviews.json`, `/search.json`, `/{slug}.md`, `/{slug}.meta.json`), and SHALL attribute Forever Healthy per Evipedia's **CC BY 4.0** license wherever its content is surfaced.

Scenario 1 · Review retrieved and attributed

- **WHEN** an intervention has an Evipedia review
- **THEN** the platform can retrieve its conclusion, metadata (alternate names, PMIDs), and full review via the Evipedia MCP or JSON/MD endpoints
- **AND** any surfaced Evipedia content carries visible attribution to Forever Healthy (CC BY 4.0)

§2.5

Content structure aligned to the Evipedia / AI4L review model

The system SHALL align its intervention-evidence structure — conclusion, graded evidence, risk-benefit, ordered citations with PMIDs, review dates — with the Evipedia/AI4L model, so content interoperates and can be reused both ways.

Scenario 1 · Interoperable structure with our differential on top

- **WHEN** the platform stores an intervention's evidence
- **THEN** its structure maps to the Evipedia review model (conclusion, grading, risk-benefit, citations, dates)
- **AND** biohack.it's differential — the Outcome↔Biomarker link and the observed-cohort column — is layered on top (see `evidence-layer`, `dashboards-and-doctor-view`)

§2.6

Previous-work and license register

The system SHALL maintain a register of reused prior work with its license and attribution, and SHALL keep every integration compatible with the project's AGPL-3.0 license.

Scenario 1 · License recorded before reuse

- **WHEN** a component or dataset is reused — Evipedia (CC BY 4.0), AI4L (MIT), `get-based` (AGPL), LOINC / UCUM / RxNorm / ATC / PubChem / UNII
- **THEN** its license and required attribution are recorded and shown
- **AND** any license incompatibility is flagged before integration proceeds

Scenario 2 · Public trust artifact

- **WHEN** a physician or researcher asks "can I trust these numbers?"
- **THEN** the published AI surfaces map + attribution register answers where AI runs, where it does not, and whose vetted knowledge is reused

§2.7

Technical evaluation of reused open-source / open-data components

The system SHALL maintain a technical evaluation of every reused open-source and open-data component — what it is, its fit, license, maintenance/health, and how it is integrated — with the rationale for use.

Scenario 1 · Component evaluated before adoption

- **WHEN** an OSS or open-data component is considered for reuse
- **THEN** it is evaluated for fit, license compatibility, maintenance and integration approach
- **AND** the evaluation and rationale are recorded (see `data-standards-and-typing`, `related-initiatives`)

§2.8

Notify upstreams of their inclusion

The system SHALL notify each upstream project whose work is included or built upon, so biohack.it acts as a **connector** of others' data effort rather than a silent consumer.

Scenario 1 · Upstream is told, credited, and given back

- **WHEN** a component or dataset is included (e.g. Evipedia, get-based)
- **THEN** its maintainers are notified of the inclusion and credited
- **AND** contributions back (e.g. LOINC mapping upstream) are offered where they help

§3 · STANDARDS

Data Standards & Typing (scientific validity)

> A dedicated section documenting the data structures and standards in use and **why** — the backbone that makes biohack.it's data typed, comparable across people and labs, aggregatable, and scientifically credible. It consolidates the coding and typing discipline otherwise scattered across the modules, for scientific and technical reviewers.

§3.1

Document each standard and its rationale

The system SHALL document each data standard in use and why it is used — HL7 FHIR (internal clinical vocabulary), OMOP CDM (research export), LOINC + UCUM (analytes and units), RxNorm / ATC / AIFA (drugs), PubChem / ChEBI / UNII (molecules), CPIC / PharmGKB (pharmacogenomics), the Hallmarks of Aging framework, and gVCF (genomics).

Scenario 1 · Rationale is explicit

- **WHEN** a standard is adopted
- **THEN** the section states what it is used for and why it matters for validity
- **AND** alternatives considered are noted where relevant

§3.2

Typing discipline for scientific validity

The system SHALL enforce typing discipline — closed enums populated by deterministic code, free text kept separate, and entities resolved to codes — so data is typed rather than free-form.

Scenario 1 · Typed, not free-form

- **WHEN** data is persisted
- **THEN** coded fields resolve to their vocabularies and free text is separate (see `interventions-and-catalog`, `biomarkers-and-labs`)
- **AND** untyped/unresolved entries are marked ambiguous, never left as silent free text

§3.3

Coding plus provenance make aggregation valid

The system SHALL treat coding (LOINC/UCUM/…) plus method/lab provenance as the precondition for cross-person / cross-lab comparability and open-data validity.

Scenario 1 · Only comparable data aggregates

- **WHEN** data is aggregated or released for research
- **THEN** comparison uses coded, provenance-carrying values (lab-relative z-scores) (see `analytics-and-open-data`)
- **AND** uncoded data is excluded from the aggregated research export

§3.4

Published and maintained as a dedicated section

The system SHALL publish this rationale as a dedicated, maintained section for scientific and technical reviewers.

Scenario 1 · Reviewer-facing section

- **WHEN** a scientist or engineer evaluates the platform
- **THEN** the data-standards-and-typing section explains the model and its validity basis

§4 · METHOD

Scientific Method & Evidence (cross-cutting)

> States publicly what this data can and cannot support, so credibility comes from declared limits rather than from claimed rigour. It carries the project's honest position — **self-managed, self-declared, self-selected experimentation, not a clinical trial** — the evidence hierarchy that ranks what we hold, the bar a study must clear before the words *distributed trial* are used, and the cultural commitment to negative results.
>
> The honest comparison is not against a controlled trial, which this loses and does not try to win. It is against what exists today: a dose written in prose, a screenshot of a lab report, a spreadsheet nobody else can open.

§4.1

Publicly declared nature and limits

The system SHALL publish, as a first-class page, that its data is self-managed, self-declared and self-selected, without randomisation, blinding or controlled conditions, and that participants change protocols mid-course, measure irregularly and drop out.

Scenario 1 · Limits are stated before claims

- **WHEN** a visitor or researcher encounters the platform's results
- **THEN** the declared nature and limits are available and linked from where results are shown
- **AND** the stated comparison is against unstructured practice (prose, screenshots, spreadsheets), never against a controlled trial

§4.2

Evidence hierarchy

The system SHALL declare and apply an evidence hierarchy — controlled trial > observational > N-of-1 > anecdote — and SHALL state when an N-of-1 is genuinely informative and when it is not.

Scenario 1 · A claim is placed in the hierarchy

- **WHEN** evidence supports a statement on the platform
- **THEN** its level in the hierarchy is shown alongside it (see `evidence-layer`)
- **AND** self-reported data is never presented at the level of controlled evidence

§4.3

Qualification bar for a distributed trial

The system SHALL publish the criteria a study must meet before it is described as a distributed trial — pre-registered endpoint, shared protocol version, declared timepoints, adherence floor, minimum cohort size, published analysis plan — and SHALL apply the term only to studies that meet them.

RESOLVED

RESOLVED (2026-08-19): distributed trials are kept as the stated destination; the word is earned by publishing the bar rather than avoided. Nothing *starts* as a trial — a study becomes trial-like only when enough people run the same protocol with endpoints declared in advance.

Scenario 1 · The term is applied only when earned

- **WHEN** a study does not meet every published criterion
- **THEN** it is presented as a self-experiment or a cohort observation, not as a distributed trial
- **AND** the unmet criteria are visible

§4.4

Observed, never caused

The system SHALL phrase aggregate findings as observation, not causation — "among qualifying users exposed to X, the observed change was Y" — and SHALL NOT state or imply that an intervention caused an outcome.

Scenario 1 · Aggregate output is phrased as observation

- **WHEN** a cohort result is rendered anywhere in the product or exports
- **THEN** it is expressed as an observed change in a self-selected group
- **AND** causal phrasing is not produced by any surface, including AI-generated narration

§4.5

Evidence confidence on every displayed number

The system SHALL accompany every meaningful number it displays or exports with four things — **value, uncertainty, provenance and evidence level** — and SHALL suppress or mark as not-determinable any number that cannot carry them.

Scenario 1 · A number without its confidence is not shown

- **WHEN** a computed or measured value is rendered anywhere in the product, the doctor sheet, or an export
- **THEN** it carries its uncertainty, where it came from, and the strength of evidence behind it
- **AND** if any of those is unavailable, the value is marked not-determinable rather than displayed bare

Scenario 2 · Evidence confidence travels into exports

- **WHEN** data leaves the platform through open data or the research API
- **THEN** the same four attributes travel with each value (see `analytics-and-open-data`)

§4.6

Negative results are first-class

The system SHALL accept, display and encourage null and negative outcomes as valid contributions, with the same standing as positive ones.

Scenario 1 · A null result is contributed

- **WHEN** a user reports that an intervention produced no detectable change
- **THEN** the result is recorded and shown as a legitimate outcome
- **AND** it counts toward cohort aggregation and reputation for data quality (see `community-and-social`)

§5 · STATS

Statistical Methodology (cross-cutting)

> Governs how the platform reasons about numbers, so that aggregate output is defensible rather than merely computed. It names the threats to validity that self-selected, self-reported longitudinal data carries, separates biological change from analytical noise, and makes uncertainty mandatory. Companion to `scientific-method-and-evidence`, which owns the declared limits and the phrasing discipline; this capability owns the arithmetic behind them.

§5.1

Named threats to validity

The system SHALL declare, and where possible surface or mitigate, the threats to validity carried by its data — confounding, regression to the mean, healthy-user bias, survivor and drop-out bias, selection at entry, measurement error, and multiple testing.

Scenario 1 · A cohort result names its threats

- **WHEN** an aggregate result is produced for a cohort
- **THEN** the applicable threats to validity are stated with the result
- **AND** the result is never presented as if these had been controlled for

§5.2

Regression to the mean and multiplicity are handled explicitly

The system SHALL account for regression to the mean when a cohort is defined by an extreme baseline, and SHALL correct or declare multiplicity when several biomarkers or interventions are examined together.

Scenario 1 · Extreme baseline cohort

- **WHEN** a cohort is selected on an out-of-range baseline value
- **THEN** the expected regression-to-the-mean component is declared alongside the observed change

Scenario 2 · Many markers examined at once

- **WHEN** an analysis spans multiple biomarkers or interventions
- **THEN** the number of comparisons is declared and correction is applied or its absence stated

§5.3

Separate biological change from analytical noise

The system SHALL distinguish intra-individual biological variability and assay drift from genuine change, using the method/assay and laboratory provenance stored with every measurement.

Scenario 1 · Laboratory or method changes mid-series

- **WHEN** a series contains measurements from different laboratories or assays
- **THEN** the change point is flagged in the series
- **AND** the comparison uses lab-relative normalisation rather than raw values (see `analytics-and-open-data`)

Scenario 2 · Change within known biological variability

- **WHEN** an observed change falls within the analyte's known intra-individual variability
- **THEN** it is presented as indistinguishable from noise, not as an effect

§5.4

Missing data is declared, never imputed silently

The system SHALL declare missingness and drop-out in any aggregate, and SHALL NOT silently impute values.

Scenario 1 · Incomplete follow-up

- **WHEN** participants lack follow-up measurements
- **THEN** the aggregate states how many contributed at each timepoint and how many dropped out
- **AND** any imputation is explicit, labelled, and separable from measured values

§5.5

Uncertainty and effect size are mandatory

The system SHALL express every reported effect with its uncertainty and in the outcome's native units, and SHALL NOT derive a trend from insufficient data.

Scenario 1 · No interval, no effect

- **WHEN** an effect is reported
- **THEN** it carries an uncertainty interval and native units
- **AND** a trend is not computed from two measurements (see `analytics-and-open-data`, "stops instead of guessing")

§5.6

Under-powered comparisons are declared, not hidden

The system SHALL state when a cohort is too small or too sparse to support a comparison, rather than presenting a precise-looking number.

Scenario 1 · Cohort too small

- **WHEN** a comparison falls below the declared minimum cohort or coverage
- **THEN** the system reports "not determinable" with what is missing

§6 · LANDSCAPE

Related Initiatives (ecosystem landscape)

> A clearly highlighted, maintained public register of similar initiatives — commercial and non-commercial — across longevity/biohacking data, protocols, and biomarkers, stating what each does, its data and licensing posture, and biohack.it's relationship to it (reuse, complement, differentiate, or contact). Honesty about the ecosystem is a credibility and acquisition surface, and the basis for being a connector rather than a walled garden. (This is the dedicated home for the landscape that was deliberately kept out of the conference deck.)

§6.1

Maintain a register of similar initiatives

The system SHALL maintain a register of similar initiatives — commercial and non-commercial — each with what it does, its data/licensing posture, and biohack.it's relationship to it.

Scenario 1 · An initiative is recorded with our relationship

- **WHEN** a similar initiative is identified (e.g. Evipedia / Forever Healthy, get-based, ProtocolEngine, MyAgingTests / Clock Foundation, Lamplit, Lucis, TruDiagnostics, GlycanAge, Bryan Johnson / Blueprint)
- **THEN** it is recorded with its description, data/licensing posture, and our relationship
- **AND** the entry declares whether it is commercial or non-commercial

§6.2

Classify the relationship

The system SHALL classify each initiative's relationship as reuse, complement, differentiate, or contact-upstream.

Scenario 1 · Reuse vs differentiate

- **WHEN** an initiative's code or data is reusable under a compatible license
- **THEN** it is classified `reuse` (see `ai-uses-and-attribution`)
- **AND** a proprietary database we only study is classified `differentiate`

§6.3

Prominently surfaced and current

The system SHALL surface this register as a clearly highlighted public section and keep it current as the field moves.

Scenario 1 · Public, living section

- **WHEN** a visitor looks for how biohack.it relates to the rest of the field
- **THEN** the register is prominently accessible and dated
- **AND** it is updated as initiatives change

§7 · ACCOUNTS

Accounts & Profiles

> Registration, the public profile, onboarding journeys, and the To seeding of curated content and invited alpha testers. Server-hosted; no local-first mode. Public-by-default posture starts here.

§7.1

Server-hosted account

The system SHALL require a server-hosted account to create or share protocols and measurements; there is no browser-only/local-first mode.

RESOLVED

RESOLVED (2026-08-04): confirmed — server-hosted, no local-first. Self-hosting a full instance is a different, supported path (below).

Scenario 1 · Account needed to contribute

- **WHEN** a visitor wants to save a protocol or a measurement
- **THEN** the system requires registration/sign-in
- **AND** read access to public profiles and protocols does not require an account

§7.2

Self-hosting a full instance (distinct from local-first)

The system SHALL be self-hostable as a full instance under AGPL-3.0 by a technically capable user, seedable from the public OpenData snapshot; this is distinct from a per-user local-first mode, which does not exist.

Scenario 1 · Power user stands up their own instance

- **WHEN** a technical user deploys the AGPL software and imports the public OpenData snapshot
- **THEN** they obtain a running server instance populated with the public protocols, treatments, measurements and public profiles
- **AND** this is a full instance, not a browser-only local-first mode

§7.3

Public-by-default profile

The system SHALL make a user's profile, protocols and biomarker outcomes public by default, and SHALL present sharing as the expected path while allowing per-item withholding.

RESOLVED

RESOLVED (2026-08-04): confirmed — public by default (supersedes spec v0.3 Principle #2). Genomics is public too, gVCF included, behind heightened consent (reverses Principle #9; see `genomics`).

Scenario 1 · Onboarding invites sharing

- **WHEN** a user completes onboarding
- **THEN** the UI explicitly invites them to make their protocol and outcomes public and explains the collective benefit
- **AND** the user can withhold specific items, but is not defaulted into a fully private profile

§7.4

Required biological profile

The system SHALL require the user to configure their biological attributes — date of birth (for age), sex, ethnicity, height, and other stable or slowly-changing traits — as specifically as possible, used for sex/age normalization and stratification.

Scenario 1 · Stable traits captured specifically

- **WHEN** a user sets up their profile
- **THEN** they configure date of birth, sex, ethnicity, height and other stable traits as specifically as possible
- **AND** these feed sex/age-normalized percentiles and cohort stratification (see `dashboards-and-doctor-view`, `analytics-and-open-data`)

§7.5

Strong invitation to link social and forum identities

The system SHALL strongly invite the user to configure their social and community identities — Telegram, WeChat, Xiaohongshu (RED / "Little Red Book"), other social profiles, and their Rapamycin News forum nickname — surfaced on the public profile.

Scenario 1 · Social and forum handles linked

- **WHEN** a user completes onboarding
- **THEN** they are strongly invited to add Telegram, WeChat, Xiaohongshu (RED) and their Rapamycin News nickname
- **AND** these appear on the public-by-default profile and link community identity (see `community-and-social`)

§7.6

Onboarding "bring your reports"

The system SHALL onboard users around importing lab reports they already have, targeting first value within ten minutes at zero spend, with a persistent multi-step progress indicator.

Scenario 1 · First chart within minutes

- **WHEN** a new user uploads a past lab report during onboarding
- **THEN** the system parses and charts it (see `biomarkers-and-labs`)
- **AND** the empty state teaches ("upload a report you already have"), never sells a test
- **AND** the original report file is retained and is public by default in its raw original format (see `biomarkers-and-labs`, `analytics-and-open-data`)

§7.7

Two independent reputation axes

The system SHALL model reputation as two separate axes — verified professional credential (who you are) and data verification (what your measurements show) — and SHALL never fuse them into one score.

Scenario 1 · A physician without data, a biohacker with data

- **WHEN** a physician verifies a credential but has no measurements
- **THEN** their credential axis is populated and their data-verification axis is empty
- **AND** discovery can filter by specialty/credential separately from data quality

§7.8

T0 alpha-tester seeding

The system SHALL support inviting a first cohort of biohackers as alpha testers with real, attributed profiles, so that curated content and early community protocols coexist at launch.

Scenario 1 · Invited alpha tester publishes

- **WHEN** an invited alpha biohacker accepts and imports their running protocol
- **THEN** their profile and protocol become part of the seeded, browsable library
- **AND** the platform never presents an empty community at launch (curated + Fabio + alphas)

§8 · DELEGATION

Management & Delegation (Mentor / Longevity Doctor / Clinic)

> Supports a third-party **managing user** — a Biohacking Mentor / Longevity Doctor / Longevity Clinic — who supervises and manages one or more users' protocols: administering, measuring, and monitoring on their behalf. Kept deliberately lean: a user signs up on their own and may then assign a manager; if assigned, it is indicated. Builds on existing hooks — TreatmentPlan `Physician-Assigned`, attributable change events, the doctor view, and the verified-credential reputation axis — so it is additive rather than a rewrite.

§8.1

Managing-user role

The system SHALL support a managing-user role — Biohacking Mentor / Longevity Doctor / Longevity Clinic — able to manage one or more users as a roster.

Scenario 1 · A manager holds a roster

- **WHEN** a managing user is set up
- **THEN** they can hold a roster of one or more managed users
- **AND** the role carries the verified-credential axis (see `accounts-and-profiles`)

§8.2

Self-signup, then assign a manager

The system SHALL let a user sign up on their own and then assign — and revoke — a manager who follows them; assignment is the user's choice and never required to use the platform.

Scenario 1 · User assigns and can revoke a manager

- **WHEN** a self-registered user assigns a manager
- **THEN** a managed-by relationship is created, revocable by the user at any time
- **AND** using the platform without a manager remains fully possible

§8.3

Management scope — administer, measure, monitor

The system SHALL let an assigned manager administer (author/adjust the managed user's TreatmentPlan as `Physician-Assigned`), measure (enter/import measurements), and monitor (view dashboards), within what the user has granted.

Scenario 1 · Manager adjusts a managed protocol

- **WHEN** a manager edits a managed user's TreatmentPlan
- **THEN** the change is applied as `Physician-Assigned` with a mandatory reason (see `protocols`)
- **AND** it is attributed to the manager, not the user

§8.4

Manager actions are attributed

The system SHALL attribute every managing action to the manager in the tracked change/measurement history, visible to the managed user.

Scenario 1 · Who did what is clear

- **WHEN** a manager performs an action on a managed user
- **THEN** the history records the manager as actor
- **AND** the managed user can see what their manager did

§8.5

An assigned manager is indicated

The system SHALL indicate, on the managed user's profile and dashboard, that a manager is assigned and who they are.

> OPEN
>
> OPEN (kept lean for now): fine-grained permission scopes, and who owns the safety-gate risk acknowledgment when a manager assigns a protocol (managed user vs manager). To be detailed later without blocking the MVP relationship.

Scenario 1 · Assignment is shown

- **WHEN** a user has an assigned manager
- **THEN** the profile and dashboard indicate the manager (name + credential)
- **AND** revoking the assignment is reflected everywhere it was shown

§9 · PROTOCOLS

Protocols (M1 — Protocol Builder & Library)

> The authoring, versioning, and lifecycle of protocols: created from scratch, **copied/forked** from another user, or **imported from an already-running regimen**. Protocols are living objects — they change, and every change is tracked. Covers the curated vs community origin and the To seeding that removes the chicken-and-egg problem. Discovery and the social act of following are specified in `community-and-social`; this capability owns the object and its history.

§9.1

Create a protocol from scratch

The system SHALL let a user author a `Protocol` by composing a `TestingProtocol`, a `TreatmentPlan`, and one or more `Goal`s, in a builder that does not require all three to be complete before saving a draft.

Scenario 1 · Draft saved incrementally

- **WHEN** a user adds a single intervention and saves
- **THEN** the protocol persists as a `Draft`
- **AND** the system surfaces what is still missing (e.g. "no efficacy markers yet") without blocking the save

Scenario 2 · System proposes the measurement side

- **WHEN** a user adds interventions to a `TreatmentPlan`
- **THEN** the system SHALL propose the efficacy and safety biomarkers implied by those interventions (see `safety-guardrails`, `measurement-planning`)
- **AND** the proposals are shown as an Accept-All / Reject-All block, never auto-applied

§9.2

Import an already-running protocol (experienced baseline)

The system SHALL support an onboarding path for an experienced biohacker who is **already running a protocol**, capturing the current regimen and its historical baseline rather than assuming a clean start. Where the analytic history is hard to reconstruct, the system SHALL accept a coarse free-text lineage note and SHALL flag the record as starting from a summarized, non-analytic history.

> RESOLVED
>
> RESOLVED (2026-08-04): record `in_effect_since ≠ tracked_since`; where reconstruction is hard, accept a coarse "how you got here" note AND flag the record as `history: synthesized` (versus `analytic`).

Scenario 1 · Baseline capture for someone mid-protocol

- **WHEN** a new experienced user declares interventions they are already taking
- **THEN** the system records `in_effect_since` separately from `tracked_since` and marks the protocol `Active` from import
- **AND** it invites (does not require) backfilling past measurements and prior changes as history

Scenario 2 · Synthesized history when reconstruction is hard

- **WHEN** the user cannot reconstruct the analytic history of how the protocol evolved
- **THEN** the system accepts a coarse free-text lineage note ("how you arrived here")
- **AND** it flags the protocol/baseline as `history: synthesized` (not `analytic`) so downstream analytics can distinguish it (see `analytics-and-open-data`)

§9.3

Copy / fork an existing protocol

The system SHALL let a user copy another user's public `Protocol`, producing a new protocol that records its `forked_from` lineage (source protocol + version).

Scenario 1 · Forking preserves lineage and safety

- **WHEN** a user copies a public protocol
- **THEN** the new protocol references the source protocol and the exact source version
- **AND** the inherited safety markers and mandatory baseline travel with it (see `safety-guardrails`)

Scenario 2 · Fork then diverge

- **WHEN** the user edits their fork
- **THEN** changes apply only to their copy
- **AND** the lineage link to the source is retained for later comparison

§9.4

Versioning and change tracking

The system SHALL version every protocol facet, and SHALL record each change as a tracked, timestamped, attributable event with a **mandatory reason**.

Scenario 1 · A dose change is a tracked event

- **WHEN** a user changes rapamycin from 6 mg/week to 8 mg/week
- **THEN** a new version is created with a change event {field, old, new, timestamp, actor, reason}
- **AND** the change cannot be saved without a reason
- **AND** the prior version remains retrievable and referenceable by measurements taken under it

Scenario 2 · Diff between versions

- **WHEN** a user or viewer compares two versions of a protocol
- **THEN** the system SHALL render a field-level diff of interventions, doses, cycles, and measures

Scenario 3 · Measurements bind to the version in effect

- **WHEN** a measurement is recorded on a date
- **THEN** it is associated with the protocol version active on that date
- **AND** later edits to the protocol never silently reinterpret past measurements

§9.5

Curated and community origin

The system SHALL tag each protocol with origin `curated` (authored/vetted by the HackingBiology team) or `community`, and SHALL make origin visible wherever a protocol is shown.

Scenario 1 · Curated seed exists at T0

- **WHEN** the platform launches
- **THEN** a set of team-curated protocols (including Fabio's own) SHALL be present as `curated`
- **AND** a first cohort of invited alpha-tester profiles can publish `community` protocols on top of them (see `accounts-and-profiles`)

§9.6

Publish with per-facet visibility

The system SHALL let a user publish a protocol public-by-default, while allowing specific measures or interventions to be withheld per-item.

Scenario 1 · Public by default, withhold one marker

- **WHEN** a user publishes a protocol
- **THEN** the protocol and its outcomes are public unless the user explicitly withholds an item
- **AND** genomic elements follow the same public-by-default rule, behind heightened consent (see `genomics`)

§9.7

Protocol states

The system SHALL support protocol lifecycle states `Draft | Active | Paused | Completed` and SHALL reflect state in discovery and dashboards.

Scenario 1 · Pausing stops overdue accrual

- **WHEN** a user sets a protocol to `Paused`
- **THEN** measurement-plan Overdue calculations for that protocol pause (see `measurement-planning`)
- **AND** the pause is itself a tracked change event

§9.8

Protocol level classification (medical | research)

The system SHALL classify a protocol by level — `medical` (established, clinician-grade) or `research` (experimental) — and display the level wherever the protocol appears, so a follower sees what kind of protocol they are copying.

Scenario 1 · Research-level protocol is labelled

- **WHEN** a protocol contains experimental interventions
- **THEN** it is classified and displayed as `research` level
- **AND** a `medical`-level protocol is labelled distinctly, with the acknowledgment implication handled in `safety-guardrails`

§10 · INTERVENTIONS

Interventions & Substance Catalog

> The catalog of substances and the entity-resolution discipline that keeps a multilingual, `.it` project from silently ignoring what users enter. Owns `Substance` records, external coding, synonyms, and the input→output diff. Extends spec v0.3 §8.5quater.

§10.1

Substance resolution on codes, not English strings

The system SHALL resolve every substance to external identifiers (RxNorm / ATC / PubChem / UNII) via a multilingual synonym table, never on raw English strings.

Scenario 1 · Italian input resolves

- **WHEN** a user enters "rapamicina"
- **THEN** the system resolves it to the same entity as "rapamycin"/"sirolimus" via synonyms
- **AND** the resolved entity carries its external codes for downstream safety and aggregation

Scenario 2 · Unresolved substance is surfaced, not dropped

- **WHEN** a substance cannot be resolved to a code
- **THEN** the entry is marked `ambiguous` and queued for review
- **AND** it is never silently discarded or matched to the wrong entity

§10.2

LLM never writes into constrained columns

The system SHALL populate closed enums by deterministic code after validation; free text goes to a separate wide notes field; out-of-enum proposals are marked `ambiguous`.

Scenario 1 · Over-long model output does not corrupt persistence

- **WHEN** the extractor proposes a value that does not fit a constrained field
- **THEN** the value is routed to review as `ambiguous`
- **AND** the insert neither truncates silently nor loses sibling entries

§10.3

Input↔output diff at end of extraction

The system SHALL always show the diff between what the user submitted and what was recognised.

Scenario 1 · 8 in, 5 recognised

- **WHEN** a user submits eight substances and five resolve
- **THEN** the system shows "8 submitted, 5 recognised, 3 not understood" with the three listed
- **AND** the user can correct or confirm the three before anything is committed

§10.4

Substance catalog metadata

The system SHALL store per `Substance`: name(s), external identifiers, class, known interactions, and associated safety markers.

Scenario 1 · Safety markers travel with the substance

- **WHEN** a substance with known safety markers is added to a plan
- **THEN** its associated safety biomarkers become candidates for the measurement plan (see `safety-guardrails`)

§10.5

Long extraction jobs are queued and resumable

The system SHALL run long parsing/extraction jobs in an idempotent queue with resume, never as an interactive request that can time out with quota consumed and no result.

Scenario 1 · Job survives a timeout

- **WHEN** an extraction job exceeds an interactive limit
- **THEN** it continues in the queue and its partial progress is retained
- **AND** re-running is idempotent rather than duplicating entries

§10.6

Seed and enrich the catalog from Evipedia

The system SHALL seed and enrich its substance/intervention catalog from Evipedia's interventions and their alternate names (see `ai-uses-and-attribution`), to bootstrap coverage and multilingual entity resolution.

Scenario 1 · Alternate names aid resolution

- **WHEN** the catalog is seeded or enriched from Evipedia
- **THEN** Evipedia intervention names and synonyms populate the synonym table alongside RxNorm/ATC/PubChem/UNII
- **AND** Forever Healthy is attributed (CC BY 4.0)

§10.7

Posology provenance (human vs animal)

The system SHALL record, per intervention posology, whether the dosing derives from `human` or `animal` experimentation, with its source, so extrapolated dosing is never presented as established.

Scenario 1 · Animal-derived dose is marked

- **WHEN** a dose is extrapolated from animal studies
- **THEN** the posology is tagged `animal-derived` with its source
- **AND** a dose from human trials is tagged `human-derived` (see `evidence-layer`)

§10.8

Principal peptides present in the catalog

The system SHALL include the principal, commonly-used peptides as known substances in the catalog (resolved to codes and synonyms), so peptide-first users find them — handled as ordinary substances, without a peptide-specific subsystem.

Scenario 1 · A common peptide is already present

- **WHEN** a user searches for a common peptide (e.g. BPC-157, ipamorelin, semaglutide)
- **THEN** it is present as a known substance with its codes/synonyms and evidence link
- **AND** peptides are treated like any other substance, not a separate subsystem

§11 · BLOOD LAYER

Biomarkers & Lab Data — the Blood Layer (M2)

> The flagship of Phase 1. Analyte registry with mandatory coding, the deterministic import pipeline, manual entry of equal standing, three ranges per analyte, and completeness/freshness. Standalone value: "upload your past lab reports, get your clinical history in graphs, free." The architecture here is what makes open data possible later.

§11.1

Coded analyte registry

The system SHALL maintain a biomarker registry where every analyte carries a LOINC code and a UCUM unit, and every stored measurement additionally carries its method/assay and originating laboratory.

Scenario 1 · Coding is mandatory for a measurement

- **WHEN** a measurement is persisted
- **THEN** it has a UCUM unit and links to a LOINC-coded analyte
- **AND** without method/assay and lab it is flagged not aggregation-eligible

§11.2

Three ranges per analyte

The system SHALL store and display three distinct ranges — laboratory reference (method-dependent), longevity-optimal (with cited source), and safety threshold — and never collapse them.

Scenario 1 · A value near the optimal edge but within reference

- **WHEN** an ApoB value sits inside the lab reference range but outside the optimal range
- **THEN** the display shows it as reference-normal yet optimal-out
- **AND** the optimal range shows its source

§11.3

Optimal-range governance

The system SHALL attach to every longevity-optimal range its **source, reference population, endpoint, evidence grade and last-reviewed date**, and SHALL make them inspectable in one interaction from wherever the range is shown. An optimal range without this provenance SHALL NOT be displayed.

> RESOLVED
>
> RESOLVED (2026-08-19): reference range and safety threshold are relatively defined; "optimal" is not. Without visible derivation it is wellness numerology — and it is the most exposed number on the platform.

Scenario 1 · Inspecting an optimal range

- **WHEN** a user opens the optimal range shown for an analyte
- **THEN** they see its source, the population it derives from, the endpoint it optimises, its evidence grade and when it was last reviewed
- **AND** a range lacking any of these is not shown as optimal

Scenario 2 · Contested optimal range

- **WHEN** an optimal range is disputed
- **THEN** the dispute is visible on the range itself and handled as a public contested claim (see `community-and-social`)

§11.4

Deterministic pipeline with a hard boundary

The system SHALL make extraction, mapping, and threshold/index computation deterministic and reproducible (same input → same output); the LLM SHALL only produce the plain-language narration, generated once and cached.

Scenario 1 · Numbers are reproducible

- **WHEN** the same report is imported twice
- **THEN** the extracted values, mappings, and computed thresholds are identical
- **AND** any LLM narration is fetched from cache keyed on a fingerprint of the underlying data

§11.5

Import from PDF, photo, and CSV with learned templates

The system SHALL import lab data from PDF, photo, and CSV; on recognising a laboratory's format it SHALL generate a deterministic template so subsequent imports from that lab bypass the LLM.

Scenario 1 · Second import from a known lab

- **WHEN** a report matches a previously learned laboratory template
- **THEN** extraction runs deterministically without an LLM call

§11.6

Manual entry of equal standing

The system SHALL support manual measurement entry as a first-class path, with a sanity check on out-of-range values.

Scenario 1 · Typed value gets the same treatment

- **WHEN** a user types a value manually
- **THEN** it is stored with the same provenance fields and a plausibility check
- **AND** it is not treated as inferior to an imported value except where verification status matters (see `community-and-social`)

§11.7

Human review before persistence

The system SHALL require human confirmation before extracted values enter the system, and SHALL fail loudly on anything not understood.

Scenario 1 · Unrecognised line item

- **WHEN** the parser cannot interpret a line
- **THEN** it declares the item and asks, rather than proceeding on an incomplete report
- **AND** the review step shows recognised / not-recognised / inferred status per item

§11.8

Completeness and freshness indicators

The system SHALL show, per organ system, how complete and how fresh the data is.

Scenario 1 · Sparse metabolic panel

- **WHEN** a system has 6 of 13 expected markers, last updated 4 months ago
- **THEN** the UI states "Metabolic: 6/13, last updated 4 months ago"
- **AND** any index computed on sparse data is annotated as such (see `dashboards-and-doctor-view`)

§11.9

Biological clocks are first-class biomarkers

The system SHALL model biological (aging) clocks as biomarkers with role and provenance, covering both computed clocks and uploaded measured clocks, and SHALL surface them within biomarkers and dashboards.

Scenario 1 · A clock behaves like a biomarker

- **WHEN** a biological-clock value exists (computed or uploaded)
- **THEN** it is stored as a biomarker with method/provider, date, unit and uncertainty
- **AND** it appears in the biomarker registry and in dashboards (see `dashboards-and-doctor-view`)

§11.10

Deterministic PhenoAge computation from the standard panel

The system SHALL compute PhenoAge (Levine 2018) deterministically from its constituent blood analytes, and SHALL include those analytes in the standard blood panel so PhenoAge is computable from routine labs.

Scenario 1 · PhenoAge from routine bloods

- **WHEN** the standard blood panel is measured
- **THEN** the system computes PhenoAge deterministically from its constituent analytes (albumin, creatinine, glucose, CRP, lymphocyte %, MCV, RDW, ALP, WBC, and chronological age)
- **AND** the computation is reproducible (same input → same output), with the LLM never involved

§11.11

Upload of measured biological clocks (epigenetic, glycation)

The system SHALL accept uploaded biological-clock results — epigenetic clocks (e.g. TruAge and others) and glycation clocks (e.g. GlycanAge and others) — as measurements, with provider, method, date and uncertainty.

Scenario 1 · Epigenetic and glycation clocks uploaded

- **WHEN** a user uploads a TruAge epigenetic result or a GlycanAge glycation result
- **THEN** it is stored as a biological-clock biomarker with its provider, method and uncertainty
- **AND** provider/method are retained because clocks are not comparable across providers

§11.12

Inventory of biological clocks with accessibility assessment

The system SHALL maintain an inventory of biological clocks classified by accessibility — computable from existing data, easy to calculate, or purchasable from a reputable provider — to guide which clocks to add.

Scenario 1 · Assess a candidate clock

- **WHEN** a new biological clock is considered
- **THEN** it is recorded with its accessibility class and provider reputation
- **AND** clocks computable from existing data are prioritized over ones that must be purchased

§11.13

Optional PII obfuscation on import

The system SHALL offer PII obfuscation of imported documents as an option with a reviewable diff; by default the original file is retained **raw** (the user chooses to obfuscate).

Scenario 1 · User opts to obfuscate

- **WHEN** a user chooses to obfuscate a report containing name and ID
- **THEN** the PII is obfuscated and the user can see what was removed
- **AND** by default, without that choice, the original file is retained raw

§11.14

Original report file retained and public raw

The system SHALL retain the original lab report file and make it public by default in its **raw** original format, as part of OpenData.

Scenario 1 · Raw original file published

- **WHEN** a user's lab report is public
- **THEN** the raw original-format file (e.g. PDF) is available exactly as uploaded
- **AND** it is part of the OpenData public snapshot (see `analytics-and-open-data`)

§12 · PLANNING

Measurement Planning (M3 — Scheduling & Measurement Planner)

> The engine that turns active interventions into "what to measure, when, and why". Computes the panel, sets phase-dependent cadence, pools analytes into a single blood draw, and treats `Overdue` as a safety signal rather than a calendar nag.

§12.1

The panel is computed, not chosen

The system SHALL derive the required panel from the protocol as `Safety Core + Δ safety-rule markers of active substances + Δ efficacy markers of goals + optional deep-dives`.

Scenario 1 · Adding a nephrotoxic compound expands the panel

- **WHEN** a user adds a compound with a renal safety rule
- **THEN** the relevant renal markers are added to the computed panel with a stated reason
- **AND** the panel updates as interventions change

§12.2

Phase-dependent cadence

The system SHALL schedule denser measurement during initiation/titration and sparser during maintenance, per intervention phase.

Scenario 1 · Titration cadence

- **WHEN** an intervention is in titration
- **THEN** its safety markers are scheduled at a tighter cadence than in maintenance

§12.3

Analyte pooling into one draw

The system SHALL pool analytes due within a near window into a single blood draw (set-cover), respecting pre-analytic requirements (fasting, wash-outs, timing).

Scenario 1 · Fewer needles

- **WHEN** several analytes fall due within the same window
- **THEN** the planner proposes one draw covering them
- **AND** it honours conflicting pre-analytic constraints or splits when it must

§12.4

Overdue is a safety signal

The system SHALL compute `Overdue` for due measurements and SHALL express it in safety terms tied to the responsible intervention, not as a generic reminder.

Scenario 1 · Overdue framed against the drug

- **WHEN** a liver panel is overdue while a hepatically-relevant compound is active
- **THEN** the system states "rapamycin active for 94 days, liver panel overdue by 34"
- **AND** this links to the corresponding SafetyRule (see `safety-guardrails`)

Scenario 2 · Overdue pauses with the protocol

- **WHEN** the protocol is `Paused`
- **THEN** Overdue accrual for its measures pauses

§12.5

Planner views

The system SHALL provide calendar, year, and table views, filterable by type × status, with any Study timeline overlaid on the personal calendar.

Scenario 1 · Study overlay

- **WHEN** a user has an active N-of-1 Study
- **THEN** its timepoints appear on the personal calendar alongside routine measures

§13 · SAFETY

Safety Guardrails (M7)

> The module that justifies the project. Baseline gating before following a protocol, automatically inherited safety markers, dose plausibility, interaction and critical-value alerts, and escalation with a ready protocol sheet. Describes what to monitor; never prescribes what to take.

§13.1

Mandatory baseline AND risk acknowledgment before following a protocol

The system SHALL require BOTH (a) the required baseline safety markers to be recorded — by import or manual entry — AND (b) an explicit, logged risk acknowledgment, before a user can set a copied/followed protocol to `Active`. Neither substitutes for the other.

> RESOLVED
>
> RESOLVED (2026-08-04): baseline values are required *and* an acknowledgment is required; acknowledgment is never a bypass for missing baseline.

Scenario 1 · Beginner cannot start blind

- **WHEN** a beginner copies a protocol containing an off-label compound
- **THEN** the system blocks activation until the required baseline markers are recorded
- **AND** the user has also explicitly acknowledged the risk (recorded, timestamped)
- **AND** the required safety markers for that compound are already attached

Scenario 2 · Acknowledgment alone is not enough

- **WHEN** a user acknowledges the risk but has not recorded the required baseline
- **THEN** activation remains blocked
- **AND** the UI states which baseline markers are still missing

§13.2

Safety markers inherited automatically

The system SHALL attach a substance's associated safety markers to the user's measurement plan whenever that substance becomes active — safety is not opt-in.

Scenario 1 · Inheriting on copy

- **WHEN** a protocol is copied
- **THEN** its safety markers travel with it into the follower's plan (see `protocols`)

§13.3

Dose sanity check

The system SHALL compare an entered dose against tolerable upper limit, usual therapeutic dose, and maximum reported-in-literature dose, and flag it on three levels: outside-common-use / above-upper-limit / potentially-toxic. This is numeric plausibility validation, not clinical advice.

Scenario 1 · The extra zero

- **WHEN** a user enters vitamin D3 at 100,000 IU/day
- **THEN** the system flags it as potentially toxic and asks for confirmation
- **AND** the same applies to e.g. selenium 2000 mcg or rapamycin 30 mg/day

§13.4

Interaction and critical-value alerts

The system SHALL raise interaction warnings between active substances and alerts on critical out-of-range measured values, with declared confidence where interaction data is uncertain.

Scenario 1 · Interaction warned with confidence

- **WHEN** two active substances have a known interaction
- **THEN** the system warns and states the confidence/source of the interaction claim

§13.5

Escalation with a ready sheet

The system SHALL provide an explicit "consult a physician" escalation that produces the protocol sheet ready to hand over (see `dashboards-and-doctor-view`).

Scenario 1 · Critical value escalates

- **WHEN** a measured value crosses a safety threshold
- **THEN** the system escalates with the protocol sheet pre-generated
- **AND** it never issues a personalised dosage or diagnosis

§13.6

Heightened acknowledgment for research-level or animal-derived dosing

The system SHALL require an explicit, logged risk acknowledgment — accepting the risk knowingly — before activating a protocol that is `research` level or whose posology is `animal-derived`, in addition to the mandatory baseline.

Scenario 1 · Accepting the risk of a research protocol

- **WHEN** a user activates a research-level or animal-derived-dose protocol
- **THEN** the system requires an explicit acknowledgment naming the level and dose provenance (see `protocols`, `interventions-and-catalog`)
- **AND** this is on top of the mandatory baseline and safety-marker inheritance

§14 · ADHERENCE

Daily Log & Adherence (M13)

> The thirty-seconds-a-day surface that keeps the platform alive between blood draws, and the source of the adherence signal that makes an N-of-1 interpretable and weights a cohort contribution. Phase 1.

§14.1

Daily check-off with dose and slot pre-resolved

The system SHALL present a daily check-off of intakes with dose and time-slot already resolved from the plan.

Scenario 1 · One-tap logging

- **WHEN** a user opens the daily log
- **THEN** each due intake shows its resolved dose and slot
- **AND** the user marks it without re-entering dose

§14.2

Four-state adherence

The system SHALL record adherence in four states: taken-as-planned | partial-dose | intentionally-skipped | forgotten, and compute daily and cumulative adherence.

Scenario 1 · Intentional skip is distinct from forgotten

- **WHEN** a user marks a dose intentionally-skipped
- **THEN** it is stored distinctly from "forgotten"
- **AND** both feed the adherence percentage that qualifies results

§14.3

Wellness check

The system SHALL capture a daily subjective wellness check — mood 1–5, energy 1–5, sleep hours, sleep quality 1–5, free notes.

Scenario 1 · High-frequency series with zero cost

- **WHEN** a user submits a wellness check
- **THEN** it is stored as a high-frequency series usable in dashboards
- **AND** it requires no purchase or lab

§14.4

Adherence qualifies data downstream

The system SHALL make adherence available to analytics so a low-adherence contribution can be down-weighted or excluded.

Scenario 1 · 35% adherence flagged

- **WHEN** a contribution has 35% adherence
- **THEN** analytics can exclude or down-weight it versus a 95% contribution (see `analytics-and-open-data`)

§15 · DASHBOARDS

Dashboards & Doctor View (M5)

> The public dashboard and the physician-facing output. Organ-system decomposition, biological-age delta shown honestly, the doctor protocol sheet, and the double-column "literature vs observed" comparison that no competitor can build.

§15.1

Public dashboard grouped by biomarker family

The system SHALL present dashboards grouped by biomarker family (lipids, glucose, kidney, liver, inflammation, hormones, body composition, biological age), each stating why the group is monitored.

Scenario 1 · Group shows its rationale

- **WHEN** a viewer opens the kidney group
- **THEN** it shows the markers and the reason they are monitored (which active compound requires them)

§15.2

Organ-system view unifies efficacy and safety

The system SHALL provide an organ-system view where the same panel that shows efficacy also shows safety gaps.

Scenario 1 · Efficacy and safety in one tile

- **WHEN** the renal tile shows a trend
- **THEN** it can simultaneously show "you take a nephrotoxic compound and have not measured creatinine in 8 months"

§15.3

Headline metric first, then organ-system grid

The system SHALL lead the personal dashboard with a single headline metric — the biological-vs-chronological age delta — shown with its uncertainty interval and the clock and laboratory declared, followed by the organ-system grid. It SHALL NOT use this metric as a reputation or ranking metric.

> RESOLVED
>
> RESOLVED (2026-08-04): headline number first, then the organ-system grid.

Scenario 1 · Headline then grid

- **WHEN** the dashboard loads and a biological age is available
- **THEN** the biological-age delta is shown first as the headline, with its uncertainty interval, clock, and lab
- **AND** the organ-system grid follows below

Scenario 2 · Graceful fallback when the clock is missing

- **WHEN** no biological age is available
- **THEN** the headline falls back to a data-coverage summary (systems covered / freshness)
- **AND** the organ-system grid still renders in full

§15.4

Biological clocks surfaced as a group

The system SHALL present the available biological clocks — computed (PhenoAge) and uploaded (epigenetic, glycation) — as a dashboard group, each with its method/provider and uncertainty; the headline biological-age delta uses a declared default clock.

Scenario 1 · Multiple clocks shown honestly

- **WHEN** more than one biological clock is available
- **THEN** the dashboard shows each clock with its method/provider and uncertainty interval
- **AND** the headline names which clock and lab it uses, and clocks are not presented as comparable across providers

§15.5

Doctor protocol sheet

The system SHALL generate a "protocol sheet" for a physician: everything taken, doses, timing, since when, markers monitored with rationale, trends, and open alerts — as a public link and a PDF.

Scenario 1 · Handover artifact

- **WHEN** a user chooses "prepare for my doctor"
- **THEN** the system generates the sheet as link + PDF
- **AND** it is legible without a platform account

§15.6

Double-column literature vs observed

The system SHALL, per intervention, place side by side what the literature predicts (Evidence Layer) and what the biomarkers of people practising it show (cohort).

Scenario 1 · Prediction beside reality

- **WHEN** an intervention has both literature evidence and a cohort
- **THEN** the view shows "literature predicts X; the N people who did it show Y"
- **AND** the comparison uses the Outcome↔Biomarker link (see `evidence-layer`, `analytics-and-open-data`)

§15.7

Cohort comparison — inline reference AND dedicated view

The system SHALL surface cohort comparison in BOTH forms: (A) a lightweight inline reference on each dashboard tile, AND (B) a dedicated "vs cohort" view with distribution, percentile, stratification, sample size (n), and adherence weighting.

> RESOLVED
>
> RESOLVED (2026-08-04): both A and B.

Scenario 1 · Inline reference on the tile

- **WHEN** a biomarker's protocol has a cohort large enough to compare
- **THEN** its dashboard tile shows a lightweight "you vs cohort" reference (e.g. vs cohort median)
- **AND** the tile links into the dedicated "vs cohort" view

Scenario 2 · Dedicated view shows the distribution honestly

- **WHEN** the user opens the "vs cohort" view for a biomarker
- **THEN** it shows the cohort distribution with n, stratification, and adherence weighting
- **AND** it declares "not determinable" when the cohort is too small rather than implying precision

§15.8

Per-biomarker sex/age-normalized percentile

The system SHALL, for every biomarker measurement, report the percentile it falls in, normalized for sex and age against **published reference distributions**, alongside the reference / optimal / safety ranges, and SHALL declare the published source used.

Scenario 1 · Percentile shown per marker

- **WHEN** a biomarker value is displayed
- **THEN** it shows its sex/age-normalized percentile with the reference population declared
- **AND** this is distinct from the cohort comparison (population percentile vs same-protocol peers)

Scenario 2 · Insufficient reference data

- **WHEN** the reference population is insufficient to place a percentile for that sex/age
- **THEN** the system declares "not determinable" rather than showing a guessed percentile

§16 · COMMUNITY

Community & Social (M6)

> Following, copying, cohorts, the calculated Evidence Badge, public claim review, two-axis reputation, and data-quality-only gamification — attached to the pre-existing Rapamycin News community rather than launched cold.

§16.1

Follow and copy

The system SHALL let a user follow another user and copy/subscribe to a public protocol, with copying invoking the protocol fork + safety inheritance path (see `protocols`, `safety-guardrails`).

Scenario 1 · Copy enters the safety gate

- **WHEN** a user copies a public protocol
- **THEN** a fork is created and the mandatory baseline/safety gate is applied before activation

§16.2

One-click copyable catalog with an informed-decision panel

The system SHALL keep every intervention — including the most sensitive (peptides, plasmapheresis, IV) — in the public catalog, one-click copyable, with nothing documentable-but-not-copyable, and SHALL present an informed-decision panel at the point of copy.

RESOLVED

RESOLVED (2026-08-04): resolves spec v0.3 §11 Q7 — everything is copyable; the mitigation is informed choice plus the safety gate, not exclusion from the catalog.

Scenario 1 · Informed one-click copy

- **WHEN** a user chooses to copy any protocol, including a peptide / plasmapheresis / IV one
- **THEN** before activation an informed-decision panel surfaces: how many others practise it and their outcomes (adoption + Evidence Badge), the evidence corpus and grading, the medical/research level and dose provenance, the potential impacts, and Rapamycin News links for deeper discussion
- **AND** the copy then proceeds through the safety gate (baseline + heightened acknowledgment for research/animal-derived), with nothing withheld from the copyable catalog

§16.3

Cohorts form around protocols

The system SHALL group users practising the same protocol into a `Cohort` and let a follower compare against it.

Scenario 1 · Compare to the cohort

- **WHEN** a user practises a protocol with others
- **THEN** they can see their own trajectory against the cohort's (see `analytics-and-open-data`)

§16.4

Integrate with Rapamycin News (Discourse Connect where feasible)

The system SHALL integrate community with the existing Rapamycin News (Discourse) — via Discourse Connect SSO where feasible, falling back to bidirectional links otherwise — and SHALL treat community as an extension of it rather than a cold-launched social feature.

> RESOLVED
>
> RESOLVED (2026-08-04): Discourse Connect SSO ideally; else link-only. A form of partnership with Rapamycin News is expected.

Scenario 1 · SSO where feasible, link as fallback

- **WHEN** Discourse Connect SSO is available
- **THEN** identity is shared via Discourse Connect and protocol↔thread links are bidirectional
- **AND** where SSO is not available, bidirectional links are still established

Scenario 2 · Thread link on every protocol

- **WHEN** a protocol references a compound discussed on Rapamycin News
- **THEN** the protocol links bidirectionally to that thread

§16.5

Host study proposals for comment and consensus

The system SHALL host N-of-1 Study proposals (research question + endpoints) as community threads open for comment, and SHALL surface the proposal discussion alongside the Study.

Scenario 1 · Proposal collects community input

- **WHEN** a Study is pre-registered (see `studies-nof1`)
- **THEN** its question and endpoints appear as a community proposal thread open for comment
- **AND** the resulting discussion is visible from the Study, and the Study from the thread

§16.6

Calculated Evidence Badge

The system SHALL compute reputation from data the system already holds and show its derivation, never a self-declared claim.

Scenario 1 · Derivation, not claim

- **WHEN** a protocol has verified measurements
- **THEN** the badge reads e.g. "Rapamycin 6mg/week, 94 days, adherence 91%; ApoB 78→91 (+17%), ALT stable, lymphocytes −18%; 3 measurements, single lab, verified reports"
- **AND** no free-form success claim is shown in its place

§16.7

Public claim review

The system SHALL make contested claims public, discussed, and traceable — not a private flag to a team.

Scenario 1 · Contested claim is visible

- **WHEN** users contest a claim
- **THEN** the item shows "contested by N people, here's why" publicly

§16.8

Two-axis reputation and specialty discovery

The system SHALL keep verified-credential and data-verification as separate axes and provide discovery by specialty, not popularity.

Scenario 1 · Discover by specialty

- **WHEN** a user browses contributors
- **THEN** they can filter by verified specialty distinct from data-quality

§16.9

Gamification only on data quality

The system SHALL gamify only data quality — adherence streaks, panel completeness, documentation continuity — and SHALL NOT rank users by health outcomes.

Scenario 1 · No outcome leaderboard

- **WHEN** gamification is displayed
- **THEN** it rewards logging/completeness
- **AND** there is no "lowest ApoB" or "best biological age" ranking

§17 · STUDIES

N-of-1 Studies

> Makes a self-experiment interpretable by forcing the question and the endpoint to be declared before it starts. The container that later composes into distributed, grassroots trials. Pre-registration costs nothing and disarms the main methodological objection in advance.

§17.1

Study with pre-declared endpoints

The system SHALL model a `Study` with a research question and endpoints declared before T0, an associated protocol, timepoints (T0 baseline, T1…Tn), a per-timepoint test battery, duration, wash-out, stop criteria, and a pre/post statistical comparison.

Scenario 1 · Endpoint declared before start

- **WHEN** a user creates a Study
- **THEN** the endpoint must be declared before the study can begin
- **AND** it is timestamped so later results cannot silently redefine success

§17.2

Timepoints drive the planner

The system SHALL feed a Study's timepoints and test battery into the measurement planner and overlay them on the personal calendar.

Scenario 1 · Battery scheduled

- **WHEN** a Study defines T1 at week 8 with a battery
- **THEN** those measures appear as due in the planner at week 8

§17.3

Pre/post read-out

The system SHALL compute the declared pre/post comparison at completion and present it against the pre-registered endpoint.

Scenario 1 · Honest read-out

- **WHEN** a Study completes
- **THEN** the system reports the endpoint result as declared, including a null result
- **AND** insufficient data yields "not determinable" rather than a guessed value

§17.4

Study question and endpoints are a community proposal

The system SHALL require a Study's research question and endpoints to be published as a **community proposal** — a forum thread open for comments (Rapamycin News / Discourse) — as the act of pre-registration, so that pre-registration serves consensus rather than a private declaration. The questions a Study seeks to answer are themselves subject to community proposal.

RESOLVED

RESOLVED (2026-08-04): pre-registration is **mandatory** and takes the form of a community proposal for comment/consensus; the research questions are themselves subject to community proposal.

Scenario 1 · Pre-registration posts a proposal

- **WHEN** a user pre-registers a Study
- **THEN** the system publishes the question and endpoints as a community proposal thread open for comments
- **AND** the Study links bidirectionally to its proposal thread (see `community-and-social`)

Scenario 2 · Consensus sought, endpoints frozen

- **WHEN** a Study's proposal is open
- **THEN** comments are collected against the declared question and endpoints
- **AND** the pre-registered endpoints are immutable — changing them requires a new proposal/version, never a silent edit

§18 · PROCUREMENT

Procurement & Inventory (M4)

> Turns a protocol into a supply plan — projected need, batch purchase, inventory, lots and expiry — and links to Pills Management for restock. Daily intake scheduling lives in `pills-management`. Phase 3. No commercial integration with suppliers in the initial phase — the main conflict-of-interest surface to keep clean.

§18.1

Projected need from protocol

The system SHALL project consumption from the active protocol and derive batch purchase suggestions with lead time.

Scenario 1 · Reorder before stock-out

- **WHEN** projected consumption will exhaust a stock before its lead time elapses
- **THEN** the system raises a reorder alert in time

§18.2

Inventory with lots and expiry

The system SHALL track stock, lots, and expiry, and alert on expiring stock.

Scenario 1 · Expiry alert

- **WHEN** a lot approaches expiry
- **THEN** the system alerts and reflects it in the projected need

§18.3

No monetised procurement initially

The system SHALL NOT monetise procurement or embed supplier affiliations in the initial phase.

Scenario 1 · Neutral sourcing

- **WHEN** sourcing information is shown
- **THEN** it carries no paid placement or affiliate incentive

§19 · PILLS

Pills Management (M4 — daily intake orchestration)

> Orchestrates the daily intake of pills and supplements: which to take, when, how to distribute them across the day, spacing between them, whether with or away from meals, and what a substance needs for absorption (e.g. dietary fat). Links to procurement for restock. This is the daily-intake half split out of the original M4; purchase and inventory live in `procurement-and-inventory`, and the schedule surfaces in the daily log and the programming calendar.

§19.1

Daily intake schedule across slots

The system SHALL allocate intakes into day slots optimising bioavailability and avoiding interference and toxicity summation (mineral separation, competing compounds, hepatic load), using a constraint solver rather than an LLM.

Scenario 1 · Competing compounds separated

- **WHEN** two compounds should not be co-administered
- **THEN** the schedule places them in separated slots with a stated reason
- **AND** the allocation is deterministic and explainable

§19.2

Meal timing and absorption requirements

The system SHALL model each substance's meal constraints — with-meal, away-from-meals, and any co-ingested substance required for absorption (e.g. dietary fat for fat-soluble compounds) — and reflect them in the schedule.

Scenario 1 · Fat-soluble compound with a fatty meal

- **WHEN** a fat-soluble compound is scheduled
- **THEN** it is placed with a meal containing sufficient fat
- **AND** a compound requiring an empty stomach is placed away from meals

§19.3

Spacing and distance rules

The system SHALL honour required spacing/distance between compounds (competition, absorption windows) and pre/post-meal distances, or flag an unresolvable conflict.

Scenario 1 · Distance respected or flagged

- **WHEN** a compound must be taken a set distance from another or from food
- **THEN** the schedule enforces that distance
- **AND** it flags the conflict when the constraints cannot all be met

§19.4

Restock link to procurement

The system SHALL link pill scheduling to procurement so that low stock triggers a restock action (see `procurement-and-inventory`).

Scenario 1 · Running low triggers reorder

- **WHEN** scheduled consumption will exhaust a compound's stock within its lead time
- **THEN** the system raises a restock action via procurement

§19.5

Scale to high-volume regimens

The system SHALL remain usable and legible for high-volume regimens — dozens to well over a hundred pills per day (e.g. 60–130) — grouping intakes and keeping check-off fast.

Scenario 1 · 130 pills a day stays manageable

- **WHEN** a user's regimen involves 60–130 pills per day
- **THEN** the schedule groups them by slot/meal for legibility and fast check-off
- **AND** performance and the daily log remain usable at that scale

§19.6

Feeds the daily log and calendar

The system SHALL surface the resolved pill schedule in the daily log and the programming calendar.

Scenario 1 · Schedule appears where the user acts

- **WHEN** the daily schedule is computed
- **THEN** each intake appears in the daily log with dose and slot (see `daily-log-and-adherence`)
- **AND** on the daily/weekly calendar alongside therapies and exercise (see `measurement-planning`)

§20 · THERAPEUTICS

Therapeutics Management (periodic treatments)

> Manages interventions that require a periodic treatment session rather than a pill — red light / photobiomodulation, HBOT (incl. hypoxia-hyperoxia), IHHT (intermittent hypoxia-hyperoxia training), sauna, cold exposure / cold plunge, whole-body cryotherapy, therapeutic plasma exchange (TPE / plasmapheresis), IV infusions (e.g. NAD+, vitamins), ozone / EBOO, PEMF, whole-body vibration, vagus nerve stimulation (tVNS), and others to add — capturing their **typed parameters**, scheduling their sessions, and logging done/not-done. Builds on the Device/Therapy and Procedure intervention subtypes (see `domain-model`).

§20.1

Model periodic therapy sessions

The system SHALL model periodic therapy interventions — red light / photobiomodulation, HBOT (incl. hypoxia-hyperoxia), IHHT (intermittent hypoxia-hyperoxia training), sauna, cold exposure / cold plunge, whole-body cryotherapy, therapeutic plasma exchange (TPE / plasmapheresis), IV infusions (e.g. NAD+, vitamins), ozone / EBOO, PEMF, whole-body vibration, vagus nerve stimulation (tVNS) — with their parameters and session cadence, and the catalog SHALL be extensible.

Scenario 1 · HBOT protocol captured

- **WHEN** a user configures an HBOT hypoxia-hyperoxia protocol
- **THEN** its parameters and session cadence are stored
- **AND** red light, cryotherapy, cold plunge, plasmapheresis and the others are represented the same way

§20.2

Typed parameters per therapy

The system SHALL type each therapy with its own characteristic parameters — for example red light (wavelength(s), irradiance, duration, distance, body area), HBOT (pressure in ATA, FiO2 profile, duration), IHHT (FiO2 high/low cycle, cycle count, session duration), cryotherapy / cold plunge (temperature, duration), sauna (temperature, humidity, duration), TPE / plasmapheresis (volume exchanged, replacement fluid, frequency), IV infusion (agent, dose, rate), whole-body vibration (frequency, amplitude, duration), tVNS (intensity, site, duration) — so each exposes only what is meaningful for it.

Scenario 1 · Parameters differ by therapy

- **WHEN** two different therapies are configured
- **THEN** each exposes only its characteristic parameters
- **AND** a cold plunge's water temperature is never conflated with a red-light wavelength

§20.3

Extensible therapy catalog

The system SHALL allow new therapy types to be added and configured with parameters and cadence, without changing the model.

Scenario 1 · Add a new therapy

- **WHEN** a new periodic therapy is introduced
- **THEN** it can be configured with its parameters and cadence like existing ones

§20.4

Schedule therapy sessions on the calendar

The system SHALL place therapy sessions on the daily/weekly programming calendar and remind the user.

Scenario 1 · Session scheduled and reminded

- **WHEN** a therapy has a due session
- **THEN** it appears on the calendar (see `measurement-planning`) with a reminder

§20.5

Log sessions done or not

The system SHALL log therapy sessions with four-state adherence (done / partial / skipped / missed), feeding adherence.

Scenario 1 · Session logged

- **WHEN** a user completes or misses a therapy session
- **THEN** it is logged with its adherence state (see `daily-log-and-adherence`)

§21 · EXERCISE

Exercise Reporting

> Lets a user define what their physical activity is and log whether they did it — **not a gym-management app**. Covers strength training, cardio, mobility, whether HIIT is performed, and VO2max. Its job is to make the activity explicit, log done/not-done, and place it in the daily/weekly programming alongside pills and therapeutics.

§21.1

Define the exercise activity profile

The system SHALL let a user declare their exercise activity — strength training, cardio, mobility, and whether HIIT is performed — at the level of what they do, not per-set gym tracking.

Scenario 1 · Activity declared, not micro-managed

- **WHEN** a user sets up their exercise profile
- **THEN** they declare strength, cardio, mobility and HIIT yes/no at an activity level
- **AND** the system does not require per-set/per-rep gym logging

§21.2

VO2max captured as a biomarker

The system SHALL capture VO2max as a measured biomarker with date and method.

Scenario 1 · VO2max recorded

- **WHEN** a user enters or uploads a VO2max value
- **THEN** it is stored as a biomarker (see `biomarkers-and-labs`) and shown in dashboards

§21.3

Log done or not-done

The system SHALL log whether a planned exercise activity was done, with four-state adherence, without becoming a gym tracker.

Scenario 1 · Weekly strength session logged

- **WHEN** a planned strength session is due
- **THEN** the user marks it done / partial / skipped / missed (see `daily-log-and-adherence`)
- **AND** no set-by-set detail is required

§21.4

Appears in the programming calendar

The system SHALL place exercise activities in the daily/weekly programming calendar alongside pills and therapeutics.

Scenario 1 · Exercise on the weekly plan

- **WHEN** the weekly plan is shown
- **THEN** exercise activities appear alongside pill intakes and therapy sessions (see `measurement-planning`)

§22 · NUTRITION

Nutrition Management

> Defines the user's dietary regime so it can be represented, scheduled and logged alongside the rest of the protocol: a computed **TDEE**, the typical eating **schedule/pattern** they intend to follow (OMAD, calorie restriction, intermittent fasting, or conventional breakfast-lunch-dinner), and the typical **macronutrient split** (carbohydrate / protein / fat). Builds on the Nutrition/Fasting intervention subtype (see `domain-model`) and connects to meal-timed intakes in `pills-management`.

§22.1

Compute TDEE

The system SHALL compute the user's Total Daily Energy Expenditure deterministically from their biological profile and activity level, and SHALL show the formula used.

Scenario 1 · TDEE from profile and activity

- **WHEN** the user's biological profile (see `accounts-and-profiles`) and activity level (see `exercise-reporting`) are known
- **THEN** the system computes TDEE deterministically with a declared formula
- **AND** it recomputes when the inputs change

§22.2

Declare the typical eating schedule/pattern

The system SHALL let the user declare their typical intended eating pattern — OMAD, calorie restriction, intermittent fasting (e.g. 16/8), or conventional breakfast-lunch-dinner — together with the eating window.

Scenario 1 · Fasting pattern declared and scheduled

- **WHEN** a user follows 16/8 intermittent fasting
- **THEN** the pattern and eating window are recorded
- **AND** they are reflected in the daily/weekly calendar and inform meal-timed intakes (see `measurement-planning`, `pills-management`)

§22.3

Typical macronutrient composition

The system SHALL capture the typical macronutrient split — carbohydrate, protein, fat — the user intends to follow.

Scenario 1 · Macro split recorded

- **WHEN** a user sets their typical macros
- **THEN** the carbs/protein/fat composition is stored and shown on the profile

§22.4

Log adherence to the nutrition plan

The system SHALL log adherence to the declared nutrition pattern (followed / deviated) with exceptions, feeding overall adherence.

Scenario 1 · Fasting exception logged

- **WHEN** a user breaks their fasting window
- **THEN** the deviation is logged as an exception (see `daily-log-and-adherence`)

§23 · EVIDENCE

Evidence Layer (M11)

> Grades literature into structured, comparable claims and — crucially — links each outcome to a biomarker, which is what lets the platform compare prediction with observation. Bootstrapped on the longevity subset first.

§23.1

Evipedia as the primary evidence source

The system SHALL use Forever Healthy's Evipedia knowledgebase as its primary source of intervention evidence reviews (via `ai-uses-and-attribution`), and SHALL run its own AI4L-audited ingestion only for interventions Evipedia does not cover.

Scenario 1 · Base on Evipedia, extend where absent

- **WHEN** evidence for an intervention is needed
- **THEN** the platform first uses the Evipedia review (conclusion, grading, PMIDs), attributed to Forever Healthy
- **AND** it runs its own AI4L-audited ingestion only where Evipedia has no coverage

§23.2

Structured evidence claims

The system SHALL ingest literature (Europe PMC / PubMed / preprints) and extract, in batch, `EvidenceClaim` records of intervention × outcome × study.

Scenario 1 · Batch extraction offline

- **WHEN** literature is ingested
- **THEN** claims are extracted in batch offline, not in an interactive request

§23.3

Confidence separate from effect size

The system SHALL grade confidence with GRADE-inspired weighted factors and report effect size separately, in the outcome's native units.

Scenario 1 · No percentages from absolutes

- **WHEN** an effect size is stored
- **THEN** it is in native units, not a percentage derived from absolute values
- **AND** confidence is a separate field from effect size

§23.4

Directional evidence with study counts

The system SHALL record directional evidence with explicit counts of favourable, null, and unfavourable studies.

Scenario 1 · The nulls are counted

- **WHEN** evidence for a compound is summarised
- **THEN** favourable/null/unfavourable counts are all shown

§23.5

Outcome↔Biomarker link

The system SHALL map outcomes to biomarkers so literature predictions can be compared with observed cohort biomarkers.

Scenario 1 · Enables the double column

- **WHEN** an outcome maps to a biomarker
- **THEN** the dashboard double-column can align literature and cohort (see `dashboards-and-doctor-view`)

§23.6

Alerts on evidence change

The system SHALL alert users when evidence for a compound they use materially changes.

Scenario 1 · New unfavourable evidence

- **WHEN** significant new evidence lands for an active compound
- **THEN** affected users are alerted

§24 · OPEN DATA

Analytics & Open Data (M8)

> Aggregation across cohorts, honest cross-person comparison, stratification, and the open-data discipline that decides whether researchers cite the dataset or ignore it. OpenData ships as **two products**: a **full clonable public snapshot** (public profiles and their public data, protocols, treatments, measurements) that lets anyone self-host an equivalent instance, and an **aggregated research export** (lab-relative z-scores, cohort thresholds, OMOP). Phase 4.

§24.1

Defined OpenData scope

The system SHALL define OpenData as everything a user has made public: public user profiles and their public data, defined protocols, treatments/interventions, measurements, **raw original lab report files**, and **public genomic data including the gVCF**. Only withheld per-item data is excluded.

Scenario 1 · What is and isn't OpenData

- **WHEN** the OpenData set is assembled
- **THEN** it includes public profiles and their public data, protocols, treatments, measurements, original lab files, and public genomics (incl. gVCF)
- **AND** it excludes only withheld items (see `genomics`, `biomarkers-and-labs`)

§24.2

Full clonable public snapshot

The system SHALL publish OpenData as a full, clonable snapshot sufficient to seed an independent self-hosted instance, in addition to the aggregated research export.

Scenario 1 · Clone the whole public dataset

- **WHEN** a third party downloads the public snapshot
- **THEN** it contains the public profiles, protocols, treatments, measurements, original lab files and public genomics as published
- **AND** it is sufficient to stand up an equivalent instance (see `accounts-and-profiles`)

Scenario 2 · One boundary for both products

- **WHEN** either the public snapshot or the aggregated research export is produced
- **THEN** neither includes any withheld item (public genomics, gVCF and original lab files are included, being public)

§24.3

Cohort aggregation

The system SHALL aggregate outcomes across a cohort practising the same protocol, weighting contributions by adherence and data completeness.

Scenario 1 · Adherence-weighted aggregate

- **WHEN** a cohort aggregate is computed
- **THEN** low-adherence and sparse contributions are down-weighted or excluded by a declared rule

§24.4

Compare on z-score, not raw value

The system SHALL compare people using the z-score relative to the originating laboratory's range, not the raw value.

Scenario 1 · Cross-lab comparison

- **WHEN** two users measured the same analyte at different labs
- **THEN** comparison uses lab-relative z-scores

§24.5

Aggregated research export uses only well-coded data

The system SHALL include in the **aggregated research export** only measurements carrying LOINC + UCUM + declared provenance, and only above minimum cohort thresholds guarding re-identification. (The full public snapshot, by contrast, carries public data as published.)

Scenario 1 · Small clean over large dirty

- **WHEN** the aggregated research export runs
- **THEN** it excludes uncoded/unprovenanced measurements
- **AND** it suppresses cohorts too small to protect identity

§24.6

Researcher access and OMOP export

The system SHALL provide a researcher endpoint (dump + API) and an OMOP CDM export for the research layer.

Scenario 1 · OMOP export

- **WHEN** the research export runs
- **THEN** it produces an OMOP-CDM dataset consumable by OHDSI tooling

§24.7

Sex/age-normalized percentile computation

The system SHALL compute, per biomarker, a percentile normalized for sex and age against **published reference distributions**, and SHALL expose it to dashboards and the doctor view (see `dashboards-and-doctor-view`), declaring the published source used per marker.

> RESOLVED
>
> RESOLVED (2026-08-04): the percentile is computed against **published reference distributions**, not the platform's own aggregate; the source is declared per marker.

Scenario 1 · Percentile with declared reference

- **WHEN** a percentile is computed for a marker
- **THEN** it is normalized for the user's sex and age band
- **AND** the reference population and its source are declared with the value

Scenario 2 · Distinct from cohort comparison

- **WHEN** both a population percentile and a same-protocol cohort comparison exist for a marker
- **THEN** they are computed and presented as two distinct references (population vs peers-on-this-protocol)

§24.8

Raw-data publication validated by a community poll at BETA

The platform's policy of publishing raw data (raw lab reports, gVCF, and all public data) SHALL be put to a community poll on Rapamycin News when the public BETA testing call opens, to gauge participant opinion, and the outcome SHALL inform the raw-publication policy.

Scenario 1 · BETA community poll

- **WHEN** the public BETA testing call opens
- **THEN** a poll on Rapamycin News asks participants their view on publishing raw data of everything
- **AND** the outcome informs the raw-publication policy (see `community-and-social`)

§24.9

"Stops instead of guessing" in aggregates

The system SHALL declare insufficiency rather than present an approximate synthetic value when coverage is inadequate.

Scenario 1 · Not enough coverage

- **WHEN** an organ-system index or projection lacks coverage
- **THEN** the system says "not determinable" instead of an approximate number

§25 · GENOMICS

Genomics (M9 — minimal interpretation, public by default)

> No heavy re-analysis pipeline: the platform imports selected interpreted variants and, where provided, retains the raw **gVCF**, focused on the actionable. Like all other data on biohack.it, genomic data — including the gVCF — is **public by default** and part of OpenData, behind a heightened, explicit consent step because it is maximally identifying.

§25.1

Accept genomic data including gVCF

The system SHALL accept genomic data — interpreted variants/reports and, where provided, the raw gVCF — store it, and treat it as publishable data.

Scenario 1 · gVCF stored and publishable

- **WHEN** a user provides a gVCF
- **THEN** the system stores it and treats it as publishable, public-by-default data
- **AND** deep re-analysis (variant calling from raw reads) is out of scope

§25.2

Actionable interpretation focus

The system SHALL surface genomic interpretation limited to the actionable — pharmacogenomics (CYP metaboliser status; APOE) — while retaining the raw gVCF for sharing and export.

Scenario 1 · Actionable variant surfaced, raw retained

- **WHEN** an actionable pharmacogenomic variant is present
- **THEN** it can inform safety context
- **AND** non-actionable interpretation stays out of scope even though the raw gVCF is retained and shareable

§25.3

Public by default with heightened consent (incl. gVCF)

The system SHALL make genomic data — including the gVCF — public by default like other data, gated by an explicit heightened-consent step stating its identifiability, the exposure of biological relatives, and the effective irreversibility of publication.

> DECISION
>
> DECISION (2026-08-04): reverses spec v0.3 Principle #9 — genomics is public by default, gVCF included. FLAG for Fabio: genomic data is maximally identifying (of the user and their biological relatives), is a GDPR Art. 9 special category, and public release is effectively irreversible. The heightened-consent gate is the mitigation — confirm it is sufficient.

Scenario 1 · Heightened consent before genomics goes public

- **WHEN** a user's genomic data (incl. gVCF) is set public
- **THEN** the system requires an explicit heightened-consent acknowledgment covering identifiability, relatives, and irreversibility
- **AND** once consented, the data is part of OpenData like other public data (see `analytics-and-open-data`)

§26 · AGENT ACCESS

Agent Access (MCP)

> Lets a biohacker query their own data with their own agent through a revocable, read-only token and a lightweight MCP gateway. Low cost, high cultural impact for the hacker audience, and something no commercial competitor will ship.

§26.1

Revocable read-only token

The system SHALL issue per-profile, read-only, revocable tokens for agent access, and SHALL never grant write access through this path.

Scenario 1 · Token queried then revoked

- **WHEN** a user issues an agent token and later revokes it
- **THEN** queries succeed before revocation and fail immediately after
- **AND** the token can never mutate data

§26.2

MCP gateway over encrypted context

The system SHALL expose a lightweight MCP endpoint returning the profile's data context, queryable by external agents (e.g. Claude, Cursor, a Nostr bot).

Scenario 1 · External agent reads biomarkers

- **WHEN** an authorised agent queries via MCP
- **THEN** it receives the read-only biomarker/protocol context for that profile

§26.3

Agent access respects per-item visibility

The system SHALL return through agent access only data consistent with the profile's visibility settings — withheld items are never returned; public data (including public genomics, where the user made it public) follows those settings. There is no special genomics carve-out.

Scenario 1 · Withheld items not exposed

- **WHEN** an agent queries a profile
- **THEN** withheld items are not returned
- **AND** public data follows the profile's visibility settings