FINDINGS

Privacy Policy

UNIT01 d.o.o.
Vodovodska 75, 11030 Belgrade, Serbia
hello@unit01.dev

Version: 8 July 2026

Canonical web version:
https://getfindings.com/privacy

1. Introduction

This Privacy Policy explains how UNIT01 d.o.o. (Vodovodska 75, 11030 Belgrade, Serbia) ("we", "us", "our") collects, uses, stores, shares, and protects personal data when you use the Findings health intelligence platform — including our website, web application, mobile applications, and related APIs (together, the "Service").

Findings helps you upload laboratory reports, track biomarkers over time, log daily health readings, manage medications, and receive AI-generated explanations and reminders. Because the Service processes health-related information, we treat privacy and security as core product requirements, not afterthoughts.

We are the data controller for personal data described in this policy. Data protection contact: Rade Joksimovic, UNIT01 d.o.o., Vodovodska 75, 11030 Belgrade, Serbia. Email: hello@unit01.dev.

This policy should be read together with our Terms of Service (https://getfindings.com/terms). If you do not agree with this policy, please do not use the Service.

2. Scope and roles

This policy applies to:

- Visitors to our public website and marketing pages
- Registered users who create an account and use Findings
- Individuals whose data is entered by a user (e.g. you should not upload another person's lab report without their permission)

We do not knowingly offer the Service to children. See Section 15. We may update subprocessors and technical providers over time; material changes are described in Section 18.

3. Categories of personal data we collect

Depending on how you use the Service, we may process the categories below. Not every user provides every category.

3.1 Account and identity data

- Email address and authentication credentials (managed via Supabase Auth — we do not store plaintext passwords)
- Account identifiers (internal user UUID)
- Optional display name and avatar URL

3.2 Profile and demographic data

- Date of birth (used to derive age for range calculations)
- Biological sex (for sex-specific biomarker reference ranges where clinically relevant)
- Country of residence (ISO code)
- Timezone, language, and unit preferences (metric/imperial, conventional/SI lab units, per-biomarker overrides)
- Chronic conditions, surgical history, and family history you select or enter
- Onboarding completion timestamps

3.3 Health and wellness data (special category)

The following are special category data under Article 9 of the UK/EU GDPR because they relate to your health:

- Laboratory report files you upload (PDF or image) and metadata (filename, file type, detected lab name, panel date, processing status)
- Biomarker values extracted by OCR, including raw text from your report, parsed numeric values, units, reference ranges, assay method where detected, and OCR confidence scores
- Confirmed biomarker entries in your longitudinal history (values, units, canonical normalised values, collection dates, notes)
- Daily tracker definitions and readings (e.g. blood pressure, blood glucose, SpO ,, weight, and custom trackers) including source (manual entry; Apple Health / Google Health Connect when enabled)
- Medications and supplements (name, dose, frequency, start/end dates, active status, instructions)
- AI-generated drug–biomarker interaction summaries associated with your medications
- Retest reminders and dismissal actions
- AI chat messages you send and assistant responses you receive
- Conversation titles and conversation identifiers
- "AI Health Context" memories (short factual statements about medications, conditions, lifestyle, goals, or context) that you add manually or that we extract from chat, including vector embeddings of those memories for semantic retrieval
- Cached AI overview summaries shown on specific screens (dashboard, biomarkers, trackers, upload flows)
- Notification preferences and delivery logs (push and email where enabled)

3.4 Subscription, billing, and gift purchase data

- Subscription status, plan type, platform (web, iOS, Android), product identifiers, billing period dates, trial status, and cancellation flags
- RevenueCat customer identifiers (we do not store full payment card numbers — payment processing is handled by Apple, Google, or our payment partners via RevenueCat)
- Gift purchase details when you buy or redeem a gift membership: purchaser and recipient names, optional recipient email, optional personal message, delivery preference, gift code, redemption status, membership length purchased, and related transaction identifiers from our billing partners

3.5 Technical and usage data

- Device push notification tokens (Expo push token) when you enable notifications
- IP address, request timestamps, API access logs, and security-related events
- Browser or app type, and basic diagnostic information needed to operate and secure the Service
- Cookies and similar technologies used for authentication sessions (see Section 12)
- Optional campaign measurement identifiers when your preference allows (hashed email, click/browser IDs such as gclid/fbp/fbc, coarse conversion events) — not lab or health profile content
- Privacy choice preferences (analytics and campaign measurement toggles, consent region) synced to your account when you are signed in

3.6 Data we do not intentionally collect

- Government-issued ID numbers, unless you voluntarily include them in an uploaded document
- Emergency contact details, unless you enter them in chat or profile fields
- Precise geolocation tracking (we use country for regional settings, not continuous GPS)

4. How we collect data

- Directly from you — registration, profile forms, lab uploads, tracker logs, medication entries, chat messages, memory edits, settings, subscription actions, and gift purchase forms
- Automatically — when you use the Service (e.g. OCR processing, trend calculations, retest schedule computation, correlation analysis, AI context assembly)
- From third parties — authentication session data from Supabase; subscription events from RevenueCat webhooks; optional email delivery via Resend; push delivery via Expo; AI inference from Anthropic and OpenAI as described below
- Health platform sync (planned) — if you connect Apple Health or Google Health Connect, we will import only the data types you authorise for the trackers you enable

5. Purposes and legal bases for processing

Where the GDPR or UK GDPR applies, we rely on the legal bases below. For health data (Article 9), we rely primarily on your explicit consent given when you create an account, upload health data, or enable health-related features. You may withdraw consent by deleting your account or specific data where the product allows; withdrawal does not affect the lawfulness of processing before withdrawal.

Purpose	Examples	Legal basis (GDPR)
Provide the Service	Account creation, biomarker storage, charts, reminders, medications module	Performance of contract (Art. 6(1)(b)); explicit consent for health data (Art. 9(2)(a))
OCR and data normalisation	Extract biomarkers from your lab PDF/image; match to our catalogue; convert units	Explicit consent (Art. 9(2)(a)); contract
AI features	Chat coach, dashboard summaries, scoped overviews, memory extraction, embeddings, drug interaction text	Explicit consent (Art. 9(2)(a)); contract
Subscriptions	Verify access, process webhooks, gift purchase and redemption, enforce preview vs paid features	Contract; legitimate interests (fraud prevention)
Notifications	Retest reminders, upload status alerts (if you opt in)	Consent / contract (depending on channel and jurisdiction)
Security & abuse prevention	Authentication, rate limits, prompt-injection screening, audit logs	Legitimate interests (Art. 6(1)(f)); legal obligation where applicable
Legal compliance	Respond to lawful requests, tax/accounting where required	Legal obligation (Art. 6(1)(c))
Improve reliability	Aggregated, de-identified analytics on errors and performance (not sale of health data)	Legitimate interests

Campaign measurement (optional)	Measure whether marketing led to account creation or membership when you opt in; browser tags (Google Tag Manager, Meta Pixel !' Stape server-side host) and matching server-side conversion forwarding to Google/Meta only when your campaign measurement preference allows	Consent where required (EEA/UK/CH); implicit allow with opt-out elsewhere until you save preferences

We do not sell your personal data. We do not use your health data for advertising or ad platform targeting. Optional campaign measurement subprocessors (Meta, Google, and our server-side tag hosting provider) receive only hashed contact information and coarse conversion events when your campaign measurement preference allows — never laboratory results, biomarker values, health profile fields, or other special category health data. Browser tags and server-side conversion forwarding use the same preference (synced from your device when you are signed in). We do not use your health data to train third-party foundation models. AI providers process your data as subprocessors solely to deliver features you request, under our instructions and contractual safeguards.

6. Artificial intelligence and automated processing

Findings uses AI extensively. This section describes what is automated, what is not, and how your data flows through AI systems.

6.1 What AI does

- OCR and structured extraction from lab reports (Anthropic Claude vision)
- Conversational health coach responses (streaming and non-streaming chat)
- Short dashboard and scoped "overview" summaries cached per screen
- Automatic conversation title generation from your first message
- Memory extraction after chat turns (ADD / UPDATE / DELETE factual statements)
- Semantic memory retrieval using embeddings (OpenAI text-embedding-3-small) stored in our database
- Drug–biomarker interaction summaries when you add a medication
- Safety screening on prompts and responses (e.g. PII detection, prompt-injection mitigation via llmverify)

6.2 What AI does not do

- AI output is not a medical diagnosis, prescription, or emergency triage system
- We do not make solely automated decisions with legal or similarly significant effects about you (e.g. insurance eligibility)
- AI suggestions about retesting are educational scheduling aids — you and your clinician remain responsible for clinical decisions

6.3 Context sent to AI providers

When you use AI features, we assemble a context bundle that may include: profile demographics and health goals; active medications; relevant memories (hybrid semantic + recency retrieval, capped); recent biomarker values and trends; daily tracker readings (typically last 14 days); cross-marker correlation summaries; and your current message. This bundle is transmitted to our AI subprocessors to generate a response.

Chat message text is stored in our database so you can view history. We do not routinely store the full context bundle alongside each message in production, though our schema may support optional technical snapshots for debugging when strictly necessary.

6.4 Memory system

After chat turns, a background process may extract durable facts (e.g. "takes metformin 1000mg") into your AI Health Context store. You can view, edit, or delete these in Profile !' Health Context. Profile fields such as age, sex, country, and chronic conditions are read from your profile directly and are not duplicated in memory. Embeddings are regenerated when memory text changes.

6.5 Rate limits and logs

AI chat is rate-limited (currently 30 messages per hour per user) to prevent abuse. API and security logs may record timestamps, user IDs, and error metadata but are not used to train public AI models.

7. Lab report files and retention

When you upload a lab report, the file is temporarily stored in encrypted object storage (Supabase Storage) while our OCR pipeline processes it. After successful extraction:

- The raw file is removed from storage
- The storage path on your upload record is cleared
- We retain structured extraction rows (biomarker names, values, units, confidence, method) and your confirmed entries in your health history

If processing fails on the final attempt, we still remove the raw file. Uploads that never finish processing (for example, if the app closes before processing starts) are purged from storage within 6 hours. You should keep your own copies of original lab reports. Upload metadata (filename, dates, status) may be retained for your upload history list.

8. Who we share data with (subprocessors)

We use trusted service providers who process data on our behalf under written terms. They may only use data as instructed by us. The table below lists key subprocessors as of the last updated date:

Some subprocessors are used only for optional campaign measurement when your campaign measurement preference allows (cookie banner, Settings !' Privacy & data). They help us understand whether marketing campaigns led to account creation or membership — not how you use health features. Browser tags and server-side conversion forwarding respect the same preference. Campaign measurement uses hashed contact information and coarse conversion events; we do not send lab results, biomarker values, conditions, medications, or other health profile content to advertising platforms.

Provider	Role	Data involved	Location / notes
Supabase, Inc.	Database, authentication, encrypted storage, Row Level Security	Account, profile, health records, chat, memories, embeddings	EU/US — DPA via Supabase dashboard (Org !' Legal documents)
Fly.io, Inc.	API application hosting (NestJS backend)	Request metadata, operational logs; health payloads in transit to/from API	US/EU regions — pre-signed DPA in Fly dashboard !' Compliance
Vercel, Inc.	Web application hosting (Next.js frontend)	Access logs, IP addresses; may proxy /api/v1 to API when configured	Global edge — DPA incorporated in Vercel terms (vercel.com/legal/dpa)

Anthropic, PBC	OCR, chat, summaries, memory extraction, titles, drug interaction text	Lab images/PDFs, chat content, assembled health context	US — DPA in Anthropic Commercial Terms; no training on API customer data
OpenAI, LLC	Text embeddings for memory search only	Memory text (category, key, value)	US — DPA at openai.com/policies/data-processing-addendum
RevenueCat, Inc.	Subscription management and webhook billing events	User ID, subscription status, product IDs	US — DPA at revenuecat.com/dpa
Apple Inc. / Google LLC	In-app purchases when you subscribe via App Store or Play Store	Payment handled by platform; we receive entitlement status via RevenueCat	Per platform policies
Resend, Inc.	Transactional email (when enabled)	Email address, notification content	US — DPA at resend.com/legal/dpa
Expo (push infrastructure)	Mobile push notifications	Push token, notification payload	US
Upstash, Inc.	Redis job queue (OCR and notification workers)	Job metadata (upload IDs, user IDs); not full health records in queue payloads by design	EU/US — DPA incorporated in Upstash terms (upstash.com/trust/dpa.pdf)
PostHog, Inc.	Optional product analytics (when you opt in)	Pseudonymous device ID, product event names, coarse usage metadata — not lab or health profile content	EU (eu.posthog.com) — DPA at posthog.com/dpa
Meta Platforms, Inc.	Optional campaign measurement (Conversions API via server-side tags when enabled)	Hashed email (SHA-256), pseudonymous event IDs, browser/click identifiers (fbp/fbc), coarse conversion events (e.g. signup, membership) — not lab values or health profile content	US — Meta Business Tools Terms; only when you opt in to campaign measurement where required
Google LLC	Optional Google Tag Manager, Google Ads, and Google Analytics measurement when enabled	Hashed contact info for enhanced conversions where used, gclid, pseudonymous analytics IDs, coarse conversion events — not lab values or health profile content	Global — Google Ads/Analytics terms; only when you opt in to campaign measurement where required
Stape.io	Server-side tag management hosting (routes measurement events when campaign measurement is enabled)	Event payloads for measurement forwarding (hashed contact info, click IDs, coarse conversion metadata) — not lab values or health profile content	EU/US — DPA at stape.io; optional subprocessor for measurement only

We may also disclose data if required by law, court order, or to protect rights, safety, and security. In a merger or acquisition, data may transfer to the successor entity subject to this policy or equivalent protections, with notice where required.

9. International data transfers

UNIT01 d.o.o. is established in Serbia. Our infrastructure and subprocessors may process data in Serbia, the European Economic Area, the United Kingdom, the United States, and other countries where our providers operate.

Where GDPR/UK GDPR applies and data is transferred outside adequacy decisions, we implement appropriate safeguards such as the European Commission Standard Contractual Clauses (2021/914), UK International Data Transfer Addendum, or equivalent mechanisms, and assess transfer risks where required.

You may request more information about safeguards by contacting hello@unit01.dev.

10. How long we keep data

- Account and health data: retained while your account is active and as needed to provide the Service
- After account deletion: personal data is deleted or anonymised via cascade deletion on our primary database, subject to backup retention windows (typically rolling backups for a limited period) and legal hold requirements
- Chat history: retained until you delete conversations (where supported) or delete your account
- AI memories: soft-deleted when you remove them; inactive rows may remain in backups for a limited time
- Subscription records: retained as required for accounting, tax, and fraud prevention
- Gift voucher records: retained for fulfilment, support, accounting, and fraud prevention until redeemed and consumed, expired, or refunded, and as required by law thereafter
- Security logs: retained for a limited operational period (typically up to 90 days unless needed for incident investigation)
- Raw lab files: removed after successful extraction; abandoned uploads purged within 6 hours (Section 7)

Data export (PDF/CSV) is available to subscribers in Settings. You may also request a copy of your data by emailing hello@unit01.dev under applicable access rights.

11. Security measures

We implement technical and organisational measures appropriate to the sensitivity of health data, including:

- Encryption in transit (TLS) and encryption at rest on our database and storage
- Row Level Security (RLS) on user tables so each user can access only their own records
- Authentication via industry-standard JWT sessions (Supabase Auth)
- Access controls and least-privilege for production systems
- Signed URLs for time-limited upload access
- API subscription and authentication guards
- AI safety middleware (prompt injection and PII screening on AI routes)
- Rate limiting on AI endpoints

No system is perfectly secure. You are responsible for safeguarding your login credentials and devices. Report suspected unauthorised access immediately to hello@unit01.dev.

12. Cookies and similar technologies

Our web application uses:

- Strictly necessary cookies / local storage for authentication session management (Supabase)
- Similar technologies in mobile apps for session persistence

We do not load third-party advertising or campaign measurement tags without your consent where the law requires it. Optional product analytics (PostHog, EU-hosted) and optional campaign measurement (Google Tag Manager, Google Analytics, Google Ads, Meta Pixel, Meta Conversions API, and our server-side tag host Stape at edge.getfindings.com) are controlled via the privacy notice, footer Privacy choices, or Settings !' Privacy & data. When you are signed in, your saved preferences are synced to our servers so browser tags and most server-side conversion forwarding follow the same rules. **Trial and membership** conversions may still be reported server-side when you subscribe (hashed contact and purchase value only) so we can attribute billing to our contract with you — never lab or health data. Product analytics uses a pseudonymous device identifier. Campaign measurement uses hashed contact information (e.g. SHA-256 email), click identifiers such as gclid/fbp/fbc, and coarse conversion events (signup, membership) only. We do not send lab values, biomarker results, health profile fields, conditions, or medications to analytics or ad platforms. Where required by law (EEA, UK, Switzerland), non-essential analytics and campaign measurement stay off until you opt in with an explicit **Accept all**; you may choose **Reject optional** with the same ease. The privacy notice cannot be dismissed without a choice in any region. Optional toggles in Customize start off until you turn them on in strict regions; elsewhere **Continue** saves accept-all preferences until you change them in Settings.

Manage preferences anytime in Settings !' Privacy & data or the site footer Privacy choices link.

13. Your privacy rights

Depending on your location, you may have the following rights regarding your personal data (subject to exceptions):

- Access — obtain a copy of data we hold about you
- Rectification — correct inaccurate profile or health data in settings
- Erasure — delete specific data without closing your account (https://getfindings.com/delete-data), or delete your entire account (Settings !' Delete account; https://getfindings.com/delete-account)
- Restriction — ask us to limit processing in certain circumstances
- Portability — receive data in a structured, machine-readable format where technically feasible
- Objection — object to processing based on legitimate interests
- Withdraw consent — where processing is consent-based (including health data), without affecting prior lawful processing
- Complaint — lodge a complaint with a supervisory authority

13.1 European Economic Area and UK

UNIT01 d.o.o. is established in Serbia, outside the EEA and UK. For privacy requests and to exercise your rights, contact our data protection contact Rade Joksimovic at hello@unit01.dev.

Where GDPR Article 27 (or equivalent UK rules) requires a local representative for non-EEA controllers, that role must be held by a person or organisation established in the EU or UK — it cannot be the same individual acting from Serbia. We will publish EU/UK representative contact details here when appointed. Until then, use the data protection contact above; we respond to all requests without undue delay.

EEA/UK users may also lodge a complaint with their local data protection authority.

13.2 Serbia

Serbian residents may contact the Commissioner for Information of Public Importance and Personal Data Protection (Poverenik).

13.3 United States (including California)

Where state privacy laws apply (e.g. CCPA/CPRA), you may have rights to know, delete, and correct personal information, and to opt out of "sale" or "sharing" — we do not sell personal information as defined by those laws. Submit requests to hello@unit01.dev.

To exercise rights, email hello@unit01.dev from your registered address. We may verify your identity before responding. We aim to respond within one month (or the period required by law).

14. Your choices in the product

- Edit profile, medications, trackers, and memories in the app
- Control notification preferences in Settings
- Delete individual AI memories in Health Profile !’ AI Health Context
- Dismiss retest reminders (subscription may be required)
- Delete your entire account in Settings — this is irreversible and removes access immediately
- First-report preview: you may use limited features without a paid subscription; see Terms for access rules

15. Children

The Service is intended for adults aged 18 and over (or the age of majority in your jurisdiction). We do not knowingly collect personal data from children. If you believe a child has provided data, contact hello@unit01.dev and we will delete it.

16. Data breaches

If we become aware of a personal data breach likely to result in risk to your rights and freedoms, we will notify the competent supervisory authority where required and inform affected users without undue delay when notification is legally required.

17. Profiling and automated decisions

The Service uses algorithms to classify biomarker status (optimal / sufficient / out of range), compute trends, detect statistical correlations, and schedule retest reminders. These are wellness and educational tools, not clinical decision systems. AI-generated text is probabilistic and may be incomplete or incorrect — always verify important health decisions with a qualified professional.

18. Changes to this policy

We may update this Privacy Policy to reflect product, legal, or regulatory changes. Material changes will be notified via the Service or email where required. The "Last updated" date at the top indicates the current version. Continued use after the effective date constitutes acceptance where permitted by law.

19. Contact

UNIT01 d.o.o.
Vodovodska 75, 11030 Belgrade, Serbia
Email: hello@unit01.dev

For data protection enquiries, privacy rights requests, and security reports, use the same address: hello@unit01.dev.

UNIT01 d.o.o. · hello@unit01.dev · Last updated 8 July 2026

UNIT01 d.o.o. · hello@unit01.dev · Last updated 8 July 2026

UNIT01 d.o.o. · hello@unit01.dev · Last updated 8 July 2026

UNIT01 d.o.o. · hello@unit01.dev · Last updated 8 July 2026

UNIT01 d.o.o. · hello@unit01.dev · Last updated 8 July 2026

UNIT01 d.o.o. · hello@unit01.dev · Last updated 8 July 2026

UNIT01 d.o.o. · hello@unit01.dev · Last updated 8 July 2026

UNIT01 d.o.o. · hello@unit01.dev · Last updated 8 July 2026

UNIT01 d.o.o. · hello@unit01.dev · Last updated 8 July 2026

UNIT01 d.o.o. · hello@unit01.dev · Last updated 8 July 2026