

Aidealy Product Privacy Policy

Effective date: 2026-09-27 · **Version:** 1.11.0 · **Last updated:** 2026-09-27

This Privacy Policy explains how **Aidealy Ltd.** ("Aidealy", "we", "us") handles personal data in connection with our **software product and service** - the Aidealy web application (admin and analytics/chat surfaces) and the account you use to access it (the "Service").

It covers people who **sign in to or administer an Aidealy account** - account owners, administrators, and the team members an account invites - and Aidealy's own processing to run, secure, bill for, support, improve and market the Service.

It does **not** cover: our **marketing website** at aidealy.ai (see the **Website Privacy Policy**); the personal data of your developers and the source code, prompts, and developer-activity data that Aidealy processes **on your behalf as a processor** (see §3 and the **Data Processing Agreement**); or the Aidealy IDE extension's developer telemetry (see the **Extension Privacy Notice**).

1. Who we are (controller) and how to contact us

Aidealy Ltd. is a company registered in **Israel**, at **Hamidron 1, Herzliya 4654110, Israel** (company registration number 517367066). For the account and product processing described in this Policy, Aidealy is the **data controller**.

- **Privacy contact / data-protection enquiries:** privacy@aidealy.ai
- **EU representative (GDPR Art. 27):** being appointed via **EuVerify** (GDPR Art. 27 service - EU + UK plan). Until that appointment completes, please contact **privacy@aidealy.ai** about anything you would otherwise raise with our EU representative. We will publish the representative's name and EU member-state address here on completion.
- **UK representative (UK GDPR Art. 27):** being appointed via **EuVerify** (same EU + UK plan). Until that appointment completes, please contact **privacy@aidealy.ai**. We will publish the representative's name and UK address here on completion.

2. Who this Policy is for

Aidealy is built for organisations and their authorised users, not for individual consumers. This Policy is written for the individuals whose personal data Aidealy controls in running the Service: the **account owner/administrator** who contracts for and manages the account, and **everyone your organisation provisions or invites onto its account** - including team members with a **pending invitation they have not yet accepted**, and developers whose seat is used **only through the IDE extension** and who may never open this web application. If you are in one of those groups, Aidealy holds account data about you (provided by your organisation - see §4) even if you never sign in, and this Policy is the notice of that processing. If your employer or another organisation gave you access to Aidealy, that organisation decides what work data is analysed and is responsible for telling you about that processing (see §3); this Policy explains Aidealy's own handling of your **account data**.

3. Our roles: controller and processor

Aidealy plays **two different roles**, and the law treats them differently:

- **As a controller** (this Policy), Aidealy decides how to handle the data needed to **run the account relationship and the Service for its own purposes** - your account identity and log-in, account security, billing, support, service improvement, and our communications with you.
- **As a processor**, Aidealy handles your organisation's **work data** - source code, AI prompts and responses, git author identity, and developer-activity metrics (including an estimate of the working time each developer spent, the split of their code between AI-assisted and human-written, and attributable AI-model usage) - strictly **on your organisation's instructions and on its**

behalf. For that data, **your organisation is the controller**, and Aidealy's obligations are set out in our **Data Processing Agreement (DPA)**, not this Policy.

The same company can be a controller for some data and a processor for other data - that is the case here. Questions about the work data Aidealy processes for your organisation should go to that organisation first.

4. What information we collect (as controller)

To provide the Service we collect and use:

- **Account identity & profile:** your name, your email address and any further email addresses linked to you in the account (one person may hold several: for example a second address under which a development tool on your machine reports your work, which we attach to your record), your role/permissions in the account, account status, the team you belong to, and (where your organisation uses single sign-on) an employee identifier supplied by your organisation. Some of this is provided directly by you; some is **obtained from your organisation** - your account administrator provides it when they invite or provision you (so the source of that data is your organisation); and a further email address is reported by the development tool on your machine that sends data to the Service.
- **Unmatched-address review list:** where an email address reported from a developer's machine does not match anyone your organisation has listed as its people, we keep that full address on a review list for your organisation, so that a person can later work out what the mismatch means; nothing is blocked or deleted on the strength of an entry. Each entry is deleted automatically **90 days after that address was first seen**, whether or not the mismatch is still occurring (an address that keeps being reported after that is recorded afresh, so it can reappear on the list for further 90-day periods), and it is cleared sooner when the person it concerns is marked as having left the company (§9).
- **Authentication data:** the credentials and sign-in method you use - password (with optional two-factor authentication), email magic-link, Google sign-in, or your organisation's SAML single sign-on - handled through our authentication provider.
- **Account-security & audit records:** records of account-security events such as sign-in success/failure, two-factor checks, single-sign-on enforcement, and sign-out, used to secure the account and detect abuse.
- **Billing & subscription data:** your subscription, plan, seat counts and billing status. **Payment is processed by Paddle as our Merchant of Record** - Paddle is the seller of record and handles checkout, card details, invoicing and sales tax; Aidealy does not store your full card or payment details. What Aidealy itself keeps for billing is: your organisation's subscription and customer references, and a **daily, pseudonymous billing-evidence snapshot** of the account's user roster: for each provisioned or invited user, an internal user code, role, seat type, account status, and joined/last-active dates (**no name, no email**), together with daily user-count records: a total-user tally, separate tallies of joined users and of invited users who have not yet joined, the same totals and splits per seat type, and a tally of the git identities our systems classify as AI agents **with git work ingested in the trailing 30 days** (the tally rises and falls with that activity: an identity with no recorded git activity is never counted, and an identity whose access has been revoked can stay counted for up to one further 30-day window; the tally is derived from the customer's git data, which is read for the count but not stored in these records, and no git handle, name or email is stored in them). We keep each daily snapshot record for **2 years (730 days)** in write-protected (tamper-evident) storage as our evidence of why each month's bill changed; **only users who have joined** count toward the billable figure, and the stored daily counts are the evidence base from which the invoiced seat quantity can be recomputed - none of them is itself the invoiced quantity, nor the contractual seat-metering record. See §9 and the **Data Retention & Deletion Policy**.
- **Communications & support:** messages you send us (e.g. support requests) and our service and product communications to you, together with a record on your account entry of the last

account-invitation email we sent you (when it was sent and whether it was a first invitation, a re-sent invitation or an administrator invitation). That record stays with your account entry - including after a de-identification (§9) - until that entry is deleted.

- **Use of the AI analytics feature:** when you use Aidealy's natural-language analytics ("chat") feature, we process the queries you submit and the results returned, to provide and improve that feature (see §6).
- **Error monitoring (including Session Replay):** we use **Sentry** to detect and diagnose errors in the application. When an error occurs, Sentry may capture a **masked "session replay"** of the interaction. By default this replay **masks all text and form inputs and blocks media** before it leaves your browser, and we do not continuously record sessions - replay is captured around errors (see §8).

We do **not** intentionally collect special-category / sensitive personal data about account users for our own purposes, and ask you not to submit it through support messages.

5. Why we use your information, and our legal basis

Purpose	Personal data	Legal basis (EEA/UK)
Provide the Service and your account	Account identity, authentication	Performance of a contract (with you or your organisation)
Secure the account; prevent fraud/abuse; keep the account membership accurate	Authentication, security/audit records, the unmatched-address review list (§4)	Our legitimate interest in security, fraud prevention and accurate account membership
Bill for the Service	Billing & subscription data	Performance of a contract ; legal obligation (tax/accounting)
Keep billing evidence (prove why each month's bill changed; defend billing disputes)	Daily pseudonymous user-roster snapshot & user counts (§4)	Our legitimate interest in billing accuracy and in establishing, exercising or defending billing claims (retained no longer than the 730-day period stated in §4)
Provide support and service communications	Communications, account data	Performance of a contract ; our legitimate interest in supporting users
Provide and improve the AI analytics feature	Query text and results	Performance of a contract ; our legitimate interest in improving the Service
Improve and develop the Service	Usage and account data	Our legitimate interest in improving our product
Keep legal-agreement acceptance records and claims-evidence records	Acceptance records; claims-evidence records (§9)	Our legitimate interest in establishing, exercising or defending legal claims
Send marketing / product-update messages	Name, email	Your consent , where required by the law that applies to you (see §7)

Where we rely on legitimate interests, you have the right to object (see §9). You can withdraw any consent at any time without affecting prior processing.

6. Artificial intelligence in the Service

- **You are interacting with an AI system.** Aidealy's analytics/chat feature is powered by AI models. We tell you when you are interacting with AI rather than a person.
- **No solely-automated decisions about you.** As required by GDPR Articles 13(2)(f) and 22 (and, in the UK, Articles 22A to 22D of the UK GDPR, and equivalent laws), we confirm that Aidealy does **not** make decisions about you based **solely** on automated processing - including profiling - that produce legal or similarly significant effects; our analytics are decision-support intended for human review. If you believe an automated decision has significantly affected you, raise it with the **organisation that deployed Aidealy for you** (your employer or engaging organisation): that organisation makes the decisions the analytics inform, and it is responsible for providing human review, hearing your view, and handling your contest of the decision. Aidealy assists that organisation in meeting those obligations; if you contact us at privacy@aidealy.ai instead, we will route your request to it.
- **Use to evaluate individuals, and your rights.** Your organisation may use Aidealy's analytics to evaluate the performance or behaviour of individuals, including you. Where it does, Aidealy is provided as a **high-risk AI system** under the EU AI Act and your organisation is the **deployer**: it must **inform you that you are subject to the system**, must apply meaningful human oversight, and **must not make a decision about you based solely on the Service's output** (see "No solely-automated decisions about you" above). You keep your rights under data-protection law - including, in the EEA, the right under GDPR Article 22 to obtain human review of, to express your view on, and to contest, a decision based solely on automated processing, and, in the UK, the corresponding safeguards under Articles 22A to 22D of the UK GDPR (substituted for Article 22 by section 80(1) of the Data (Use and Access) Act 2025, with effect from 5 February 2026): to be informed about such decisions, to make representations, to obtain human intervention, and to contest them - and your organisation must give you the transparency information required of it (in the EEA/UK, GDPR Articles 13-14). Similar workplace-AI protections are emerging outside the EU. For example, **Illinois** requires employers to tell employees when AI is used in employment-related decisions. From 2027, **Colorado** adds rights to a disclosure, correction and human review of certain adverse AI-influenced consequential decisions. These are examples, not a complete list. Your organisation, as the deployer/employer, is responsible for the employer-side duties that apply to it. **Korea's** AI Framework Act also imposes AI-transparency duties, but most of those fall on **Aidealy** as the AI operator (see the AI Addendum).
- **Message classification, including tone.** Where the analytics include automated classification of typed messages (the prompts developers type to their AI coding tools), each message is labelled in general terms for the part of the system it concerns, the kind of work it asks for, any quality concern it raises, and its apparent tone. This is **profiling** derived from **message text only** - each message is read together with up to three of your previous messages in the same session, as context only, and never with voice, face, keystroke dynamics, or any biometric signal - is subject to human review, and is **not** used to infer any health or mental-health condition.
- **Concerns about our AI.** Separately from the rights above, if you have a concern about how Aidealy develops or uses AI - for example the fairness of AI-assisted evaluations of individuals, the guardrails around them, or possible misuse - you can raise it with us at **ai-concerns@aidealy.ai**. This is an additional way to reach us; it does not replace your rights or the human-review and contest routes described here, and it is open to individuals whose work is evaluated through the Service whether or not they use it directly. If you raise a concern, we do not disclose your identity to your organisation without your agreement, except where the law requires it or where the concern cannot be assessed or acted on without contacting the organisation - in which case we tell you so before proceeding. We will acknowledge a report to this channel within **ten (10) Business Days** (days other than Friday, Saturday, or a public holiday in Israel); beyond that acknowledgement, this channel does not carry a service level, and we do not commit to a particular investigation, outcome, or response timeline.

- **We do not train AI models on your data.** Aidealy does not use your account data or your organisation's work data to train AI models. Our AI providers are **contractually committed not to train their models on our customers' content: Anthropic** ("Anthropic may not train models on Customer Content from Services" - a contractual prohibition) and **OpenAI** ("OpenAI will not use Customer Content to develop or improve the Services, unless Customer explicitly agrees to such use" - OpenAI Services Agreement s. 4.2); **Aidealy has not agreed, and will not agree, to any use of customer content for training.**
- **How our AI providers handle your AI queries and your organisation's work data.** When an AI feature processes your activity, the relevant content is sent to one of our AI model providers, which **processes it to generate a result.** Aidealy uses two AI model providers for the AI features of the Service (the analysis and scoring of source code, the classification of developers' typed messages, and the natural-language analytics assistant, which is the analytics (chat) agent described in this Policy): **OpenAI OpCo, LLC** (an AI model API provider in the United States) and **Anthropic, PBC** (an AI model API provider that stores API data in the United States and offers no EU option). Either provider may perform any of those steps; Aidealy chooses the routing and may change it at any time, including as a fallback when one provider is unavailable, without adding a provider. Both are named, with their legal entities and processing locations, on our dated **Sub-processor List** (<https://aidealy.ai/legal/sub-processors>). **Before a provider that is not already named on that list is used for your data, Aidealy tells your organisation at least thirty (30) days in advance, as the Data Processing Agreement requires; that agreement allows one exception: where a provider suddenly becomes unavailable or must be replaced urgently to keep the Service running, secure or intact, Aidealy may make the change first, must then tell your organisation without undue delay, and your organisation keeps its right to object. If a processing location stated here changes, we update this Policy before the change takes effect where we make the change, and as soon as we learn of it where a provider makes it. Where the change concerns the work data we process for your organisation, your organisation, as the controller of that data, passes it on to you; for the account data covered by this Policy we post the updated Policy (§18).** Whichever provider runs a step, it keeps what it receives for a limited period under its own terms, does not use it to train its models, and remains subject to the contract (DPA) we have in place with it: OpenAI may retain API inputs and outputs for up to 30 days to provide its services and identify abuse, and data sent through its batch and file interfaces persists until deleted, so where a step runs through those interfaces Aidealy sends the requests with storage switched off, deletes the files it uploads and the files the batch produces once the results are stored, and in any case sets each of those files to expire 48 hours after it is created (the batch's own status record, which cannot be deleted through OpenAI's interface, holds no request or response content); Anthropic deletes inputs and outputs within 30 days of receipt or generation, subject to the exceptions its policy states (for example content flagged for usage-policy enforcement), and data sent through its batch interface is stored for up to 29 days after the batch is created. One feature keeps state with a provider: where the analytics (chat) agent is routed to OpenAI, it **stores its multi-turn conversation history server-side at OpenAI, in the United States, and relies on that stored history for multi-turn context** - your earlier messages in a chat conversation are held there, rather than duplicated on our own servers; Anthropic holds no server-side conversation state. When an individual who used the feature is de-identified (for example, after leaving your organisation), a **nightly deletion job removes that person's stored conversations through OpenAI's interface** (see §9); when your organisation's account is offboarded, we **delete all of its stored conversations the same way before the environment is torn down** - which also covers any conversations the nightly job cannot remove. Standalone AI responses not attached to a stored conversation cannot be individually deleted and are retained by OpenAI for **at least 30 days under its own policy** (a stated minimum, not a promised deletion date); they hold no stored conversation content. Separately, for security auditing and as our evidence of what the Service delivered, we keep our own **AI-interaction audit log** of each AI request and response - including the content of AI analytics queries and answers and the AI

model used - as a write-locked, tamper-evident record for **seven (7) years** per entry (deleted in full when your organisation's environment is decommissioned); see §9 and the **Data Retention & Deletion Policy**.

7. Marketing communications and how to opt out

With your consent where the law requires it, we may send you product news, feature announcements, tips and offers. You can opt out at any time using the unsubscribe link in every marketing message, or by contacting privacy@aidealy.ai. Specifically:

- **EEA/UK, Israel and Korea:** we send marketing email on an **opt-in** basis (and, in Korea, we ask for your marketing consent **separately** from other consents and will not refuse you the Service if you decline it).
- **Israel:** marketing emails are labelled as advertising ("פרסומת") and identify us and an easy way to unsubscribe, as required by Israeli law.
- **United States:** marketing emails identify themselves as advertising, include our postal address, and honour opt-out requests promptly (within 10 business days), as required by the CAN-SPAM Act.
- **Canada, Australia and New Zealand:** we send marketing email on an **express opt-in** basis, identify ourselves as the sender, and provide a functional unsubscribe in every message, honoured promptly (within **10 business days** in Canada and **5 working days** in Australia and New Zealand).
- **Service / transactional messages** (e.g. invitations, password resets, security and billing notices) are **not** marketing and are sent as part of operating your account; you cannot opt out of these while you have an account.

8. Cookies and similar technologies in the Service

The signed-in application uses **strictly-necessary** cookies and similar storage only - to authenticate you and keep you signed in, and to route you to the correct regional service (`aidealy_region`). Our sign-in pages also use **Cloudflare Turnstile** to protect the Service against automated sign-in abuse; any cookies or storage it sets serve that security purpose only. These are essential to provide the Service you request and to keep it secure, and do not require a consent banner; we do not use advertising, analytics, or cross-site tracking cookies in the application.

Our error-monitoring tool (Sentry) may store and access information on your device as part of capturing a **masked replay around an error**. Where consent is required for this in your region, we will obtain it; the replay is configured to mask text, inputs and media by default. See §9 and §12 for your choices.

9. Your rights

Depending on where you live, you may have the right to: **access** your data; **correct** it; **delete** it; **restrict** or **object** to processing; **data portability**; and to **withdraw consent** at any time. You also have an **absolute right to object to direct marketing** - if you object, we will stop using your data for marketing. To exercise any right, contact privacy@aidealy.ai. We will respond within the time limits set by the law that applies to you, and you will not receive discriminatory treatment for exercising your rights. Where a legal preservation obligation binds us (for example, a litigation hold or a preservation order), a deletion that would otherwise be due yields to it - to the narrowest extent, and for no longer than, necessary.

If your access to Aidealy was provided by your employer or another organisation, some requests about your **work data** are best directed to that organisation, which is the controller of that data (see §3); we will help route your request appropriately. **If that organisation has ceased to exist without a successor, or cannot be reached for a prolonged period despite our reasonable attempts, we will give effect to your request directly to the extent our systems allow** - on a verified **access** request, by providing you with a copy of the personal data we still hold about you (including your identity-linked derived metrics and scores, and, **if you used the AI analytics feature**, your stored

AI-chat conversations), prepared and delivered through the same support-executed export mechanics described in §14; on an **erasure** request, by de-identifying your identity records and, **if you used the AI analytics feature**, deleting your stored AI-chat conversations - rather than leaving it unanswered; see the DPA's controller-cessation provision.

Verification, authorised agents, and if we decline your request. To protect your data we **verify** requests (typically by confirming control of the email on the account, or equivalent evidence); where your local law allows an **authorised agent** to act for you, we accept agent requests with proof of your authorisation and verification of your identity. If we **decline** a request in whole or part, we tell you why and how to take it further: you may **appeal** by replying to our response or writing to privacy@aidealy.ai with "Appeal" in the subject, and we will have the appeal reviewed by someone other than the original decision-maker and respond **within 60 days** (or any shorter period your local law requires). If your appeal is unsuccessful and you are in a US state whose privacy law provides it, you may **contact your state Attorney General** (for example, through the Attorney General's consumer-complaint channel in Virginia and the other states with appeal-referral rights); wherever you live, you can also complain to your local data-protection authority (§19).

The reach of deletion and de-identification. If your organisation asks us to erase you (for example, when you leave the company), or you exercise a valid right to erasure, we remove your **directly identifying details** (such as name and email) from the identity records in our operational databases - the account store and the analytics graph database - and replace them with a non-identifying placeholder. We treat this as **de-identification**: we do not claim it makes the data fully anonymous, so while any residual record could still be linked to you it remains personal data and stays protected under this Policy. The following limits on its reach are stated plainly rather than overstated:

- **AI-provider conversation history (§6) - reached by a separate, automatic deletion job rather than by the rename.** If you used the AI analytics feature, when you are de-identified a nightly job **deletes your stored AI chat conversations from our AI provider**, so deletion completes on a **nightly cycle** - normally within about a day. Any conversations that job cannot remove - and in any case your organisation's whole conversation history - are deleted when its account ends (§14). Standalone AI responses not attached to a stored conversation cannot be individually deleted; OpenAI retains those response records for **at least 30 days under its own policy** (a stated minimum, not a promised deletion date), and they hold no stored conversation content.
- **Historical work-data records in our analytics data lake and its query layer.** Your real name and email persist in those historical records until they are deleted - **five (5) years after each stored copy was written** (a copy re-written when we re-organise our storage starts a new five-year period; a deleted copy's residual storage version is purged within about ninety (90) days after that) - or earlier when your organisation's account is offboarded. The de-identification does not rewrite these historical records by design; the scheduled deletion and our yearly retention review are the mechanisms that purge them. (These records are your organisation's work data - see §3 and the **Data Retention & Deletion Policy**.)
- **Account-security event records.** Events such as inviting you to an account or adding you as an IDE user are written to a tamper-evident security register that records your real email address. We do not edit these records - including after a de-identification - because they are kept as security evidence; they are retained and deleted on the security-log schedule in the **Data Retention & Deletion Policy**.
- **Legal-agreement acceptance records.** If you use the Aidealy IDE extension, we keep an identity-bound, versioned record of your acceptance or decline of the Extension EULA and related legal dialogs (described in the **Extension Privacy Notice**). It is held in write-protected, tamper-evident storage, is not edited or de-identified - including after a de-identification - because it is kept as evidence to establish, exercise or defend legal claims, and is retained for **seven years after the later of** the end of your organisation's relationship with Aidealy or your deactivation as a user.

- **Claims-evidence records.** Our evidence of what the Service delivered to your organisation is our **AI-interaction audit log** (§6): the write-locked (tamper-evident) record of each AI request and response, including the content of your AI analytics queries and answers and the AI model used. We do not edit or remove individual entries before their period ends - including after a de-identification - because the log is kept as security and legal-claims evidence; your identity therefore persists in that stored content for the log's retention period. Each entry is retained for the period during which a related legal claim could be brought (in Israel, generally **seven years**), and the whole log is deleted earlier, in full, when your organisation's account is offboarded - it does **not** survive account closure. Separately, where a contested AI output is human-reviewed, we keep a short account-side **contest-review note** of that review (the date, what was reviewed, and the outcome); like the acceptance records above, these notes are a legal-evidence holding that we do not edit or de-identify - including after a de-identification - and they are retained for the same seven-year claims horizon and then deleted; a documented legal hold can extend a specific record (see the preservation caveat earlier in this §9, and §14). The work-data retention periods above are unaffected; the full schedule is in the **Data Retention & Deletion Policy**.
- **Billing-evidence snapshot records.** The daily billing-evidence snapshots described in §4 hold a pseudonymous roster of your organisation's users - your internal user code, role, seat type, account status, and joined/last-active dates, with **no name and no email**. These records are **never edited after a de-identification**: snapshot records written *before* you were de-identified keep your pre-departure role, seat and status history, under your internal user code, for each record's full 730-day period, because they are our evidence of why past bills were correct - records we keep to establish, exercise or defend legal claims (in the EEA/UK, the erasure exception for legal claims; for California residents, we may decline a deletion request for this data on the grounds that it is reasonably necessary to **complete the transaction** for which it was collected and for **internal uses aligned with your reasonable expectations** as disclosed in §4). Because your internal user code could still be linked back to you, these records remain **pseudonymous personal data** - not anonymous data - and stay protected under this Policy; where a record notes that an account's status is "anonymized", that is a status flag describing the account, not a claim about the data. Each daily record is purged automatically when its own 730-day period expires; in normal operation nothing else deletes it.
- **Backup copies of the account database.** Your details remain in any backup copy taken before a de-identification until that copy is deleted on the backup schedule in our **Data Retention & Deletion Policy**. Until then the copy is **kept beyond use** - access-restricted and not actively processed - and if we restore from a backup we re-apply the erasures and deletions completed since it was taken, so a restore does not undo them.
- **The reach of correction (rectification).** Correction has the same reach as de-identification: your **identity fields** (a misspelled name, a wrong email) are correctable - your organisation's administrator can fix them directly in the Admin application; the **derived analytics are recomputed from the source data as new data arrives**, so a correction made at the source (for example, in your organisation's git provider) flows into outputs computed from then on; and where you **contest a derived metric or score** held in one of the stores we do not edit (listed above), the correction takes the form of your organisation recording your contest and its decision alongside the output it relied on - ask your organisation, which owns those decisions, and see §6 for your review-and-contest route.
- **De-identification does not stop new collection.** If your development environment keeps sending data to the Service after you were de-identified, the Service treats it as a new person joining your organisation's team and creates a **new, active record carrying your email address and a display name taken from it**. The de-identified record is never reversed or reused, but fresh personal data about you is collected again; ask your organisation to revoke your access (for example, uninstalling or signing out the IDE extension) as part of its own offboarding.

10. Who we share your information with

We share personal data with service providers ("processors"/"sub-processors") who act for us under a contract that prohibits them from selling or sharing it, or using it outside our instructions:

- **Supabase** - database and authentication.
- **Amazon Web Services (AWS)** - cloud hosting and storage.
- **Vercel** - application hosting; **Cloudflare** - bot protection (Turnstile) on the sign-in pages.
- **Sentry** - error monitoring (including masked Session Replay).
- **Mailgun** - sending account and product emails.
- **Paddle** - payment processing as Merchant of Record; Paddle's handling of your payment data is described in Paddle's Privacy Notice (paddle.com/legal/privacy).
- **OpenAI OpCo, LLC** (United States) and **Anthropic, PBC** (United States, no EU option) - the AI model providers behind the AI features of the Service (§6); either may perform any AI step, and both are named, with their legal entities and processing locations, on our dated **Sub-processor List** (<https://aidealy.ai/legal/sub-processors>).
- **LangSmith** (LangChain, Inc.) - LLM tracing / observability of the AI features. The integration is **wired but dormant**: no data flows to LangChain unless and until we activate it, which we will do only once LangChain's data-processing terms are executed and on file. Once activated, it would process the prompts and responses sent to the AI models, in the **United States**.
- Additional providers used to operate the Service are listed in our **Sub-processor List**.

We do **not sell** your personal information, and we do **not share** it for cross-context behavioural advertising.

Legal process and authorities. We may also disclose personal data to **courts, litigants, and regulatory or law-enforcement authorities** where valid legal process or applicable law requires it, limited to what the process or law requires. Where such a demand concerns your organisation's **work data**, it is handled under the **Data Processing Agreement** (including challenging or narrowing overbroad demands, and notice to your organisation where lawful).

Business transfers. If Aidealy is involved in a **merger, acquisition, financing, corporate reorganisation, or sale of some or all of its business or assets**, the account personal data we hold as controller may be disclosed to the prospective or actual acquirer, investor, or lender (and their professional advisers) as part of that transaction - during due diligence only to the minimum extent necessary and under confidentiality obligations - and may be transferred to the **successor entity** when the transaction completes. Where GDPR or UK GDPR applies, we rely on our **legitimate interests** in conducting and completing corporate transactions as the legal basis for these disclosures. Any successor may use your personal data only **consistently with this Policy as in force at the time**; if a successor intends to materially change that, you will be given the notice (or, where required, asked for the consent) that applicable law requires before the change applies to you. Under US state privacy laws such as the California CCPA, a transfer of personal information as an asset in such a transaction is not a "sale" or "sharing" of personal information. Your organisation's **work data** is not disclosed in due diligence; in a completed transaction it remains governed by the **Data Processing Agreement** with your organisation, which binds any successor to the same obligations.

11. International data transfers

Aidealy operates **separate EU and US regions**. Each customer's account is assigned to **one** region, and we do not move a customer's data between regions. Aidealy is based in **Israel**, and some of our providers are located in the United States or elsewhere. We rely on the following safeguards:

- **From the EEA:** Israel benefits from a **European Commission adequacy decision**, so transfers to Aidealy in Israel are permitted without additional safeguards. For providers in the United States, we rely on the **EU Standard Contractual Clauses**, together with the **EU-US Data Privacy Framework** where the provider holds a live certification. You can request a copy of the safeguards using the contact in §1.

- **From the UK:** the UK recognises **Israel as adequate** ("data bridge"), so UK→Israel transfers need no additional safeguards. For US providers we rely on the UK-US data bridge or the UK International Data Transfer Agreement / EU SCCs with the UK Addendum.
- **From Israel:** we transfer data abroad in accordance with the Israeli Protection of Privacy (Transfer of Data Abroad) Regulations, relying on your consent and/or contractual undertakings from recipients.

The region your data sits in follows **your organisation's** choice, not your own location. So if your organisation selected the **US region**, your account data is stored and processed in the United States even if you are located in the EEA or UK - your data is collected directly into that US region and is not held in Europe first. European (and UK) data-protection law still applies to you in that case, and you keep all the rights described in this Policy. Where your organisation provides your details to us (for example when your administrator invites or provisions you - see §4), that disclosure by your organisation to Aidealy is covered by the Israel adequacy decisions described above. Where we then use US-based service providers to process that data on our behalf, we rely on the **EU Standard Contractual Clauses** in our agreements with those providers (and the EU-US Data Privacy Framework where the provider is certified) to safeguard it.

AI processing - one exception to region residency. When an AI feature processes a prompt, the content submitted for that request leaves your organisation's region for whichever of our two AI model providers runs that step (§6), even if your organisation is on our **EU region**. **Anthropic, PBC** does not offer EU data residency: a prompt sent to Anthropic is transferred to Anthropic in the **United States**, where Anthropic stores its API data, and that transfer is safeguarded by the **EU Standard Contractual Clauses** in our agreement with Anthropic (see the safeguards described above in this §11 and our DPA). Today we call **OpenAI OpCo, LLC** at its default endpoint (OpenAI is a **United States** recipient) (where the analytics (chat) agent is routed to OpenAI, it also stores that agent's conversation history server-side in the United States): OpenAI's EU regional option is supported in our software but is not yet switched on, so until it is, the OpenAI leg is likewise a transfer to a US recipient safeguarded by the **EU Standard Contractual Clauses**. Whichever provider runs a step, this affects only the **content submitted for that request**, which the provider holds under the retention terms described in §6 and **does not use to train any AI model** - the data we **store** for you remains in your organisation's region.

12. Your US privacy choices

We collect **identifiers** (name, email, and the internal user codes in our billing-evidence snapshots), **professional/employment information** (your role, seat type and organisation), **internet/usage and security activity** (audit logs, error data, joined/last-active dates), **commercial information** (subscription/billing references and the billing-evidence snapshots described in §4), and the **contents of your communications** (support messages and AI queries). For each we describe the sources, purposes, retention and recipients in this Policy.

We do **not sell** your personal information and do **not share** it for cross-context behavioural advertising, so no "Do Not Sell or Share" link is required for the Service. Certain login credentials and message contents (including the queries you submit to the AI feature) are treated as **sensitive personal information**; we use them to provide, operate and secure the Service. If you would like to limit our use of your sensitive personal information, contact privacy@aidealy.ai and we will honour that right where it applies to you. US residents also have rights to know, delete, and correct their personal information - see §9.

13. Region-specific information

13.1 EEA / European Union

You may lodge a complaint with your local **supervisory authority**. The status of our EU representative, and how to reach us in the meantime, are set out in §1. We rely on the legal bases in §5, and provide the information required by Articles 13 and 14 of the GDPR, including - where we

receive your data from your organisation rather than from you - the categories and source of that data.

13.2 United Kingdom

UK residents may complain to the UK data protection regulator, the **ICO**, at ico.org.uk - please contact us first so we can try to resolve it. We will acknowledge a privacy complaint within **30 days** and respond without undue delay. Our UK representative is in §1.

13.3 Israel

As an Israeli company, we provide notice under the Protection of Privacy Law. When you provide personal data, we tell you whether it is required and the consequence of not providing it, the purpose, our identity and contact details, who we share it with, and your rights to **access** and **correct** your data. Providing account data is necessary to use the Service; **if you do not provide it, we cannot give you access.**

13.4 South Korea

If you are in South Korea, the Personal Information Protection Act (PIPA) applies. We obtain your **separate, itemised consent** at the point of collection for: (1) collection and use of your data; (2) provision to third parties where applicable; and (3) **marketing** (sought separately, and you may decline without losing the Service). For the **transfer of your account data outside Korea, we do not rely on your consent**: we rely on **Article 28-8(1)3 of PIPA** - the overseas outsourcing of processing, and storage, of personal data necessary to perform our contract with you - and we make the disclosure PIPA requires for that route here: **(a) what is transferred** - the account data described in §4; **(b) to whom, and where** - to Aidealy Ltd. in **Israel**, and to the service providers named on our **Sub-processor List** (which states each recipient's name, role, location, and how to reach it), principally in the **EU** and the **United States** (§10, §11); **(c) when and how** - continuously over encrypted connections, for as long as you use the Service; **(d) the recipients' purpose and retention** - solely to provide the Service on our behalf, for the retention periods in §14 and the Data Retention & Deletion Policy; and **(e) how to refuse, and the effect** - you may refuse the overseas transfer by contacting privacy@aidealy.ai; because the Service is operated from outside Korea, refusal means we cannot provide the Service to you. If any of these matters change, we will notify you and, where PIPA requires, seek consent. Korea has designated only the EU/EEA as an adequate destination - not Israel or the United States - so we apply the protective measures PIPA requires to these transfers through our contracts with each recipient.

13.5 Japan

If you are in Japan, the Act on the Protection of Personal Information (APPI) applies to our handling of your account data. We specify and make public our purpose of use, which covers our product communications. Your account data is handled outside Japan: by us in **Israel** and by the providers listed in §10, principally in the **EU** and the **United States**. **We do not ask for your consent for these cross-border transfers.** Before we begin serving customers in Japan, we will put in place with each recipient of personal data outside Japan (other than recipients in the EU or the United Kingdom, which Japan's Personal Information Protection Commission has designated as equivalent) **contractual safeguards that continuously secure protections equivalent to those the APPI requires of a Japanese business operator** (APPI Article 28(1); PPC Rules Article 16(i)). Under those safeguards we will: (1) confirm at least annually that each recipient implements the protections, and monitor the laws of each recipient country that could affect them; (2) suspend transfers to a recipient that can no longer secure them; and (3) on your request, tell you how each recipient's safeguards are established, an outline of the protections, how we verify them, the recipient's country, and any impediments and our response (APPI Article 28(3); PPC Rules Article 18). The countries where your account data is stored, and our security measures, are described in §11 and §15 and available on request.

13.6 United Arab Emirates / Gulf

If you are in the **United Arab Emirates**, UAE Federal Decree-Law No. 45 of 2021 (the "UAE PDPL") applies, and it restricts transfers of personal data outside the UAE. The UAE has not yet issued the PDPL's Executive Regulations, so no adequacy approvals under Article 22 of the PDPL have been published. Where we receive personal data of individuals located in the UAE, we therefore rely on **Article 23(1) of the PDPL, and principally on Article 23(1)(a)**: our agreements - including our Data Processing Agreement and its contractual safeguards - obligate the recipients of the data to apply protections, measures and controls consistent with the requirements of the PDPL. Where we act as controller of your account data, we additionally rely on **Article 23(1)(d)**, because the transfer is necessary to perform the contract under which you receive the Service. When the Executive Regulations are issued, we will review this basis and update this section where required. We honour your right to object to processing for direct-marketing purposes.

13.7 Canada (including Quebec)

If you use the Service from Canada, Canada's federal private-sector privacy law (**PIPEDA**) applies to your account data, because we collect it in the course of the commercial account relationship with you or your organisation. If you are in Quebec, the Quebec **Act respecting the protection of personal information in the private sector** (as amended by "Law 25") also applies.

- **Who is responsible.** Aidealy Ltd. is responsible for your account data, including when our service providers process it for us; we use contracts requiring them to protect it to a comparable level. For Quebec, our **person in charge of the protection of personal information** (responsable de la protection des renseignements personnels) is our founder, acting as **Privacy Officer**, reachable at privacy@aidealy.ai.
- **Your rights.** You may ask us to confirm what account data we hold about you, access it, and correct it if inaccurate or incomplete, and you may withdraw consent at any time (subject to legal or contractual restrictions and reasonable notice). We answer access and correction requests from Quebec within **30 days**.
- **Processing outside Canada and Quebec.** Your account data is stored outside Canada - in our EU or US region (§11) - and handled by the providers in §10. While your data is in another jurisdiction, it may be accessible to the courts, law enforcement, and national-security authorities of that jurisdiction. For Quebec residents, we carry out the privacy assessment Quebec law requires before communicating personal information outside Quebec.
- **No solely-automated decisions.** We do not make decisions about you based solely on automated processing; the AI features support decisions made by people (§6).
- **Marketing.** Marketing email to Canadian addresses is sent only with your express opt-in, identifies us, and includes an unsubscribe honoured within **10 business days** (CASL) - see §7.
- **Complaints.** Contact us first at privacy@aidealy.ai. You may also complain to the **Office of the Privacy Commissioner of Canada** (priv.gc.ca) and, in Quebec, to the **Commission d'accès à l'information** (cai.gouv.qc.ca), including by a "demande d'examen de mécontentement" within 30 days of our answer (or refusal, including a deemed refusal) to an access or correction request.

13.8 Australia and New Zealand

Australia. The Privacy Act 1988 (Cth) and the Australian Privacy Principles (APPs) may apply to our handling of your account data when we carry on business in Australia. Whatever our formal status under that Act, we handle Australian users' account data consistently with the APPs:

- **Access and correction:** you may request access to, and correction of, the account data we hold about you using the contact in §1. We respond within a reasonable period (as a guide, no more than **30 days**) and do not charge you for making a request. If we refuse, we will tell you why in writing, and you may ask us to attach a statement of your position to the record.
- **Where your organisation provided your details:** when your account administrator invites or provisions you (§4), we receive your details from your organisation; this Policy (with the in-product notices) is how we make you aware of that collection, its purposes, and your rights.

- **Overseas disclosure:** your account data is stored in our EU or US region (§11) and handled by the providers in §10. We take reasonable steps, including contractual safeguards, to ensure these overseas providers handle your data consistently with the APPs, and we remain accountable for their handling of it.
- **Marketing and automated decisions:** marketing email is sent only with your express opt-in, with a working unsubscribe honoured within **5 working days** (§7); and we do not use your account data in automated decision-making that has significant effects on you - the AI features support decisions made by people (§6).
- **Complaints:** please complain to us in writing first (we will respond within about 30 days); if you are not satisfied, you can complain to the **Office of the Australian Information Commissioner (OAIC)** at oaic.gov.au.

New Zealand. The Privacy Act 2020 may apply to us as an overseas agency carrying on business in New Zealand. This Policy, together with the notices at the point of collection, tells you what we collect, why, who receives it, and that Aidealy Ltd. (Hamidron 1, Herzliya 4654110, Israel) collects and holds it - including where we receive your details from your organisation rather than from you.

- **Your rights:** you may ask us to confirm whether we hold account data about you, to access it, and to correct it (or attach a statement of correction), using the contact in §1, normally at no charge.
- **Overseas disclosure:** your account data is processed outside New Zealand (our EU or US region and the providers in §10). We rely on agreements requiring those providers to protect your data with safeguards comparable to New Zealand's Privacy Act.
- **Marketing:** we only email you marketing if you opted in; every message identifies us and includes a free, working unsubscribe honoured within **5 working days**.
- **Complaints:** please raise any concern with us first; if we cannot resolve it, you can complain to the **Office of the Privacy Commissioner** (privacy.org.nz).

14. How long we keep your data

We keep account personal data for as long as your account is active. After your account is closed, we keep **only** what we are required to retain by law (for example, tax and accounting records - many of which are held by Paddle as Merchant of Record) for the period the law requires, together with the **legal-agreement acceptance records and contest-review notes described in §9** (each kept as evidence to establish, exercise or defend legal claims; the AI-interaction audit log in §9 is **not** in this list - it is part of your organisation's environment and is deleted at offboarding), the already-written **billing-evidence snapshot records described in §4 and §9** (pseudonymous; account closure does not delete them - each daily record simply ages out on its own 730-day schedule) and the limited **account-security records** kept on the security-log schedule in our **Data Retention & Deletion Policy**, and we **delete the remaining account personal data within 24 months** of account closure, keeping it in the meantime only while we still need it (for example, to deal with an open dispute).

Your organisation's **work data** is governed by the **Data Processing Agreement** and is kept for the duration of the subscription; each stored copy is deleted **five (5) years after it is written** (a copy re-written when we re-organise our storage starts a new five-year period; a deleted copy's residual storage version is purged within about ninety (90) days after that), the analytics graph database holds derived records for the duration of the subscription, and we review these stores at least once a year and delete records that are no longer needed. The full schedule is in our **Data Retention & Deletion Policy**. On termination your live access to the Service ends; for **thirty (30) days** after termination your organisation may **request a data export**, and we prepare and deliver it in **CSV/JSON** format - covering your organisation's data including derived metrics and scores (raw git inputs are not included, as during the subscription) - within **ten (10) Business Days** of the request (a "Business Day" is a day other than Friday, Saturday, or a public holiday in Israel), the one standard export run at no charge; after that window we **permanently delete** your work data from our live systems and decommission your environment - including any per-customer record of AI

prompts and responses, and the **AI-conversation history stored at our AI provider**, which we delete through the provider's interface before the environment is torn down (§6). Any residual copy that remains in our **routine backups** is **kept beyond use** - access-restricted, never restored into active use - and is removed when that backup cycle completes, as described in our **Data Retention & Deletion Policy**. We keep only **aggregated statistical data that no longer relates to any identifiable person** (and is therefore no longer personal data - unlike the individual-level de-identification in §9, which remains personal data) and the limited **account-security records, legal-agreement acceptance records, contest-review notes and billing-evidence snapshot records (§4, §9)** described in our **Data Retention & Deletion Policy** (held in our own systems, not as part of your environment, which is destroyed at offboarding - the AI-interaction audit log is part of that environment and is deleted with it, notwithstanding its seven-year per-entry period), which sets out the full schedule.

When we destroy personal data, we do so by permanently deleting the electronic records, subject to the handling of backup copies described in this Policy.

All deletion and retention timelines in this section yield, to the narrowest extent and for no longer than necessary, to legal preservation obligations that bind us (such as a litigation hold or a statutory preservation order).

15. Security

We use appropriate technical and organisational measures to protect personal data, including **encryption in transit (HTTPS/TLS)** and **encryption at rest** (including per-customer encryption keys), **access controls** limiting who can see the data, separate EU and US regions, and use of reputable providers who maintain their own security safeguards.

If a personal-data breach affects your data, we will notify you and/or the competent supervisory authority where, and within the timeframes, required by the law that applies to you.

16. Data protection officer

Given the nature of our processing, we are appointing a **Data Protection Officer (DPO)** and are committed to maintaining one. You can reach our DPO and privacy team at privacy@aidealy.ai.

17. Children

The Service is a business product, is not directed to children, and we do not knowingly collect personal data from children under 16.

18. Changes to this Policy

We may update this Policy. We will post the new version here with a new "Last updated" date and, where required, notify you.

19. Contact and complaints

Questions or complaints about this Policy or your data: privacy@aidealy.ai, Aidealy Ltd., Hamidron 1, Herzliya 4654110, Israel. For concerns specifically about how we develop or use AI - for example the fairness of AI-assisted evaluations - you can also write to ai-concerns@aidealy.ai (see §6). You may also contact your local supervisory authority (see §13).