

Aidealy AI Addendum

Effective date: 2026-09-06 · **Version:** 1.4.0 · **Last updated:** 2026-09-06 This AI Addendum (this "**AI Addendum**") supplements, and forms part of, the Aidealy Master Subscription Agreement / Terms of Service (the "**Agreement**") between **Aidealy Ltd.** ("**Aidealy**", "**we**", "**us**", "**our**") and the customer that subscribes to the Service ("**Customer**", "**you**", "**your**"). It sets out terms and disclosures that apply specifically to the **artificial-intelligence features** of the Service. Capitalised terms used but not defined in this AI Addendum have the meanings given to them in the Agreement.

This AI Addendum **references, and does not replace**, the related terms in the Agreement, the Acceptable Use Policy (the "**AUP**"), the Data Processing Agreement (the "**DPA**"), and the Privacy Policy. Where it refers to one of those documents, that document governs the referenced subject. If there is a conflict, the **order of precedence in the Agreement applies**, and this AI Addendum governs only matters specific to the AI features of the Service.

1. The AI features and the AI providers we use

1.1 What the AI features do. Parts of the Service use artificial-intelligence and machine-learning models to provide analytics and natural-language features over Customer Data - including automated code-quality analysis and scoring of source code, the computation of team- and codebase-level development metrics, and a natural-language assistant ("chat") that lets Authorised Users ask questions about Customer's own code and analytics.

1.2 The AI providers. Aidealy provides these features using a small number of specialist AI sub-processors, which process Customer Data **only to provide the Service to Customer** and under terms consistent with this AI Addendum and the DPA:

AI sub-processor	Role in the Service	Processing region(s)
Anthropic, PBC	AI model provider - the AI features of the Service (analysis and scoring of source code, classification of developers' typed messages, and the natural-language analytics assistant). Aidealy routes each AI step to either of its two listed AI model providers and may change that routing at any time, including as a fallback, without adding a sub-processor; each provider's own retention terms are disclosed in DPA Section 10.	United States (no EU option; stores API data in the United States)
OpenAI OpCo, LLC	AI model provider - the AI features of the Service (analysis and scoring of source code, classification of developers' typed messages, and the natural-language analytics assistant). Aidealy	EU/US (per project or per request); processing at OpenAI's default endpoint (a United States recipient; OpenAI's EU regional option is supported in Aidealy's software but not switched on)

AI sub-processor	Role in the Service	Processing region(s)
	routes each AI step to either of its two listed AI model providers and may change that routing at any time, including as a fallback, without adding a sub-processor; each provider's own retention terms are disclosed in DPA Section 10. Stores the analytics assistant's conversation history server-side where the assistant is routed to it.	
LangChain, Inc. (LangSmith)	Tracing and observability of the AI features (operational quality; integration wired but dormant - no data flows until activated)	United States

1.3 The authoritative, current list. The table above identifies the AI providers, any of which may perform any of the AI steps it describes. The **authoritative, maintained list** of Aidealy's sub-processors - including each one's identity, purpose, location, and the cross-border-transfer safeguards that apply to it - is the **Sub-processor List referenced in the DPA (DPA Annex III and Annex IV)**. Aidealy may change its AI providers in accordance with the change-of-sub-processor process in the DPA (advance notice and a right to object); that process, not an amendment to this AI Addendum, governs such changes. For clarity, **routing among the AI providers already on the Sub-processor List satisfies that process without further notice** (Section 1.6); the DPA's notice-and-objection mechanics attach to the **addition or replacement of a sub-processor entity**, not to routing among already-authorized entities.

1.6 Routing among listed AI providers; new providers; Aidealy-operated models. (a) Aidealy may **route any request among the AI providers identified on the Sub-processor List at any time** - including automated failover when a provider is unavailable, and routing for quality, capacity, or cost optimisation - **without further notice**, provided each provider processes Customer Data only within its disclosed role and safeguards. For clarity, either AI model provider identified in Section 1.2 may perform any of the AI steps described in that table (the analysis and scoring of source code, the classification of developers' typed messages, and the natural-language analytics assistant), and this AI Addendum does not commit Aidealy to any particular allocation of those steps between them: Aidealy chooses the routing and may change it at any time, including as a fallback when one provider is unavailable, and such a change is routing among listed providers under this Section 1.6(a), not the addition or replacement of a sub-processor. (b) A **new external AI provider** (one not on the Sub-processor List) is a new sub-processor: the DPA's advance-notice and objection process applies before it processes Customer Data, and the DPA's emergency-replacement provision governs a sudden, unplanned cutover. (c) **AI models Aidealy itself operates on its own infrastructure** (including open models Aidealy hosts, and models invoked through an infrastructure service inside Aidealy's own cloud account where the model provider has no access to Customer Data) are **Aidealy's own processing, not a sub-processor**; Aidealy will keep the list of AI providers, and the safeguards disclosed for each of them, in this Section and the Sub-processor List accurate as its architecture evolves.

1.4 Model and provider changes; no score comparability. The AI features are built on machine-learning models that Aidealy and its AI providers update and replace over time. **A change of an underlying model, model version, or AI provider is not, by itself, a material reduction of the core**

functionality of the Service for the purposes of the "Changes to the Service" provision of the Agreement (Section 2.5 of the Agreement), and provider changes are governed by the DPA process described in Section 1.3. If such a change **materially reduces the functionality of the AI features** (including a material reduction of a core AI capability described in Section 5.2, not only of the AI features taken as a whole), Aidealy will notify Customer, and Customer may terminate the affected subscription within **thirty (30) days** of the notice and receive a pro-rata refund of pre-paid, unused fees for the terminated portion (executed by Paddle as Merchant of Record for Paddle transactions); this is Customer's remedy for such a change. Customer acknowledges that **model behaviour, scores, and output style may change over time and across model versions**, and that Aidealy does **not warrant that scores, metrics, or other Outputs are comparable period-over-period or across model versions**.

1.5 If an AI provider becomes unavailable. If a third-party AI provider used for the AI features becomes unavailable, terminates or materially restricts Aidealy's access, or materially changes the terms on which it provides its models, Aidealy is **excused from performance of the affected AI features to the extent, and for so long as, that event prevents them from operating**, and such non-performance is not a breach of the Agreement or of the SLA. Aidealy will in that case **use commercially reasonable efforts to substitute a materially equivalent model or provider without undue delay** (following the DPA sub-processor process where the substitution engages it), and will keep Customer informed of material developments. If Aidealy cannot restore materially equivalent AI-feature functionality within **thirty (30) days** of the start of the event, Customer may terminate the affected subscription and receive a pro-rata refund of pre-paid, unused fees for the terminated portion (executed by Paddle as Merchant of Record for Paddle transactions). Where a subscription is terminated under this Section 1.5 after the affected AI features have ceased to function, the pro-rata refund is **computed from the date the affected AI features ceased to function** (or, if earlier, from the date that event began to prevent them from operating), not from the date of termination. This Section does not limit Section 1.4 or the Agreement's force-majeure provision.

Upstream cost pass-through. If a **documented increase in the costs charged to Aidealy by its upstream AI providers** makes continued provision of the affected AI features at the current pricing **commercially impracticable**, and **no materially equivalent cheaper substitution is reasonably available** (including by routing under Section 1.6), Aidealy may notify Customer of a surcharge on the affected subscription, **capped at the portion of the documented cost increase attributable to Customer's subscription**, and will provide Customer, with the notice or on request, reasonable supporting documentation of the increase and its attribution (redacted as needed to protect confidential terms), taking effect **no sooner than thirty (30) days** after the notice. Within that notice period, Customer may instead **terminate the affected subscription and receive a pro-rata refund of pre-paid, unused fees** for the terminated portion (executed by Paddle as Merchant of Record for Paddle transactions). If, by the end of the notice period, Customer has neither agreed to pay the surcharge nor terminated, Aidealy may terminate the affected subscription with the same pro-rata refund. The surcharge **lapses when, and to the extent that, the underlying upstream cost increase reverses**.

2. No use of Customer Data to train AI models

Aidealy's commitment that **Customer Data is not used to train, fine-tune, or improve any AI or machine-learning model** - whether Aidealy's own or a third party's - is set out in **Section 4.3 of the Agreement**, and is repeated here because it is central to the AI features:

- Aidealy does not use Customer Data (including source code, prompts, and developer-activity data) to train, fine-tune, or improve any model.
- Aidealy does not authorise its AI sub-processors to use Customer Data to train their models, and **does not enable** any optional model-training feature those providers may offer. Aidealy engages the AI providers above under terms in which **each provider contractually commits not to train its models on customer content; Aidealy has not agreed, and will not agree, to any use of customer content for training**.

- This commitment is **separate from, and is not satisfied by**, any short retention period a provider may apply to operate the service or detect abuse (for example, a provider's limited API-input retention). Retention for service-provision is not training.

Aidealy makes this commitment as to the substance it controls; it does not warrant the future terms or conduct of its AI providers. The provisions of the DPA govern Aidealy's processing of personal data within Customer Data.

3. Ownership of inputs and outputs

As between the parties, **Customer owns the inputs it submits to the AI features and the outputs the AI features generate for Customer ("Outputs")**, as set out in **Sections 4.1 and 8.2 of the Agreement** - which also contain Aidealy's **present assignment to Customer of any right Aidealy may hold (if any) in the Outputs, and its non-assert commitment**, together with the parties' acknowledgment that some Outputs may attract no copyright or other exclusive right under applicable law. Aidealy holds only the limited licence to host and process Customer Data needed to provide, secure, support, and maintain the Service. Given the nature of generative AI, **Outputs may not be unique**: the AI features may generate the same or similar output for other customers from different inputs. Customer's ownership, as between Customer and Aidealy, is of the Outputs generated for Customer, and does not extend to any similar output the AI features independently generate for another customer.

4. AI transparency

4.1 You are interacting with an AI system. Where an Authorised User interacts directly with the Service's natural-language assistant, Aidealy will make clear, at or before the point of that interaction, that the user is **interacting with an artificial-intelligence system** and not with a human. This reflects the transparency duty on providers of AI systems intended to interact directly with people under the EU AI Act (Article 50(1)), which applies from 2 August 2026. This AI Addendum, together with the in-product notice, is also the **advance notice that the Service uses generative artificial intelligence** required for users in South Korea by the AI Framework Act (Article 31(1)); that Act's implementing decree permits the notice to be given in the contract or terms of service.

4.2 AI-generated output. Outputs of the AI features are **generated by machine-learning models** and are presented as AI-assisted. The AI features produce analysis and decision-support about Customer's own code and metrics; they are not designed to create stand-alone synthetic audio, image, or video, or text passed off as authored by a person. Aidealy identifies AI-assisted output in the Service by a persistent, **human-readable** notice displayed with the chat interface, which states that the narrative explanations are AI-written (while metrics and charts are computed by deterministic code from Customer Data and Aidealy's stored analytics, which include the AI-assisted code-quality scores described in Section 1.1), consistent with the generative-AI output-labelling duty under Korea's AI Framework Act (Article 31(2) and its implementing decree, which accepts a human-recognisable indication) as well as the EU transparency rules described above.

4.3 Our role, and your own transparency duties. Aidealy operates the AI features as their provider. The Service does **not** perform "emotion recognition" or "biometric categorisation" **within the meaning of the EU AI Act** - that is, it does not infer emotions or categorise individuals **on the basis of biometric data**. The Service's AI features may classify a user's typed messages **from their text content only** (the message itself, read together with up to three of that user's previous messages in the same session as context; not from biometric data such as voice, facial expressions, or keystroke dynamics), labelling the part of the system a message is about, the kind of work being asked for, any quality concern it raises, and its apparent tone (its **sentiment**); under the European Commission's Guidelines on prohibited AI practices (C(2025) 5052 final), text/content-based sentiment analysis is **not** biometric-based emotion recognition and falls **outside** the Article 5(1)(f) prohibition. The Service does **not** generate deepfakes or publish AI-generated text to inform the public on matters of public interest. If Customer's own use of the Service, or of its Outputs, would trigger a transparency or disclosure duty that falls on Customer - for example, as the deployer of

an AI system, or under data-protection law - **meeting that duty is Customer's responsibility**, consistent with Section 6 of this AI Addendum and the AUP.

5. Intended purpose, capabilities, and limitations

This Section describes the AI features so that Customer can interpret and use their Outputs appropriately. It is **designed to align** with the instructions-for-use content for high-risk AI systems under the EU AI Act (Article 13); the **expected-accuracy metrics and the remaining technical content will be added as Aidealy's high-risk-provider compliance programme produces them**.

5.1 Intended purpose. The AI features are intended to provide **software-development analytics** - including team- and codebase-level metrics and metrics about individual contributors - and a natural-language assistant over Customer's own code and development data, as **decision-support, subject to human oversight**, for engineering leaders and teams. The intended purpose **includes use to evaluate the performance or behaviour of individual developers, and use as decision-support - subject to human oversight - for decisions about an individual's remuneration, bonuses, and pay progression**; where the AI features are used for any of those purposes they are a **high-risk AI system** under the EU AI Act, and Customer's responsible-use (deployer) obligations in **Section 4 of the AUP** (including, for pay decisions, the pay-transparency and equal-pay duties described in Section 4.4(k) of the AUP) and Section 6 apply (see Section 6). The effort-hours metric is excluded from the pay use: see Section 5.6.

5.2 Capabilities. The AI features can, among other things, analyse and score source code for quality signals, summarise and surface patterns in development activity at team and codebase level, compute per-developer contribution metrics - including an estimate of the time each developer actively spent on their changes (their "effort hours"), computed from IDE and git activity signals only, and the split of their contributed code between AI-assisted and human-written - and answer natural-language questions about Customer's own code and analytics.

5.3 Limitations. As described in **Section 9.3 of the Agreement**, AI Outputs are **probabilistic and may be inaccurate, incomplete, or unsuitable for a particular purpose**. They can reflect limitations or biases in the underlying models or in the input data, and they should not be relied on without Customer's own review. Sentiment labels may be **less reliable for non-native English speakers**. Activity-based metrics can **read lower for individuals working part-time or atypical schedules**; where an activity-derived metric feeds a pay outcome, that skew can produce a **less favourable pay result for part-time and atypical-schedule workers**, which Customer must be able to justify on objective, non-discriminatory criteria under applicable equal-pay law - under which the burden of proving the justification can fall on Customer as the employer. The AI features are not a substitute for professional judgement. The effort-hours metric is computed from **IDE and git activity signals only**, in two variants - **gross, and net of extended idle periods within the observed IDE activity**: it does not and cannot measure an individual's overall working time, presence, meetings, reviews, design work, or any activity outside those signals. Customer must not use it as a measure of an individual's total working time, presence, or inactivity, and must not use it as a criterion for determining an individual's pay, bonuses, or pay progression (Section 5.6).

5.4 Human oversight. The AI features are designed to be used with **meaningful human review**; Section 6 sets out Customer's human-oversight responsibilities.

5.5 Provider and representatives. The provider of the AI features is **Aidealy Ltd.**, Hamidron 1, Herzliya 4654110, Israel (company registration number: 517367066); contact **legal@aidealy.ai**. Aidealy's EU and UK representatives for data-protection purposes are **being appointed** and will be identified in the Privacy Policy on completion; in the meantime, data-protection enquiries may be directed to **privacy@aidealy.ai**.

5.6 Known inappropriate uses. As part of the Instructions for Use, Aidealy states the known inappropriate or harmful uses of the AI features. The AI features are not designed or intended to be used, and must not be used: as the **sole basis** for a decision that produces legal effects concerning an individual or similarly significantly affects an individual (Section 6); for any practice prohibited

under **Section 5 of the AUP**, including inferring the emotions of an individual on the basis of biometric data, social scoring, or biometric categorisation; to **infer an individual's health or mental-health condition**; on **biometric input signals** (the AI features process text and code content only); or otherwise outside the intended purpose described in Section 5.1. In addition, and as with the working-time restriction in Section 5.3: **the effort-hours metric must not be used as a criterion for determining an individual's pay, bonuses, or pay progression**. It measures observed IDE and git activity signals only, and Section 5.3 discloses that activity-based metrics can read lower for individuals working part-time or atypical schedules; used as a pay criterion, it would import that skew into pay outcomes Customer must justify under equal-pay law. This restriction does not limit the intended pay-decision-support use described in Section 5.1: the Service's other analytics may inform pay decisions subject to Section 6 and the AUP; the effort-hours metric specifically may not.

5.7 Information for Customer's own notices; updates. This AI Addendum - in particular this Section 5 - together with the further information Aidealy makes available **on request**, provides the information Customer reasonably needs for its own notices and disclosures to individuals about its use of the AI features - including the **trade name** under which the AI features are provided (they are provided as part of the Service under the trade name "**Aidealy**") and any **version identifier** Aidealy assigns to them; the identity of the provider (Section 5.5); a description of the AI features, their intended purpose and their role in decisions Customer makes (Sections 5.1 to 5.3); the categories of personal data they analyse or process, the **sources** of that personal data (the Customer-connected data sources described in Annex I to the DPA, as connected and configured by Customer), and **how that data is assessed** in reaching an Output (Sections 5.1 to 5.3); and the **explanation inputs Customer needs to give an affected individual a clear and meaningful explanation of the role the AI features played in a decision about them** (the elements of Section 5.1 to 5.3 plus, on request, a description of the main analytics elements that fed the relevant Output). Where Customer must answer an individual's request under applicable pay-transparency or equal-pay law for the **criteria used to determine their pay or pay progression** (see Section 4.2(i) of the AUP), that same information serves the request so far as it concerns the Service: which analytics elements fed the Output Customer used, and what the Output does and does not measure (Sections 5.1 to 5.3, including the measurement-scope and known-limitation statements of Section 5.3); the pay criteria themselves, and their justification, are Customer's as the employer. That information also includes, **to the extent known to Aidealy**, a description of the **categories of data, including personal data, used to train the AI features' underlying models**: Aidealy does not use Customer Data to train AI models (Section 2); the AI features use third-party foundation models from the AI providers identified in Section 1, and the categories of data used to train those models are known to Aidealy only as those providers publish them - Aidealy will identify or provide that published information on request. The individuals covered by Customer's notices are **not limited to current employees**: the Service's ingested repository and development history can identify former personnel, contractors, and external contributors (such as open-source committers), and the information in this Section serves Customer's transparency duties to that whole class (see Section 4.2(d) of the AUP). This information also includes a channel through which an individual can raise a concern about the AI features - including their fairness - directly with Aidealy at **ai-concerns@aidealy.ai**; Customer may include this address in the information it gives to the individuals subject to the AI features (in particular the workers it evaluates, who may have no other contact point with Aidealy), and using it does not replace the human-oversight, contest, and worker-information protections Customer must provide under Section 6 and Section 4 of the AUP. **Identity handling**: Aidealy does not disclose the identity of an individual who raises a concern through that channel to Customer (or to any other organisation) without the individual's agreement, except where the law requires disclosure or where the concern cannot be assessed or acted on without contacting the organisation - in which case Aidealy tells the individual so before proceeding. Aidealy will **acknowledge a report made to that channel within ten (10) Business Days** ("Business Day" has the meaning given in the SLA) of receiving it; beyond that

acknowledgement, the channel remains an invitation to report, and Aidealy does not commit to any further service level or to a particular investigation, outcome, or response timeline. Aidealy will inform Customer of **material updates, intentional and substantial modifications, and changes to the intended use of, limitations of, or risk mitigation for** the AI features within a reasonable time - for example, by release notes together with notice to the Customer's administrator contact.

5.8 Beta AI features. An AI feature identified as beta, preview, early access, experimental, pilot, or proof-of-concept is a **Beta Service** under the Agreement (Section 2.6 of the Agreement): it is provided as-is for evaluation and is **outside the capability, limitation, and accuracy descriptions of this Section 5 and Section 7**, which describe the generally available AI features. The transparency commitments in Section 4, and the responsible-use and human-oversight requirements of Section 6 and the AUP, **continue to apply to a beta AI feature that is deployed to Customer, regardless of its label.**

5.9 Audit-log export; provider information; bias-audit support.

(a) **AI-interaction audit log and AI-system event logs.** On Customer's reasonable request within the applicable retention window, Aidealy will provide Customer with a **read-only export of the entries of the AI-interaction audit log, and of the other automatically generated AI-system event logs, relating to Customer's own tenant**, as those logs are identified in Section 7 of the Data Retention & Deletion Policy. Before those logs are deleted as part of offboarding, Aidealy will **offer Customer such an export**. Exports are read-only extracts: the underlying logs are maintained as write-protected, tamper-evident records, an export does not modify them, and Aidealy does not edit them.

(b) **AI-provider information.** On Customer's reasonable request, Aidealy will inform Customer of the **AI providers, and the model families, currently serving Customer's tenant**, and will provide summaries of the upstream provider documentation that Aidealy holds and is permitted to share (the authoritative provider list remains the Sub-processor List, Section 1.3).

(c) **Bias-audit support.** Where Customer is legally required to conduct or commission a **bias audit or similar fairness assessment** of its use of the AI features - including a **pay-gap report, equal-pay analysis, or joint pay assessment** in which Customer must examine or justify a pay criterion derived from the AI features - Aidealy will, on Customer's reasonable request: provide Customer with **Customer's own Output and score distributions** and with information about the methodology of the AI features reasonably needed for the audit or assessment, and **reasonably cooperate with Customer's independent auditor**, subject to confidentiality obligations protecting Aidealy's Confidential Information. Aidealy does not hold demographic or other protected-attribute data about the individuals Customer evaluates, and does not hold Customer's pay, salary, or compensation data at all, and cannot provide data it does not hold; obtaining such data, and the audit or assessment itself, remain Customer's responsibility as the employer and deployer (Section 6.3 and Section 4 of the AUP). Support beyond the reasonable-effort assistance described in this Section 5.9 is available as paid professional services under the Agreement (Section 2.9 of the Agreement).

5.10 Message-classification labels, including sentiment: what Customer sees; objections. Where the AI features classify typed messages (Section 4.3: the part of the system, the kind of work, any quality concern, and the apparent tone or sentiment), Customer's authorised viewers see the resulting **per-message labels only: the classification analytics do not display the raw text of the classified messages**. (The AI-interaction audit log, which separately records the content of AI analytics (chat) queries and answers as a security record, remains available as described in Section 5.9(a).) The Service has **no per-user or tenant-wide control that stops message classification, or its sentiment label, alone** while other collection continues, and Aidealy does not commit to build one. If an individual's objection to this processing is upheld (including a data-protection objection that Customer, as controller, routes to Aidealy), it is honoured by **excluding that individual from collection entirely** - for example, by deactivating that individual's seat or removing the Extension from that individual's machine - not by a classification-only stop.

6. Responsible high-risk worker-evaluation use; human oversight

6.1 Responsible high-risk use. The AI features may be used to monitor or evaluate the performance or behaviour of individuals, in which case they are a **high-risk AI system** under the EU AI Act and Customer is a **deployer**. Customer's deployer obligations - using the Instructions for Use, assigning competent human oversight, informing workers' representatives and the affected individuals that they are subject to the system, carrying out a data-protection impact assessment, and the other requirements - are set out in full in **Section 4 of the AUP** (and the hook in Section 3.3 of the Agreement), and apply to use of the AI features. Laws outside the EU that regulate AI in employment decisions may also bind Customer; the jurisdiction-specific position - including where a law places its duties on Aidealy rather than on Customer - is summarised in **Section 4.4 of the AUP**. This AI Addendum does not narrow them.

6.2 Meaningful human review. As set out in **Section 6 of the AUP**, because AI Outputs may be inaccurate, Customer will apply meaningful human review before relying on them, and will **not use any Output as the sole basis for a decision producing legal effects concerning an individual, or that similarly significantly affects an individual**. This requirement is in addition to, and does not limit, Customer's own obligations under data-protection law for automated decision-making and profiling (including, in the EEA, GDPR Article 22 and, in the UK, 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)).

6.3 Customer is the employer; no agency. As set out in Section 3.4 of the Agreement, **Customer is the sole decision-maker for all employment and personnel actions** informed by the AI features. Aidealy provides the AI features as decision-support under Customer's human oversight (Sections 6.1 and 6.2); Aidealy **exercises no delegated authority over, and takes no part in, any employment or personnel decision**, and does not act as Customer's agent (or as an agent of any employer) in making or communicating such decisions. Worker notices and consultation duties, any legally required bias audits or impact assessments, and the other employer- and deployer-side duties described in Section 4 of the AUP are **Customer's responsibility**. This Section allocates responsibility between the parties; it does not limit Section 12.3 of the Agreement or any liability of Aidealy that applicable law does not permit to be allocated away.

6.4 Aidealy's provider obligations under the EU AI Act. Aidealy provides the AI features as their **provider** under the EU AI Act where they are used as described in Section 5.1. **When and as the corresponding high-risk provider obligations under the EU AI Act apply to the AI features, Aidealy will have completed the conformity steps applicable to it** - and will provide Customer, as deployer, with the documentation, information, and cooperation Customer reasonably needs to meet its own deployer obligations for its use of the AI features (building on the Instructions for Use in Section 5 and the information channel in Section 5.7). Aidealy's high-risk-provider compliance programme is maintained outside this Addendum; its status is available to Customer on reasonable request.

7. AI output: accuracy, reliance, and intellectual property

7.1 Accuracy and no sole reliance. Except for the express warranty in **Section 9.2 of the Agreement**, and as stated in **Sections 9.3 and 10 of the Agreement**, the AI features and their Outputs are provided **"as is"**, and Aidealy does not warrant that any Output is accurate, complete, or fit for a particular purpose. Customer is responsible for evaluating Outputs and deciding whether and how to use them.

7.2 No warranty of non-infringement of Outputs. Because generative-AI Outputs are produced by models from large bodies of data, Aidealy **does not warrant that Outputs do not infringe the intellectual-property or other rights of any third party**, and Customer is responsible for reviewing Outputs before relying on or distributing them.

7.3 Outputs are outside the IP indemnity. Aidealy's intellectual-property indemnity in **Section 12.1 of the Agreement** covers Customer's authorised use of the Service; it **does not cover any claim**

arising from Outputs (including a claim that an Output infringes a third party's intellectual-property rights), which are excluded from that indemnity in the same way as Customer Data and Customer's own use. The allocation of responsibility for Outputs in this Section reflects that Aidealy is an application-layer service built on third-party AI models: **any intellectual-property protections Aidealy's AI providers offer run to Aidealy as their direct customer, on those providers' own terms and subject to their exclusions, and are not passed through to Customer under this Addendum.**

7.4 Corrective action for discovered output defects. If Aidealy discovers, after delivery, that a **defect** - whether in an upstream model, in Aidealy's own scoring or analytics code, or in the data pipeline - caused Outputs already delivered to Customer to be **materially wrong**, and those Outputs are of the kind used to evaluate or make decisions about individuals (Section 5.1), Aidealy will **notify the affected customers without undue delay after the discovery, identifying the affected period and the nature of the defect**, so that Customer can re-examine decisions it made in reliance on Outputs from that period (applying its human-oversight duties under Section 6 and the AUP). Aidealy will also cooperate with Customer's reasonable requests for the information it needs to assess the impact, and will perform the incident- and serious-incident duties that apply to Aidealy as a provider when and as the corresponding obligations under the EU AI Act apply. **This Section is not triggered by deliberate improvements:** planned changes, enhancements, or optimisations of models, scoring formulas, or analytics - including changes that alter scores going forward - are governed by Section 1.4 (model and provider changes; no score comparability), and require no notice under this Section before or after shipping. A notice under this Section is given so that Customer can re-examine affected decisions: it is **not an acknowledgement or admission of fault or liability** by Aidealy, and it is Aidealy's **Confidential Information** under Section 7 of the Agreement.

8. Liability for AI errors

Liability arising out of or relating to the AI features - including any liability for an inaccurate, incomplete, or infringing Output - is subject to the **limitation of liability in Section 11 of the Agreement**. This AI Addendum creates **no separate, additional, or higher liability cap** for the AI features. For the avoidance of doubt, the uncapped and excepted liabilities in **Section 11.3 of the Agreement** - including breach of the **no-training commitment (Section 4.3)** arising from **Aidealy's own or authorised conduct** (where a breach consists of an upstream AI provider's use of content in breach of that provider's own commitments, without Aidealy's authorisation, it is governed by the liability cap in Section 11.2 of the Agreement, as Section 11.3 of the Agreement provides) and the **intentional, or knowingly unauthorised, disclosure or use of Confidential Information (Section 7)** - continue to apply.

9. Relationship to the Agreement; changes

9.1 Order of precedence. This AI Addendum is part of the Agreement and is subject to it. If there is a conflict between this AI Addendum and another document of the parties' agreement, the **single ranked order of precedence in Section 15 of the Agreement applies**: under it, a signed Order Form and the DPA (with its SCCs, on personal data) and the Extension EULA (for the IDE extension) rank above this AI Addendum, **this AI Addendum ranks above the Agreement for matters specific to the AI features**, and the Agreement otherwise prevails.

9.2 Changes. Aidealy may update this AI Addendum from time to time as the AI features and the applicable law evolve. Because this AI Addendum is incorporated into, and forms part of, the Agreement, changes to it are made, notified, and take effect as set out in the Agreement's **"Changes to this Agreement; version binding; records"** provision (Section 15 of the Agreement): the dated version of this AI Addendum in effect at the start of a subscription term governs that term; a change that materially reduces Customer's rights or materially expands its obligations takes effect at the start of Customer's next renewal term, after reasonable advance notice (for example, by email or in-product); and Aidealy maintains dated archives of each published version and will make the version applicable to Customer's subscription available on request.

10. How to contact us

Aidealy Ltd., **Hamidron 1, Herzliya 4654110, Israel**. Questions about the AI features or this AI Addendum, and legal notices: **legal@aidealy.ai**. Privacy questions: **privacy@aidealy.ai**.