

Aidealy Vulnerability Disclosure Policy

Effective date: 2026-09-06 · Version: 1.2.0 · Last updated: 2026-09-06

1. Our commitment

Security is foundational to Aidealy. We build the Service around protecting our customers' most sensitive asset - their source code, their developers' work, and the people behind it - and we welcome the work of security researchers who help us keep it safe.

This Vulnerability Disclosure Policy (the "**Policy**") explains how to report a security vulnerability you believe you have found in Aidealy, what you can expect from us in return, and the rules that keep your research lawful and safe for everyone. It applies to **Aidealy Ltd.** ("**Aidealy**", "**we**", "**us**"), the company behind the Aidealy product and website.

This is a **coordinated vulnerability disclosure** program, modelled on the recognised international practice for receiving and handling vulnerability reports (ISO/IEC 29147 and ISO/IEC 30111) and on the guidance of national cybersecurity authorities. It is **not a paid bug-bounty program** - see §9 (Recognition).

2. Scope

In scope

This Policy applies to the systems Aidealy operates and controls:

In scope	Notes
The Aidealy website - aidealy.ai and *.aidealy.ai	The marketing site and its subdomains
The Aidealy Service - the web application (admin and chat)	The product you sign in to
The Aidealy IDE extension - the VS Code / Cursor extension	Both regional editions
The Aidealy Claude Code collector - the standalone program installed on developers' machines	Including its installer and its update mechanism

Out of scope

The following are **out of scope**, and testing them is **not authorised** by this Policy:

- **Third-party platforms and services we build on** - including our cloud, infrastructure, payment, email, AI-model, source-control and monitoring providers. These have their **own** vulnerability disclosure programs; please report issues in their systems **to them**. We cannot authorise you to test systems we do not own.
- **Social engineering** of Aidealy staff, customers, users or vendors (phishing, pretexting, vishing).
- **Physical** attacks against Aidealy property, offices or personnel.
- **Denial-of-service (DoS/DDoS)**, volumetric, or any resource-exhaustion or load testing.
- Any activity that could **degrade, disrupt or damage** the Service or other users' use of it.

Findings we generally consider non-qualifying

Reports limited to the following, without a demonstrated, realistic security impact, will usually be acknowledged but **not treated as actionable vulnerabilities**: missing security headers or cookie flags; output of automated scanners without a working proof of concept; absence of rate-limiting where no concrete abuse is shown; theoretical or best-practice issues; reports about software

versions alone; and self-inflicted issues (e.g. requiring a fully compromised device or a man-in-the-middle you control).

3. How to report

Email your report to **security@aidealy.ai**.

- You may report **anonymously** if you prefer; note that we will then be unable to follow up with you or acknowledge you.
- We also make this Policy and our security contact discoverable through a machine-readable **security.txt** file at <https://aidealy.ai/.well-known/security.txt> (per RFC 9116).
- If your report contains sensitive details, you may ask us for an encrypted channel and we will arrange one.

4. Safe harbor

We want you to be able to do good-faith security research without fear of legal consequences. When you make a good-faith effort to comply with this Policy, then with respect to systems within scope (§2):

- We consider your research to be **authorised**. Many computer-misuse laws turn on whether access was "without authorization" or "unauthorised" - for example the U.S. Computer Fraud and Abuse Act, the U.K. Computer Misuse Act 1990, the Israeli Computer Law, 5755-1995, and EU Member States' laws implementing the EU directive on attacks against information systems. By this Policy, **we authorise** the testing activities it permits, so that authorised, in-scope research is **not** "unauthorised" access to our systems.
- We will treat your research as **authorised under applicable anti-circumvention laws** (such as the U.S. Digital Millennium Copyright Act) to the extent it is necessary to test our systems in scope.
- It is **exempt** from the restrictions in our Terms of Service and Acceptable Use Policy that would otherwise prohibit security testing, and we waive those restrictions for your in-scope, good-faith research on a limited basis.
- We will **not pursue or support legal action** against you - civil, or a referral for criminal prosecution - for security research carried out in good faith and in compliance with this Policy.

"**Good faith**" means, in line with recognised standards: you access only what you need to identify and demonstrate a vulnerability; you avoid harm to people, to data, and to the availability of the Service; you do not use, retain, or disclose data you encounter beyond what is necessary to report; and your purpose is to improve security, not to extort, defraud, or benefit improperly. Activity that is not in good faith - for example, using a "research" label to cover extortion - is not protected by this Policy.

Limits of this safe harbor - please read. We can only speak for **ourselves**. This Policy:

- does **not** authorise you to access, test, or affect systems, data, or accounts that belong to **third parties** - including our service providers, other customers, or other Aidealy tenants - and it does not waive anyone else's rights. If your testing would touch a third party's systems or data, you need that party's own authorisation; if in doubt, **ask us first**.
- does **not** bind courts, regulators, or public prosecutors, and is **not a guarantee of immunity**. Some protections for security researchers are a matter of prosecutorial discretion or are narrow and condition-laden, and not every country provides a statutory research defence.
- protects only activity that stays **within scope**. Going outside the authorised scope - even with good intentions - can fall back outside this safe harbor and back within those laws.

If a third party brings a claim against you for activity that complied with this Policy, we will, where we lawfully can, take reasonable steps to make clear that your activity was authorised by us.

5. What to include in your report

To help us triage and reproduce the issue quickly, please include where you can:

- a clear description of the vulnerability and the affected system, URL, or component (and version/edition);
- step-by-step instructions to reproduce it, and a proof of concept;
- the impact you believe it has, and any suggested remediation;
- your contact details, if you would like us to follow up or credit you.

6. Rules of engagement

Because the Service processes highly sensitive customer data - source code, AI prompts and responses, and personal data - these rules are conditions of the authorisation and safe harbor in §4. When testing, you **must**:

- stay **within scope** (§2) and comply with this Policy and applicable law;
- use **only your own accounts or test accounts** - never access, or attempt to access, another user's or another organisation's account or data;
- **minimise your interaction** with data: access only the minimum needed to demonstrate the issue, and do **not** copy, store, transfer, modify, or destroy any data;
- **stop immediately and report to us** if you encounter any personal data or any customer content (such as source code) - do not access it further, do not retain it, and tell us in your report so we can handle it. Aidealy processes such data under its Privacy Policy and Data Processing Agreement, and you must not undermine those protections;
- not perform social engineering, physical attacks, or denial-of-service, and not degrade the Service (§2);
- give us a reasonable opportunity to investigate and remediate before any public disclosure (§8).

7. What you can expect from us

When you report in line with this Policy, we will:

- **acknowledge** your report within **5 business days** (a business day is a day other than Friday, Saturday, or a public holiday in Israel; our standard business week is Sunday to Thursday);
- **triage and validate** the report, and let you know whether we consider it a vulnerability we will act on;
- keep you **reasonably informed** of our progress through to remediation; and
- **handle your report confidentially**, and not share your identity or details outside Aidealy without your permission, except where we are required to by law.

Regulatory reporting (EU Cyber Resilience Act). Separately from the commitments above, we assess every reported vulnerability against the reporting triggers of the EU Cyber Resilience Act (Regulation (EU) 2024/2847). Where a vulnerability in the Aidealy IDE extension or in the Aidealy Claude Code collector (in each case including the Aidealy backend services that software depends on to work) is being **actively exploited**, or a **severe incident** impacts the product's security, Aidealy reports it - as the Act requires from **11 September 2026** for software made available in the EU, including versions already on the market - to the designated national CSIRT coordinator and to ENISA through the EU's single reporting platform: an **early warning within 24 hours** of our becoming aware, a **fuller notification within 72 hours**, and a **final report** on the Act's timeline. Where the Act requires it, we also **inform impacted users** of the vulnerability or incident and of available corrective or mitigating measures. These are Aidealy's own regulatory duties to the authorities: they operate alongside - and do not change - the acknowledgement and handling commitments made to reporters in this section.

8. Coordinated disclosure

We follow a **coordinated disclosure** approach: we ask that you give us a reasonable opportunity to remediate a reported vulnerability **before** you disclose it publicly or to any third party.

- Please allow up to **90 days** from your report before any public disclosure, to give us time to develop and ship a fix. We will work to remediate within that window and to coordinate disclosure

timing with you, and we can agree a **longer or shorter** period by mutual agreement where the circumstances call for it.

- We are happy to coordinate the timing and content of any public advisory with you, and - where appropriate
 - to seek a CVE identifier.

9. Recognition and rewards

Aidealy does **not** currently operate a paid bug-bounty program, and this Policy makes **no offer of monetary reward** and **no formal public-recognition ("hall of fame") commitment**. By submitting a report you agree that you are **not entitled to compensation** for it.

We are genuinely grateful for responsible reports, and **may**, at our discretion, thank or publicly credit a researcher whose report we act on - but we do not promise to do so.

10. Privacy and related policies

If, despite the rules in §6, your research involves personal data, you must handle it consistently with this Policy and applicable data-protection law, and must not retain or further process it. How Aidealy itself handles personal data is described in our Privacy Policy and, for customer work data, our Data Processing Agreement; our broader security posture is described on our Security & Trust page. This Policy should be read alongside our Terms of Service and Acceptable Use Policy.

This Policy covers **security** vulnerabilities. If your concern is instead about how Aidealy develops or uses AI - for example the fairness of AI-assisted evaluations - rather than a security vulnerability, please write to ai-concerns@aidealy.ai. The AI-concerns channel is separate from this Policy, which continues to cover security vulnerabilities only.

11. Changes to this Policy

We may update this Policy from time to time as our practice and the legal landscape evolve. The version and "last updated" date at the top reflect the current version. Operating coordinated vulnerability disclosure is good security practice that EU cybersecurity law encourages (such as the NIS2 framework's coordinated-disclosure provisions) - and, in part, now requires: from **11 September 2026** the EU Cyber Resilience Act's vulnerability and incident reporting duties described in §7 apply to the Aidealy IDE extension and, from the day it is first made available to customers in the EU, to the Aidealy Claude Code collector, and the Act's remaining vulnerability-handling requirements - including a formal coordinated vulnerability disclosure policy and a stated support period - apply to products placed on the market from 11 December 2027.

12. Contact

Aidealy Ltd.

security@aidealy.ai

Hamidron 1

Herzliya 4654110

Israel