Halo Benefits

POLICY

Responsible AI Policy

Sets out how we use AI safely, fairly and with people in control.

UPDATED	VERSION	NEXT REVIEW
23 November 2026	2.0	November 2027

Halo Benefits Ltd · Owner: Alexander Dick, Head of Security & Operations · Classification: Internal & Public

Document control

Policy owner	Alexander Dick, Head of Security & Operations
Approved by	Board of Directors
Date approved	23rd November 2026
Next Review Date	23th November 2027
Review frequency	Annually
Applies to	All directors, employees and contractors, and all use of AI by Halo Benefits Ltd
Related policies	Information Security Policy, Data Protection Policy, Supplier Security Policy

Version history

Version	Date	Summary of changes	Author
1.0	30 Oct 2026	First issue	Alexander Dick
2.0	23 Nov 2026	Rebranded and restructured for clarity and readability	Dionne Clarke

Contents

1. Purpose of the policy 3
2. Scope 3
3. Accountability 3
4. Protecting your data 3
5. People stay in control 4
6. Accuracy and fairness 4
7. Security of AI features 4
8. AI coding assistants 5
9. Transparency and the law 5
10. Compliance and exceptions 5
11. Review 5

1. Purpose of the policy

Halo Benefits Ltd uses AI in our platform, in our internal operations and in software development. This policy makes sure that use is accountable, secure, lawful and honest about its limits, and that people, not AI models, remain responsible for decisions about individuals.

2. Scope

This policy applies to **Halo Benefits Ltd** (company number 17577358) and covers all use of AI and machine-learning services, including AI features in the Halo Benefits platform, speech and transcription services, AI used in internal operations and AI coding assistants.

3. Accountability

- **Board of Directors:** is accountable for our use of AI and approves any high-risk use.
- **Head of Security & Operations:** owns this policy and our register of approved AI providers and uses, and reviews any incident involving AI.
- **Feature owners:** every AI use has a named owner who tests it, monitors it and is responsible for its output.
- **Everyone:** uses only approved AI tools with approved data, and checks AI output before relying on it.

Every AI use is recorded, including what data goes in, which provider and model is used, where it is processed, how long data is kept, whether a person reviews the output, and why it is used.

4. Protecting your data

- **Approved providers only:** AI providers must be approved through our Supplier Security Policy before use, and only receive the minimum data their task needs.
- **No training on customer data:** customer data must never be used to train or improve AI models. We confirm this in each provider's terms and settings before approval.
- **Employer choice:** employers can switch off AI features for their organisation.
- **No unapproved tools:** customer personal data, children's data, payment information, passwords and other confidential information must never be entered into an AI tool that is not approved for that data. This includes consumer chatbots, browser extensions and free versions of tools without business terms.

5. People stay in control

- AI is never used to make decisions with legal or similarly significant effects on an individual. It is not used to decide benefit eligibility, calculate benefits or make safeguarding judgements. These remain rule-based or human decisions.
- Permitted uses include drafting, summarising, searching, forecasting, transcription and conversational help, all under human control.
- AI-drafted content sent to a customer, such as letters, documents or emails, is reviewed by a person first.
- Live AI assistants are clearly labelled as AI, are limited to defined tasks and hand over to a person on request. AI is never presented as a human.

6. Accuracy and fairness

- Before an AI feature is released, it is tested against realistic examples, including unusual cases and inputs from different groups of users, and the results are recorded.
- Features that summarise documents link back to their sources so a person can check them.
- Known limitations are documented and, where relevant, shown to users.
- Errors reported by users or staff are logged and reviewed. A feature with significant or repeated errors is corrected or withdrawn.
- Changes to the AI models or instructions behind customer-facing features are treated as changes under our Secure Development and Change Management Policy and are retested.

7. Security of AI features

- AI features use the same security controls as the rest of our platform, including sign-in, separation of each employer's data and limits on request rates.
- AI features can only see the data the task needs, and can never act beyond the user's own permissions.
- Risks specific to AI, such as attempts to manipulate its instructions or leak data, are considered in design and review.
- AI inputs and outputs are only logged as far as needed to run and fix the service.

8. AI coding assistants

Our developers use AI coding assistants. All AI-written code goes through exactly the same review, testing and security checks as any other change, and the developer who submits it is responsible for it. AI assistants are never given access to live systems or customer data.

9. Transparency and the law

- Details of which AI providers we use, what they do and where they process data are included in our sub-processor list, available on request.
- Privacy notices explain any AI processing that affects individuals.
- Our use of AI follows UK GDPR (including the rules on automated decision-making), ICO guidance on AI and data protection, and new UK requirements as they emerge.

10. Compliance and exceptions

Using an unapproved AI tool with customer data is treated as a security incident. Any breach of this policy may lead to disciplinary action or termination of contract.

Any exception (for example, a short pilot using made-up data) needs written approval from the Head of Security & Operations and is recorded.

11. Review

This policy is reviewed annually, and sooner after any significant change, including any new AI provider, model or customer-facing AI feature, or a relevant change in the law or guidance. The Head of Security & Operations proposes changes and the Board of Directors approves them.

> QUESTIONS ABOUT THIS POLICY?
>
> Contact Alexander Dick, Head of Security & Operations, at security@halobenefits.co.uk.