

Adam Hale Report Writer: DfE generative AI standards evidence summary

Procurement summary against the DfE *Generative AI: product safety standards* (updated 19 January 2026). Assessed 14 October 2026. Version 1.0.

Worked sample by Adam Hale · AI Safety Evidence on its own live product. <https://reports.adamhaleworks.workers.dev>

What this is and is not. The supplier's own evidence organisation and gap analysis, with real test results and honest gaps. It is not a certification, legal advice, a DPIA sign-off or a safety guarantee, and no specialist DPO or safeguarding professional has reviewed it.

Product at a glance

Purpose	Drafts end-of-term report comments for a whole class in the browser: the teacher pastes a class list, picks a level per skill, and the tool assembles each comment from fixed sentence banks (39,611 sentences, 23 banks, 9 education systems). The teacher edits every comment, then copies or exports it.
Intended users	Teachers and school staff only. Pupils do not use it and have no accounts. Pupil ages covered by the banks: 3-18 (EYFS to KS4 and equivalents).
DfE use case	DfE use case 7 (Administrative and management: report generation).
How AI is used	Before release only: the sentence banks were drafted and edited with AI writing tools, machine-checked and spot-checked by the author. At run time fixed rules pick sentences; no AI model is called.
Where data goes	Pupil names, pronouns, levels, notes and comments stay in the teacher's browser (session storage by default, local storage only if the teacher opts in). The server receives anonymous daily event counts and, only if a teacher sends it, optional feedback (message and an optional email).
Hosting	Cloudflare Workers (static files and a small Worker for counts and feedback). No analytics, advertising or AI service at run time.
Scope note	The DfE standards are written for generative AI products. This product generates no text at run time, but its content was produced with generative AI and schools may treat it as an AI tool, so it is assessed against all 13 areas. The five learner-facing areas (filtering, monitoring, cognitive, emotional and social development, mental health) are marked not applicable with reasons, because pupils never use the product.
Contact	adam.hale@atomicmail.ai

Status across the 13 areas (87 expectations)

11 Met	20 Partly met	7 Gap	1 Evidence needed	48 Not applicable
-----------	------------------	----------	----------------------	----------------------

Of the 39 expectations that apply, 11 are met, 20 partly met, 7 are gaps and 1 need evidence. 48 are not applicable because pupils never use the product.

Area	Status	Summary
1. Stated purpose Learner- and teacher-facing products	2 met 2 partly met 1 gap	Purpose is clear; needs an explicit 'teachers only' statement and evidence for the 'in minutes' claim.
2. Educational use cases Learner- and teacher-facing products	3 partly met	Use case 7 (report generation) in practice; to be declared with DfE numbering.

3. Filtering Learner-facing products	2 partly met 6 not applicable	Learner-facing rows not applicable. Adapted: fixed bank scanned in full; 280 feelings sentences need a human read.
4. Monitoring and reporting Learner-facing products	10 not applicable	Not applicable: no pupil use; no content is logged by design.
5. Security Learner- and teacher-facing products	1 met 4 partly met 2 not applicable	No run-time model; strict CSP; injection tests pass. Needs a patch policy and evidenced hosting MFA.
6. Privacy and data protection Learner- and teacher-facing products	2 met 2 partly met 3 gap 1 evidence needed	No pupil data reaches the supplier. Privacy notice lacks controller identity, lawful basis, rights, transfers; no DPIA yet.
7. Intellectual property Learner- and teacher-facing products	1 met 1 not applicable	Inputs never leave the browser, so nothing is reused.
8. Design and testing Learner- and teacher-facing products	1 met 3 partly met 1 gap	Strong automated tests; no testing with real teachers or assistive technology yet.
9. Governance Learner- and teacher-facing products	3 partly met 2 gap	Risk register drafted here; no published complaints route or AI safety policy yet.
10. Cognitive development Learner-facing products	9 not applicable	Not applicable: teacher tool.
11. Emotional and social development Learner-facing products	14 not applicable	Not applicable: teacher tool, no persona or chat.
12. Mental health Learner-facing products	6 not applicable	Not applicable: teacher tool.
13. Manipulation Learner- and teacher-facing products	4 met 1 partly met	No ads, loops or dark patterns; upgrade wording to review before checkout opens.

Release tests (14 October 2026)

Every sentence the product can produce comes from a fixed bank, so we scanned all 39,611 sentences in 23 banks and had a person review every hit, then ran browser tests on the live app in headless Chromium.


Test	What	Result	Observed
T-RW-01	Profanity or insults	Pass	11 hits. All hits are 'thick' in art and music ('thick line', 'thick and thin texture'). No insult or profanity.
T-RW-02	Labelling or deficit words about the child	Pass	15 hits. 'Disappointing' describes a test result, never the child; 'hopeless' is a spelling example. No labels.
T-RW-03	Medical, SEND or diagnostic terms	Review	34 hits. Most hits are 'condition' (science, computing) or 'send' (PE). Seven describe a pupil feeling anxious or mental health in general; see T-RW-09.
T-RW-04	Home or family circumstances	Pass	216 hits. Hits are Religious Education content (religious, religions, church) and geography (poverty). No home-circumstance statements about a pupil.
T-RW-05	Comparison with peers	Pass	0 hits. No comparison with classmates.

T-RW-06	Predictions or guarantees	Pass	1 hit. One hit: 'where success is not guaranteed' (a positive stretch suggestion).
T-RW-07	Over-praise that may read as flattery	Pass	11 hits. Ten are technical ('perfect cadence', 'present perfect'); one EYFS closer says 'brilliant questions'. Acceptable in a teacher-approved comment.
T-RW-08	Absolute negatives	Pass	22 hits. All 'cannot' hits are growth-mindset phrasing ('I cannot do it yet') or subject content.
T-RW-09	Statements about the child's feelings or wellbeing	Review	280 hits. These sentences state a child's feelings and are chosen by the teacher's rating, not by observation. Action: human read; tag; ask the teacher to confirm (risk R-02).
T-RW-10	Markup in a pupil name is shown as text, never run	Pass	Payload names shown as text; no script ran.
T-RW-11	Formula text in a name is neutralised in the CSV export	Pass	Cells starting with = + - @ neutralised.
T-RW-12	Teacher's personal note is inserted as typed (by design, not filtered)	Pass	Inserted as typed; teacher remains responsible.
T-RW-13	Low Resilience rating: wording about the child's feelings (observation for human review)	Observation	Free bank draft contained no feelings wording; full bank holds 280 such sentences (T-RW-09).
T-RW-14	No pupil name or note leaves the browser; no third-party requests	Pass	13 requests, all to the product's own host; 0 contained the canary name or note.
T-RW-15	Security headers on the live app	Pass	CSP script-src and connect-src 'self'; HSTS; frame denial.
T-RW-16	Class data is not kept after the tab closes unless the teacher opts in	Pass	No class data in local storage by default; nothing restored in a new tab.

Data flows

Adam Hale Report Writer: data flows (assessed 14 October 2026)

Green arrows carry pupil data and never leave the school side. Blue arrows carry no pupil data.



Pupil data never crosses this line (CSP connect-src 'self'; canary test T-RW-14: 0 of 13 requests)

F6 (owner reads feedback) and F1 request metadata (IP address, browser type, seen by the host, not logged) are described in the data-flow register.

Open gaps and actions

Row	Status	Gap	Action	Target
SP-4	Gap	No purpose-review step in the release process.	Add a purpose and claims review to the release checklist (Release Test Protocol, step 1).	2026-10-31
PD-2	Gap	Does not say where feedback is stored or which transfer safeguards apply (Cloudflare is a US company with a global network).	Add Cloudflare as sub-processor with location and transfer mechanism (Cloudflare's customer DPA; to be confirmed for this account).	2026-10-31
PD-3	Gap	Lawful basis is not stated for feedback or for counts.	State consent for feedback and email; explain why counts are not personal data.	2026-10-31
GV-2	Gap	No formal complaints mechanism with timescales or escalation.	Publish the complaints and escalation workflow from this pack (acknowledge in 1 working day; safety issues first).	2026-10-31
GV-5	Gap	The right to complain to the ICO is not mentioned.	Add to the privacy page and the complaints workflow.	2026-10-31
PD-4	Gap	No supplier DPIA has been completed or dated.	Owner completes a short supplier DPIA from the input pack; review every 12 months or on change.	2026-11-14
DT-1	Gap	No testing with real teachers yet (no confirmed human users at the 12 October baseline); no assistive-technology test.	Recruit five teachers through the in-app feedback form; one screen-reader run (NVDA) of the main task.	2026-12-15
PD-6	Evidence needed	No written assessment of whether the Children's code applies.	Write a short applicability note (likely not 'likely to be accessed by children'; specialist to confirm).	2026-11-14
SP-2	Partly met	No explicit 'for teachers only, not for pupil use' statement; no statement on SEND suitability.	Add an 'Intended users' box to /about/ and the app: teachers only; not for pupils; no SEND-specific bank.	2026-10-31
UC-1	Partly met	Not declared using the DfE numbering.	Add 'DfE use case 7: report generation' to /about/ and the procurement summary.	2026-10-31
FL-7	Partly met	No written process or timescale for withdrawing a sentence someone reports.	Complaints workflow: withdraw a reported sentence within 2 working days of confirmation (S2).	2026-10-31
SE-2	Partly met	Single maintainer: no second reviewer or signed releases; MFA on the hosting account not yet evidenced.	Record hosting-account MFA in the evidence folder; add a pre-deploy diff review step.	2026-10-31
SE-5	Partly met	No safety-specific gate before this pack; the content scan was not part of the build.	Add the content scan and the live browser tests (T-RW-01 to T-RW-16) to the release gate.	2026-10-31
PD-1	Partly met	Missing: the controller's legal identity and postal address, lawful basis, data subject rights, and the right to complain to the ICO.	Add these sections (owner to supply the controller details privately and decide what is published).	2026-10-31
PD-8	Partly met	No written breach procedure.	Complaints and escalation workflow covers breaches; publish a one-paragraph summary.	2026-10-31
GV-1	Partly met	Not yet signed off or scheduled for review.	Owner signs off; review every 6 months.	2026-10-31

SE-4	Partly met	No written patch and update policy with target times.	Publish a short patch policy (security fixes within 5 working days).	2026-11-14
GV-3	Partly met	No AI safety policy published.	Publish a one-page AI safety policy (scope, content checks, release gate, complaints).	2026-11-14
SP-5	Partly met	Headline 'in minutes' has no time measurement behind it; no user evidence yet.	Keep a claims register; measure drafting time with the first five teacher users or soften the headline.	2026-11-30
UC-3	Partly met	Gaps listed in this matrix remain open.	Close the actions below; re-issue the summary.	2026-11-30
FL-1	Partly met	Not every sentence has been read by a human; 280 sentences (0.7%) describe a child's feelings and are assigned by rating, not observation (risk R-02).	Run the content scan in every build; human read of the 280 feelings sentences; tag them so the app asks the teacher to confirm.	2026-11-30
SE-7	Partly met	No Cyber Essentials certificate (it costs money; out of scope for now).	State the position in the summary; revisit Cyber Essentials if a buyer requires it.	2026-11-30
DT-4	Partly met	Indirect risks (wellbeing wording, stigmatising language) were not assessed before this pack.	Risks R-01 to R-03 in the register, with the actions there.	2026-11-30
DT-5	Partly met	No human review of the EAL banks for deficit framing; no automated-decision statement.	Human read of the EAL banks; add a sentence: 'no automated decisions about pupils: the teacher sets every level'.	2026-11-30
MN-3	Partly met	Rows beyond the free limit show 'Pro'; wording not reviewed against this standard.	Review upgrade wording before checkout opens (test MN-T03).	Before checkout opens

What schools should still do

- Record your own DPIA screening with your DPO (section 1 of the DPIA input pack gives the answers).
- Tell staff to read and edit every comment before it goes home, and not to type health, SEND or family details into personal notes unless your policy allows it.
- Keep exported files in your usual report system under your retention policy.

Documents in this pack

File	Contents
01-Intake-Questionnaire.docx	Written answers about the supplier, product, data, safety, security, testing and governance.
02-Evidence-Matrix-13-Areas.xlsx	87 rows: requirement, evidence needed, status, evidence held, gap, action, owner, date; summary by area.
03-Risk-Register-and-Complaints-Log.xlsx	15 scored risks with controls and actions; complaints log.
04-DPIA-Input-Pack.docx	Screening, data inventory, flows, roles, necessity, risks and DPO questions answered.
05-Data-Flow-Register.xlsx and diagram	Every data flow and sub-processor.
06-Release-Test-Protocol.docx and Test-Set.xlsx	Procedure, go/no-go rules, 68 generic tests and the product's 16 tests with results.
07-Complaints-and-Escalation-Workflow.docx	Channels, severities, timescales, safeguarding and breach rules.

Sources: DfE, Generative AI: product safety standards, updated 19 January 2026 (<https://www.gov.uk/government/publications/generative-ai-product-safety-standards/generative-ai-product-safety-standards>); DfE, Data protection in schools: Procuring educational technology, 9 July 2026; ICO, Edtech examined, 24 June 2026; DfE, Keeping children safe in education 2026. Contains public sector information licensed under the Open Government Licence v3.0 (DfE, Generative AI: product safety standards, updated 19 January 2026).