MARLUC

CASE STUDY · MARLUC SYNOPSIS

Home Office asylum AI summaries

IMPLEMENTED 2024–2026 · GOVERNMENT AND PUBLIC SERVICES · UK · BAND D · PARTLY PROVED

VERSION 1.0 · PUBLISHED 16 NOVEMBER 2026 · UPDATED 16 NOVEMBER 2026 · FACTS CHECKED 16 NOVEMBER 2026

The Home Office uses generative AI to summarise asylum interview transcripts and search country policy information. Its 2024 pilot found time savings with no change in decision quality; after rollout in 2026, figures obtained by Open Rights Group show 22% of expert-checked summaries failed initial checks.

ABOUT THIS SYNOPSIS

This is MarLuc's summary and grading of publicly available sources, for information only. It is not investment, legal, financial or professional advice, and should not be relied on for any decision. Readers must validate it against the original sources and, where appropriate, challenge its conclusions.

Facts and figures are as reported by the sources listed below, which may be incomplete, self-reported or out of date. Check the original source before relying on any figure.

The impact band, proof marker and SEAM view are MarLuc's opinion. The organisation is not a MarLuc client and has not reviewed or endorsed this synopsis. Names and trademarks belong to their owners.

NOTICE WORDING LAST UPDATED 16 NOVEMBER 2026

AT A GLANCE

AI used: Generative AI summarisation of interview transcripts and a chat-based policy search assistant.

Implemented: 2024–2026. Grade: band D, Task support, Partly proved.

Headline result, as reported: Pilot: 23 minutes saved per case on transcript review; after rollout, 45 of 203 expert-checked summaries failed initial checks, per Open Rights Group

1 · The case

The facts, as reported by the 5 sources listed at the end. Where a source is silent, the field says "Not reported".

Organisation

Home Office (UK). UK government department responsible for immigration and asylum. The tools are used by asylum decision-makers in its Decision Making Units. Canvas block: 4. Operations and process. Also 8. Governance and risk, 5. People and skills.

Business challenge

The challenge. The Home Office said decision-makers spend a substantial amount of time analysing asylum interview transcripts, which can run to 50 pages or more, and finding country policy information. The tools were trialled as part of a wider asylum system programme of change to speed up these steps.

Who implemented it

OpenAI (Model or cloud provider): Supplied the GPT-4o model behind both tools, per Open Rights Group.
Microsoft (Model or cloud provider): Hosts the model through Azure AI Foundry, per Open Rights Group.
In-house (In-house): Home Office: designed, piloted and evaluated the tools.

Technology (SEAM pillar E)

SEAM types: T3 API integrations, T4 RAG / vector search over your own data. Deployment: Cloud APIs. Vendors and products: OpenAI, Microsoft, GPT-4o, Azure AI Foundry, Asylum Case Summarisation (ACS), Asylum Policy Search (APS).

Where it was built in: Used by asylum decision-makers: ACS summarises interview transcripts; APS finds and summarises country policy and information notes in a chat interface.

What changed

Before: decision-makers read full transcripts and searched policy notes themselves, and technical specialists reviewed and summarised cases, per the Home Office's data protection assessment cited by Open Rights Group. After: decision-makers receive an AI summary and an AI policy search assistant. Open Rights Group says the summarisation tool automates work previously done by technical specialists. [1] [2] [3] [4] [5]

What it cost

Not reported

What it saved

Pilot, by the Home Office's evaluation: 23 minutes per case on transcript review (32%) and 37 minutes per case on policy research. Open Rights Group's analysis found monthly initial asylum decisions since April 2026 down about 42% on the same period in 2025 and interview-to-decision time up 29%, though these figures do not show the tools caused the change. [1] [2] [3] [4] [5]

Governance and control

The tools were designed so they cannot be used alone to make a decision; technical specialists reviewed all pilot summaries before use, and equality impact assessments were completed. After rollout, experts checked about 0.3% of summaries, and the Home Office confirmed in August 2026 that expert testing of the policy search tool had stopped. Applicants are not told when AI is used. The tools are not on the Algorithmic Transparency Recording Standard Hub, per Open Rights Group.

S2 · Accountable owner · Not reported
Not reported.

A1 · Safe use · Missing (independent source shows it was absent)
Open Rights Group found the Home Office did not apply the government's AI Playbook and refused information on staff training and prompts. [3]

A2 · Human in the loop · Claimed
The Home Office says the tools cannot be used to make a decision and trained caseworkers decide every claim. [1] [5]

A3 · Monitoring and assurance · Missing (independent source shows it was absent)
Post-rollout expert checks covered about 0.3% of cases and testing of the policy search tool stopped in 2026, per Open Rights Group's figures. [2] [3]

A4 · Vendor AI · Not reported
OpenAI models via Microsoft's Azure AI Foundry, per Open Rights Group; vendor assurance is not reported. [2]

A5 · Agent controls · Not reported
No agents described.

A6 · Standards and regulation · Not reported
No regulator or court has acted; the tools are not on the Algorithmic Transparency Recording Standard Hub, and the borders inspectorate is considering them for 2027. [3] [5]

What worked, what did not, what is unproven

WORKED

– Time saved in the pilot: 23 minutes per case on transcripts and 37 minutes per case on policy searches, against comparison groups.

– Decision quality in the pilot's dip sample was similar between test and comparison groups.

– Most pilot users said they would keep using the summarisation tool (65%).

DID NOT

– Accuracy after rollout: 45 of 203 sampled summaries failed initial expert checks, against 9% removed in the pilot, per Open Rights Group.

– Second-stage reviews: 30 of 56 summaries failed to meet the required standard, per The Register.

– Ongoing testing: expert testing of the policy search tool stopped, confirmed by the minister in August 2026, per EIN.

UNPROVEN

– Whether the tools save time at scale: decision volumes fell and times rose after rollout, per Open Rights Group, with no causal link shown.

– Whether inaccurate summaries have affected any asylum decision.

– Whether the use is lawful: a legal opinion commissioned by Open Rights Group says not telling applicants is likely unlawful; no court has ruled.

Timeline

May 2024 Small-scale pilots of the Asylum Case Summarisation and Asylum Policy Search tools begin [1]

Apr 2025 Home Office publishes its pilot evaluation [1]

Mar 2026 Open Rights Group publishes a legal opinion that not telling applicants about AI use is likely unlawful [4]

Oct 2026 Open Rights Group report finds 22% of sampled summaries failed initial expert checks and calls for suspension [3]

2 · Grade and SEAM assessment

MarLuc's opinion, made on the evidence above. How cases are graded: marluc.co.uk/case-studies/how-we-rate

Level of change and impact

Claimed: Not applicable. Grade: D, Task support. Proof: Partly proved. Level of change: Task support.

Task support (Level 1, band D). The tools help decision-makers read and search faster; the Home Office says no step of the decision is replaced. Open Rights Group argues the summaries remove a layer of specialist review, which would be a process change.

Graded D. The tools support decision-makers in reading transcripts and finding policy, and the Home Office says they do not replace any part of the decision. The pilot showed time savings against comparison groups; post-rollout quality figures obtained by a campaign group show higher failure rates than the pilot, and no published evaluation of the full rollout exists.

The SEAM view and lessons learnt

The pilot measured time saving with comparison groups and a quality sample. The gap is in what followed: wider rollout with expert checks on about 0.3% of cases, a higher failure rate than in the pilot, and testing of one tool stopped. For any organisation using AI summaries in high-stakes decisions, the quality checks need to scale with the rollout, and the people affected need a way to see and correct what the AI wrote about them.

1. Keep accuracy sampling in proportion to volume after a pilot ends.
2. Tell the people affected when AI summarises their case and let them correct it.
3. Publish a full evaluation before claiming time savings at scale.

Evidence

Score 2 of 4. Partly independent: the Home Office's own comparison-group pilot evaluation, and post-rollout figures from Home Office sampling obtained and analysed by Open Rights Group, a campaign group; no independent evaluation.

Evaluation method. Before and after comparison: the pilot compared test and comparison groups on self-logged case times, with a dip sample of decision quality; the evaluation notes its methods rely on self-reporting. Post-rollout quality data come from Home Office expert sampling released to Open Rights Group.

Sources disagree. The Register and EIN say both tools use OpenAI's GPT-4o; Open Rights Group's earlier release says the summarisation tool uses ChatGPT-4, and IBTimes describes the GPT-4 family. IBTimes says the tool was rolled out to every caseworker in April 2026; The Register says

3 · Original sources

Read these in full. To challenge a fact or a grade, use the contact form at marluc.co.uk or the link on https://marluc.co.uk/case-studies/home-office-asylum-ai-summarisation

[1] GOV.UK (Home Office): Evaluation of AI trials in the asylum decision-making process
https://www.gov.uk/government/publications/evaluation-of-ai-trials-in-the-asylum-decision-making-process/evaluation-of-ai-trials-in-the-asylum-decision-making-process

[2] The Register: Home Office asylum AI summaries fail expert checks
https://www.theregister.com/a/5302649

[3] Electronic Immigration Network: Open Rights Group raises concerns over Home Office AI asylum tools
https://www.ein.org.uk/node/810461

[4] Open Rights Group: Home Office use of AI in asylum cases likely to be unlawful, legal opinion finds
https://www.openrightsgroup.org/press-releases/home-office-use-of-ai-in-asylum-cases-likely-to-be-unlawful-legal-opinion-finds/

[5] IBTimes Singapore: UK uses AI to summarise asylum interviews; one in five checked summaries failed
https://www.ibtimes.sg/uk-uses-ai-summarise-asylum-interviews-one-five-checked-summaries-failed-95621

Home Office (UK) website: https://www.gov.uk/government/organisations/home-office

Version 1.0 · Published 16 November 2026 · Updated 16 November 2026 · Facts checked against sources 16 November 2026 · Downloaded 27 November 2026. This PDF may be out of date: the latest version, and a way to report an error or a better source, is at https://marluc.co.uk/case-studies/home-office-asylum-ai-summarisation

VERSION HISTORY

v1.0 16 November 2026 First published