Citizens Advice Caddy adviser assistant

IMPLEMENTED 2022–2025 · CHARITIES AND NONPROFITS · UK · BAND D · CLAIMED ONLY
VERSION 1.0 · PUBLISHED 25 NOVEMBER 2026 · UPDATED 25 NOVEMBER 2026 · FACTS CHECKED 25 NOVEMBER 2026

A generative AI assistant drafts answers for trainee advisers from trusted sources, and a supervisor approves each one before the adviser sees it.

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Facts and figures are as reported by the sources listed below, which may be incomplete, self-reported or out of date.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 assistant (large language model) drafting answers for supervisor approval.

Implemented: 2022–2025. Grade: band D, Task support, Claimed only.

Headline result, as reported: Pilot in six local offices: 80% of drafts supervisor-approved and response time cut from about 10 minutes to 4, by Citizens Advice's account

1 · The case

The facts, as reported by the 4 sources listed at the end. Where a source is silent, the field says "Not reported".

Organisation

Citizens Advice. Charity network giving free advice in England and Wales, with 270 local organisations in 2,540 locations and 2.8 million people advised in 2024, per Stanford. Caddy was built by the Stockport, Oldham, Rochdale and Trafford office. Headcount: Not reported. Canvas block: 4. Operations and process. Also 5. People and skills, 3. Customers and proposition.

Business challenge

The challenge. Trainee advisers asked supervisors for guidance on complex cases. As demand rose, supervisors became a bottleneck and clients waited longer, by the charity's account.

How it was addressed. Caddy answers an adviser's question using only trusted sources, GOV.UK and Citizens Advice's own content. The draft goes to a supervisor, who edits and approves or rejects it. Only approved answers reach the adviser, who relays them to the client. Caddy never talks to the public. [1] [2] [3] [4]

Who implemented it

In-house (In-house): Built Caddy and ran the pilot..
I.AI (UK government) (Industry partner): Collaborating since 2024 on a scalable version for other advice providers, per SCVO..
Anthropic (Model or cloud provider): Claude is the language model behind Caddy, per SCVO..

Technology (SEAM pillar E)

SEAM types: T4 RAG / vector search over your own data, T3 API integrations. Deployment: Not reported. Vendors and products: Anthropic, Amazon Web Services, Claude, AWS Serverless.

Where it was built in: Built on AWS Serverless architecture and integrated with Google Workspace chat; answers draw only on GOV.UK and Citizens Advice content, and an approval step sits between the model and the adviser.

What changed

Before: advisers searched alone or waited for a supervisor. After: advisers get a sourced draft that a supervisor approves. Measured results, by the charity's account: 80% of responses approved; response time 4 minutes against 10 (Stanford) or 9 (Charity Digital). [1] [2] [3] [4]

What it cost

Not reported

What it saved

Not reported in money. Stanford says response times fell by 50%, from 10 minutes to 4; Charity Digital says 4 minutes including supervisor approval against 9 for an adviser searching alone. [1] [2] [3] [4]

Governance and control

Supervisors review every draft before advisers see it. The charity set four principles: clients always speak to a human, every response is reviewed, only pre-approved sources are used, and advisers know when AI is used. Accountable owner and vendor assurance: Not reported.

S2 · Accountable owner · Not reported
Not reported who is accountable at the charity; the project documentation says supervisors verify responses. [4]

A1 · Safe use · Claimed
The charity set four principles, including that clients always speak to a human and Caddy only uses pre-approved sources. [1]

A2 · Human in the loop · Claimed
Supervisors review every draft; advisers see only approved responses and a supervisor's comments if a draft is rejected. [2] [4]

A3 · Monitoring and assurance · Claimed
A dashboard tracks adviser questions and subject areas; the charity ran a 4-6 week pilot measuring accuracy and time. [1] [2]

A4 · Vendor AI · Not reported
Not reported how the model provider is assessed; SCVO says Caddy uses Anthropic Claude. [2]

A5 · Agent controls · Not reported
No agents described; Caddy drafts guidance for advisers and does not communicate directly with the public. [1]

A6 · Standards and regulation · Claimed
Stanford says the approach aligns with the UK Algorithmic Transparency Recording Standard; the charity also consulted a public panel. [1]

What worked, what did not, what is unproven

WORKED

– Faster responses in the pilot, by the charity's account.

– Advisers appreciated the efficiency, per Stanford.

– The design keeps a person in every answer: supervisor review plus an adviser who speaks to the client.

DID NOT

– About one in five drafts was not approved; Stanford asks where the 20% came from.

– Caddy struggled with unclear or incomplete questions.

– Two sources were not always comprehensive for welfare or debt cases, and some advisers wanted to skip supervisor approval in low-risk cases.

UNPROVEN

– Accuracy and time savings measured by anyone outside the charity.

– Effects on client outcomes and on supervisor workload over time.

– Results beyond the pilot offices.

Timeline

2022 Citizens Advice Stockport, Oldham, Rochdale and Trafford starts looking at AI for service delivery after ChatGPT launches [3]

2024 Citizens Advice and the government's I.AI team begin building a scalable version for other advice providers [2]

Mar 2025 Charity presents pilot results: 80% of responses supervisor-approved, response times roughly halved [1]

Mar 2025 SCVO reports Caddy runs on Anthropic Claude and keeps supervisors reviewing every draft [2]

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: Claimed only. Level of change: Task support.

Task support (Level 1, band D). Advisers get drafts, and a supervisor still reviews each one. Any move to supervisor-free approval would change the process.

The SEAM view and lessons learnt

The gain looks tied to the design: narrow trusted sources and a supervisor reviewing every draft. The sources do not separate AI from other changes such as training in prompts. Before copying it, a leadership team should ask how it measures accuracy beyond supervisor approval, and

1. Limit the sources an assistant may use.
2. Keep a named reviewer on every answer until accuracy is measured.
3. Record time and approval rates before expanding.

Evidence

Score 1 of 4. Self-reported: the pilot figures come from the charity's own presentations, relayed by Stanford and Charity Digital.

Evaluation method. Self-reported pilot figures from a 4-6 week trial in six offices, relayed by Stanford and Charity Digital. No independent evaluation was found.

Sources disagree. Stanford gives response times of 10 minutes against 4; Charity Digital gives 9 minutes against 4. Charity Digital names the Stockport, Oldham, Rochdale and Trafford office; SCVO says Greater Manchester.

Status. Ongoing: SCVO reports collaboration with the government's I.AI team since 2024 on a scalable version; current scale could not be confirmed from sources opened.

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/citizens-advice-caddy

[1] Stanford Justice Innovation: How AI is Augmenting Human-Led Legal Advice at Citizens Advice (20 March 2025)
https://justiceinnovation.law.stanford.edu/how-ai-is-augmenting-human-led-legal-advice-at-citizens-advice/

[2] SCVO: Building Caddy, an AI support tool for adviser teams
https://scvo.scot/p/97562/2025/03/10/building-caddy-an-ai-support-tool-for-advisor-teams

[3] Charity Digital: AI and the future of service delivery
https://charitydigital.org.uk/essentials/ai-and-the-future-of-service-delivery-11843

[4] Caddy documentation: Introducing Caddy
https://i-dot-ai.github.io/caddy-chatbot/docs/about_caddy.html

Citizens Advice website: https://www.citizensadvice.org.uk

Version 1.0 · Published 25 November 2026 · Updated 25 November 2026 · Facts checked against sources 25 November 2026 · Downloaded 26 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/citizens-advice-caddy

VERSION HISTORY

v1.0 25 November 2026 First published