

A sample of what you receive.

This is the compressed version of an audit hand over. Clients receive the full pack, which carries every workflow we looked at and the reasoning written next to each one.

DOCUMENT

Sample deliverable

PREPARED BY

OpenKit, Cambridge and Durham

DATE

August 2026

What our audits keep finding.

We have run seven audits so far. The sectors have little in common, and the findings below still repeat.

Someone is already paying for AI on a personal card.

3 OF 3

In all three of the engagements we have documented in full, people were buying their own AI subscriptions before anyone had agreed a rule about it. At one firm that meant five or six individual accounts running on customer, supplier and financial data with nothing managing them centrally.

People ask for search over their own documents before anyone suggests it.

3 OF 3

It was among the first things people asked for in all three. At one consultancy, five of the ten we interviewed raised it independently and unprompted, and at another it surfaced in every single session we ran.

Where the data lives decides the recommendation.

4 OF 7

Data sovereignty or data readiness shaped the core recommendation in at least four of the seven. EMQN needed sovereign hosting, Thorlux needed the UK region, and an international oil and gas service provider went fully on premises.

Every audit produced a counted number.

7 OF 7

Each one put a figure on what the work costs today, taken from the client's own operation rather than a benchmark. Stow Brothers' admin came to 172 hours a month, Thorlux validated around 6.5 million occupancy records against the live dashboard, and at EMQN the projection was more than 200 hours of assessor time a year.

The audit names where not to use AI.

2 OF 7

A written boundary has shipped in at least two of the seven, and in both of those we recommended against AI in specific workflows. One of those lists ran to eight items.

Nothing scales until a pilot has passed a gate.

2 OF 7

In at least two of the seven the recommendation was a pilot with a go or no go decision at the end of it, and the rollout only started once that decision had been taken.

GOVERNANCE

Governance vacuum

Ungoverned personal accounts

DATA

Data readiness

Data sovereignty

SYSTEMS AND PEOPLE

Memory locked in heads and inboxes


Closed systems that will not join up

Single points of failure

The blockers named across the set. Each one has to be settled before a workflow can run on its own.

The route we suggest.

Most engagements start with the AI Audit and Transformation. What follows depends on what it surfaces: many teams keep us on as their Embedded AI Lead, and some commission a bespoke build where nothing off the shelf will fit.



The audit.

It runs three to four weeks, and your first automation is live before it ends. You leave with a written report your board can read in one sitting alongside a prioritised 12 month roadmap. Your fee is fixed and agreed before anything starts.

What follows it.

OpenKit becomes your business's embedded AI team; the people who ran your audit stay on to deliver your 12 month roadmap. That runs as a monthly retainer after the audit, with an OpenKit team working inside yours. Where the audit surfaces something nothing off the shelf will fit, a bespoke build is scoped per project, surfaced by the audit, and owned outright by you.

Governance underneath.

Governance runs beneath all three. The rules for what AI may and may not do are written while the audit is still running, as an AI Charter: four decisions with names against them, short enough that people actually read it.


What we rule out.

The work we judged and turned down is a named part of what you get, with the reason written next to each item. A workflow that fails on data does not reach your roadmap, however good the idea behind it was, and telling you that in the first week is part of what the audit is for.


How the weeks run.

The audit runs three to four weeks, and the roadmap that comes out of it is ordered over twelve months. The shape below is typical, and your own ordering comes out of what your audit finds.

THE AUDIT, THREE TO FOUR WEEKS



THE 12 MONTH ROADMAP



WORTH KNOWING

Prove and Scale land together when the scope allows, so week four is often the same team carrying on into the next workflow with nobody new to bring up to speed.

Week one needs the people who actually do the work being mapped. After that we need one point of contact and not much more of your time.

The roadmap is ordered by what the audit found and costed against it, so you can take it to a board without us in the room.

What the platform costs.

Platform seats are the one cost you can look up before you speak to anyone. What a team actually spends turns mostly on how many people need the heavier seat, so the mix matters more than the vendor.

VENDOR PRICING AT THE TIME OF WRITING

CLAUDE TEAM PLAN	BILLED ANNUALLY	BILLED MONTHLY
Standard seat	\$20 per seat per month	\$25 per seat per month
Premium seat	\$100 per seat per month	\$125 per seat per month

Anthropic publishes these in US dollars, and they were the listed prices in August 2026. At the exchange rate that month a standard seat came to roughly £16 to £20 and a premium seat to roughly £80 to £100. Check the vendor's own page before you budget against it.

Seats are not the whole bill.

Three other things carry a monthly cost, and they are small next to the seats.

- Model usage beyond the seat, billed on what you send it.
- Somewhere for the automation to run.
- A log someone can read afterwards.

For a single workflow those usually come to tens of pounds a month, and the audit costs them against the hours the work takes today so the comparison is a real one.

Where the money often comes from.

Consolidating the subscriptions people have already bought on their own cards pays for a good part of this, which is why the first finding on page two is a commercial one as much as a governance one.

OUR OWN FEE

You will not find it in this document, or anywhere on our site. Your fee is fixed and agreed before anything starts, and what sets it is the scope: which workflows you want in range, and the systems they touch.

Start your **audit.**

Tell us which workflow eats your team's week and we'll come back with how we would scope it. We start with a conversation, audit where AI actually pays back, and build the first automation into how your team already works. We reply within one working day.

WHAT THE AUDIT
COVERS

openkit.co.uk/ai-services/ai-audit

START THE
CONVERSATION

openkit.co.uk/contact

You do not need to know what you want before an audit. Working out where the value sits is the audit's first week, not something you have to bring with you.