

THE AI OPPORTUNITY AUDIT · THE EXAMPLE PACK

This is what lands on *day 14*.

A complete audit for Elmbrook Care Ltd — a fictional 40-staff care provider — produced exactly the way yours would be. Real method, made-up company.

ELAPSED TIME

2 weeks

YOUR TEAM'S TIME

~4 hours

SYSTEM ACCESS

None

Every number in this pack is illustrative. Elmbrook Care Ltd does not exist. Your pack will carry your own numbers, from your own rota, invoices and interviews.

HOW TO READ THIS PACK

Nine deliverables. None of them a *forty-page report*.

The audit answers one question: where does AI pay in this specific business, and where does it not. Two weeks, four hours of the team's time, no system access, no resident data handled. What follows is everything Elmbrook's owner received at playback.

01 Readiness Scorecard

A score out of 100 across six dimensions, built from evidence rather than self-assessment. Page 3.

02 Opportunity Register

Every AI use case found, scored on impact, effort and risk — including the rejected ones. Pages 4–5.

03 Costed Quick Wins

The top three, costed against Elmbrook's own rota and invoices. Page 6.

04 Shadow AI & Ground Rules

What staff were already doing, and the one-page rules that fixed it. Page 7.

05 90-Day Roadmap

Prioritised, sequenced, an owner and a date on every line. Page 8.

06 The Re-Score Rubric

How Elmbrook measures itself again at day 90, without paying anyone. Page 9.

Also delivered, not reproduced here: the 14-slide findings deck (PowerPoint and PDF), the recorded playback session, and one round of amendments.

THE HOUSE RULE

Sometimes the right AI spend is zero. If that's the finding, that's the finding — in writing.

Elmbrook scores 58 / 100. Here is why.



Data		6/10
Tools		7/10
Skills		4/10
Leadership		8/10
Governance		3/10
Culture		7/10

The reading, in plain English

Leadership is bought-in and the team is willing. But nobody has been shown how to use the tools properly, and there are no ground rules — five staff were already pasting resident details into free AI accounts after long shifts. Fix governance and skills first. The time savings then land safely, and the score moves fastest where it is currently lowest.

Scored from
One questionnaire, one kickoff,
three staff interviews, desk
review of the tool stack.

Not scored from
Self-assessment. The free
online check asks what you
think; this measures what
happens.

Re-measurable
Same rubric, day 90, no
consultant required. Target after
this roadmap: 70+.

Eleven use cases found. *Not all survive.*

Every candidate scored 1–5 on impact and 5–1 on effort and risk (higher is better). The verdict column is the point of the audit: a ranked shortlist, not a brainstorm.

#	USE CASE	IMPACT	EFFORT	RISK	SCORE	VERDICT
01	Care plan first drafts & reviews Drafts assembled from assessment notes; senior signs off	5	4	4	13	DO NOW
02	Daily notes & handover summaries Shift notes to clean records, reviewed before filing	4	4	4	12	DO NOW
03	Local authority invoice chasing Remittance matching plus drafted, escalating reminders	4	4	3	11	DO NOW
04	Supervision & spot-check write-ups Voice notes to formatted drafts in house style	3	4	4	11	QUARTER 2
05	Training expiry register Certs, DBS and registration dates warned early	3	5	5	13	QUARTER 2
06	Job advert drafting Adverts from role templates, manager edits	2	5	5	12	QUARTER 2
07	CQC evidence mapping Existing evidence mapped to framework, gaps named	4	3	4	11	QUARTER 2

Impact 1–5

Hours saved or cash brought forward, measured against Elmbrook's own rota and invoices — never a vendor's brochure claim.

Effort 5–1

Higher is easier. Uses tools Elmbrook already pays for wherever possible; anything needing new systems scores low.

Risk 5–1

Higher is safer. Anything touching resident data, clinical judgement or regulated records is marked down hard.

The ones I *recommended against*.

Most audits only sell you things. The rejected list is usually worth more than the shortlist — it is the spending you avoid, with the reason in writing.

#	USE CASE	WHY NOT	VERDICT
08	Medication ordering automation GP ↔ pharmacy ↔ agency loop	The biggest prize in the sector and an integration problem, not an AI one. Three parties, no shared system, clinical risk. Anyone selling this as a quick win has never sat in a care office.	REJECTED
09	AI rota optimisation Full auto-scheduling	Elmbrook's rota problem is cover at 6:30am, not optimisation. A ranked cover shortlist with drafted messages delivers 80% of the value at 10% of the risk. The full version waits.	REJECTED
10	Resident-facing chatbot Family enquiries out of hours	Low volume (11 enquiries/week), high reputational risk, and families phoning at 2am need a person. A better voicemail script won this one.	REJECTED
11	Predictive falls analytics Sensor + notes analysis	Clinical judgement territory, thin data, and the sensor estate does not exist. Revisit in a year if the notes triage (02) proves itself.	REJECTED

The house rule

Sometimes the right AI spend is zero. If that is the finding, that is the finding — in writing. Rejections above saved Elmbrook an estimated £9,000–£15,000 in year-one spend on tools that would not have paid back.

Three wins, costed against *Elmbrook's own numbers*.

Rates from Elmbrook's actual payroll (£18/hr office-loaded). Hours measured in the interviews, not estimated in a workshop.

CARE PLANNING

~9 hrs/wk

est. £8,400 / yr

First drafts and review updates assembled from assessment notes. A senior carer checks and signs every plan. Minutes, not an evening.

DAILY RECORDS

~5 hrs/wk

est. £4,700 / yr

Shift notes turned into clean daily records and handover summaries, reviewed before filing. Ends the end-of-shift write-up pile.

FINANCE

~11 days DSO

cash-flow, not headcount

Invoices matched to commissioned hours; polite, escalating reminders drafted and scheduled, for councils and private payers alike.

The year, added up

£13,100

est. annual value of wins 1–2 alone



Set against the audit fee, the first win alone repays it inside eight weeks of going live. All figures illustrative.

The AI already running *without permission.*

Found in the interviews, handled without blame. This is the finding that most surprises owners, and the cheapest to fix.

WHAT WAS HAPPENING	WHO	THE FIX
Resident details pasted into free ChatGPT accounts to draft care notes after long shifts	5 care staff	Red line in ground rules; approved tool with a data agreement instead
Personal Gemini used to draft family emails, names included	1 deputy manager	Same job moved onto the approved account; templates provided
CV screening through a free AI site, candidates unaware	Office admin	Stopped; sifting criteria written down; human decision recorded

GROUND RULES · EXCERPT OF THE ONE-PAGER

Never

- Any personal or private AI account, for any work task
- Anything identifying a resident into any AI tool without a data agreement
- AI output filed unread — a human signs everything

Always

- The approved tools list, kept to one page
- A named owner (deputy manager) for questions
- Say when AI helped write something official

The full editable one-page policy and staff ground-rules sheet ship with the pack, ready to sign. Elmbrook’s went out in week one of the roadmap.

A name and a date on *every line*.

One page in the real pack. Sequenced so governance lands before automation, and nothing depends on a tool Elmbrook does not already pay for.

WHEN	ACTION	OWNER	DONE WHEN
Weeks 1–2	Ground rules signed and shared. Approved tools list issued; one absolute rule on resident data. Shadow AI comes in from the cold.	Registered Manager	All staff signed
Weeks 1–2	Baseline captured. Hours on care plans and notes logged for two weeks, so savings are measured, not guessed.	Deputy Manager	Baseline sheet filled
Weeks 3–6	Care-plan drafting live with deputy and two seniors. Every output human-checked. Weekly 15-minute review.	Deputy Manager	First 20 plans signed
Weeks 5–8	Daily-notes summarising rolled out across shifts, one wing at a time.	Senior Carers	All wings live
Weeks 7–10	Invoice reminders switched on inside the existing accounts package. Escalation ladder agreed with the office.	Office Manager	First cycle run
Weeks 11–12	Re-score readiness with the rubric on page 9. Compare against baseline. Decide quarter two from evidence.	Owner	Score ≥ 70

Measure yourselves again at day 90. *Free.*

The same six dimensions the audit scored, with what each score means written in plain English. No consultant required to re-run it — that is deliberate.

DIMENSION	3 OR LESS LOOKS LIKE	7+ LOOKS LIKE
Data	Nobody trusts the numbers; three versions of every list	One list per thing, an owner per list, spot-checked monthly
Tools	Free accounts, personal logins, no data agreements	Paid tools with agreements; one approved list, one page
Skills	One enthusiast; everyone else guesses or avoids	Every office role has done the hour; two champions named
Leadership	"We should look at AI" said quarterly, nothing owned	A backlog, a monthly review, a named owner
Governance	No policy, no ground rules, shadow use unmanaged	Signed policy, ground rules on the wall, incidents route somewhere
Culture	Fear or hype; nobody admits what they use	Staff volunteer what works and what worries them

Elmbrook's day-90 targets

Governance 3 → 7 (policy signed, rules live). Skills 4 → 6 (AI hour delivered, champions coached). Overall 58 → 70+. Everything else holds; nothing is allowed to slip while the new habits bed in.

Every business is on the same ladder. The audit is the *rung 1 → 2 move*.



Start smaller if you want: the £450 Health Check is 90 minutes and a four-page brief, fully credited against the audit. One rule, one number — each step is credited into the next: the £450 check comes off the audit, and £450 of the audit comes off your first build.

Book, or just ask

consulting.ai-sustained.com
kevin@ai-sustained.com

Judge the thinking first

ai-sustained.com
Proof published in the open

Elmbrook Care Ltd is fictional. The method, the deliverables and the discipline are not. © AI Sustained · Independent AI & data consultancy · UK