AI Opportunity Audit.

Prepared for: Meridian Digital — a 14-person marketing agency.

A complete example of the final audit deliverable, built on a fictional client so you can see the standard before you write your own. Structure, depth, and ROI math are exactly what a real $2–5k audit should hand over.

EXECUTIVE SUMMARY

Findings at a glance

Meridian's new-business growth is limited by three constraints in its outbound lead generation process — **lead quantity**, **lead quality**, and **response rates** — compounded by slow internal communication. This audit identified six opportunities worth **$76,100 per year** in recovered time against a one-off build cost of **$16,200**, with outreach capacity roughly tripled.

$76.1k	$16.2k	3×	~3
ANNUAL SAVING IDENTIFIED	ONE-OFF BUILD COST	OUTREACH CAPACITY	MONTHS TO PAYBACK

The three constraints

CONSTRAINT	TODAY	FIX
Quantity	~240 leads/month ceiling — one sourcing seat, lists built by hand	Scale to several sourcing accounts + sending infrastructure
Quality	Firmographics only; 30% of outreach effort spent on poor-fit leads	AI lead scraping with intent signals
Response rate	2.5% replies — templated emails, generic follow-ups	AI-personalised outreach from each lead's last 6 months of activity

Recommended sequence

PHASE	ITEMS	TIMELINE
Phase 1	Intent-based AI lead sourcing · multi-account sourcing & sending infrastructure	Weeks 1–3
Phase 2	AI-personalised outreach · local AI knowledge model in Slack · automated client reporting	Weeks 4–8
Parked	Website chatbot · full proposal automation	—

SECTION 1

Company snapshot

THE BUSINESS

Meridian Digital is a 14-person marketing agency serving 22 retained clients — paid social, SEO, and content. Growth depends on outbound: cold outreach to companies matching their ICP. The team is at capacity, and outbound results have plateaued.

AUDIT SCOPE

Three days on site. Nine stakeholder interviews recorded and transcribed with Tactiq. Three processes mapped end to end in Mermaid.ai.

PEOPLE INTERVIEWED

ROLE	FUNCTION
Owner / MD	Sales, final approvals
Ops manager	Delivery, reporting
3 account executives	Sourcing, outreach, reporting
2 account managers	Client relationships
Studio lead	Content production
Finance admin	Invoicing

TOOL STACK INVENTORY

FUNCTION	TOOL	NOTES FROM INTERVIEWS
Lead sourcing	LinkedIn Sales Navigator (1 seat), Apollo (1 seat)	Single seats cap sourcing volume; lists exported by hand
CRM	HubSpot (Starter)	Used at handoff only; outbound runs in spreadsheets
Email / outreach	Google Workspace, one sending domain	Outreach written one email at a time in Gmail
Internal comms	Slack	Questions answered from memory; long waits for answers
Reporting	Google Sheets + Looker Studio	Numbers pasted in by hand monthly
Ad platforms	Meta, Google Ads, LinkedIn	Data exported manually for client reports

PROCESSES MAPPED

#	PROCESS	WHY IT'S IN SCOPE
1	**Outbound lead generation** (ICP-to-CRM handoff)	The primary constraint on growth — mapped in detail in Section 2
2	**Monthly client reporting**	22 near-identical reports assembled by hand every month
3	**Internal communication & knowledge flow**	How answers about clients, campaigns, and the ICP move around the team

SECTION 2

Process map: outbound lead generation

Define ICP
Owner + AEs | Reviewed quarterly | ICP doc

Source leads — three channels
- LinkedIn Sales Navigator **(1 seat)**
- Apollo / ZoomInfo **(1 seat)**
- Referrals

List building & enrichment — by hand
BOTTLENECK 1

Account execs | 10 min / lead · 40 hrs / mo | Sheets + Apollo export

Qualification decision — firmographics only
BOTTLENECK 2

AE + owner | Batch, weekly | Sheets

- PASS → continues to outreach
- FAIL → removed from list

Write "customized" emails — templates, name swapped
BOTTLENECK 3

Account execs | 15 min / lead · 42 hrs / mo | Gmail

Send outreach
Account execs | Daily batches | Gmail, one domain

Response received? — 2.5% reply rate
BOTTLENECK 4

- NO RESPONSE → generic follow-up sequence, loops back
- YES → response type

Response type
- INTERESTED → manual follow-up
- BOOKING → booked discovery call
- NOT INTERESTED → archive

CRM handoff
Account execs | On booking | HubSpot

Secondary processes

Monthly client reporting

Export performance data from Meta, Google Ads, LinkedIn, GA4 — BOTTLENECK
Account exec | 30 min / client | 4 ad platforms

Paste into Sheets, fix formulas, write the same commentary structure — BOTTLENECK
Account exec | 75 min / client | Sheets + Docs

Account manager reviews → sends to client
Account manager | 15 min / client | Email

~1 hr 45 min per client × 22 clients ≈ **38 hours every month**, concentrated in the first week — the same week outbound activity stalls.

Internal communication & knowledge flow

Not a linear process — a pattern that surfaced in seven of nine interviews. Knowledge about clients, past campaigns, and the ICP lives in people's heads and scattered docs. Questions go into Slack DMs and wait on whoever holds the answer.

WHAT THE INTERVIEWS SURFACED	COST
"What did we run for [client] last quarter?" — answered from memory, hours later	Average ~3 hr wait on campaign-history questions
AEs re-ask the same ICP and positioning questions each month	Interruptions land on the two most senior people
Searching the drive, old decks, and Slack history for answers	Est. ~2 hrs/day team-wide ≈ **40 hrs/month**

WHY IT COMPOUNDS

The communication drag makes every other process slower — outreach waits on ICP answers, reporting waits on campaign history, onboarding waits on "how do we usually do this." Fixing it lifts everything else.

SECTION 3

The three constraints

Everything the interviews and process maps surfaced in outbound reduces to three numbers.

CONSTRAINT 1 · QUANTITY

Outbound is capped at ~240 leads a month

One Sales Navigator seat, one Apollo seat, one sending domain — and 40 hours a month of manual list building. The ceiling isn't demand or ICP size; it's seats and hands. The team cannot physically source or send more, so pipeline is flat regardless of effort.

CONSTRAINT 2 · QUALITY

30% of outreach effort goes to poor-fit leads

Lists are built on firmographics alone — industry, size, location. No intent signals: nothing about whether a company is hiring marketers, raising funding, switching agencies, or engaging with relevant content. Poor-fit leads absorb enrichment and writing time, then fail — after the effort is spent.

CONSTRAINT 3 · RESPONSE RATES

Reply rate is 2.5% — and falling

"Customized" emails are templates with the name and company swapped. Follow-ups are generic nudges. Recipients can tell — roughly 4 replies per month from ~170 sends, converting to about 2 discovery calls. The follow-up loop runs, but it adds noise, not signal.

THE CHAIN

These compound in sequence: limited quantity × diluted quality × weak response rate = a pipeline that can't grow. The recommendations attack all three at once — and each fix makes the next one work harder.

SECTION 4

The opportunity matrix

Every candidate scored on cost to implement vs business impact. Full ROI math per item in Section 5.

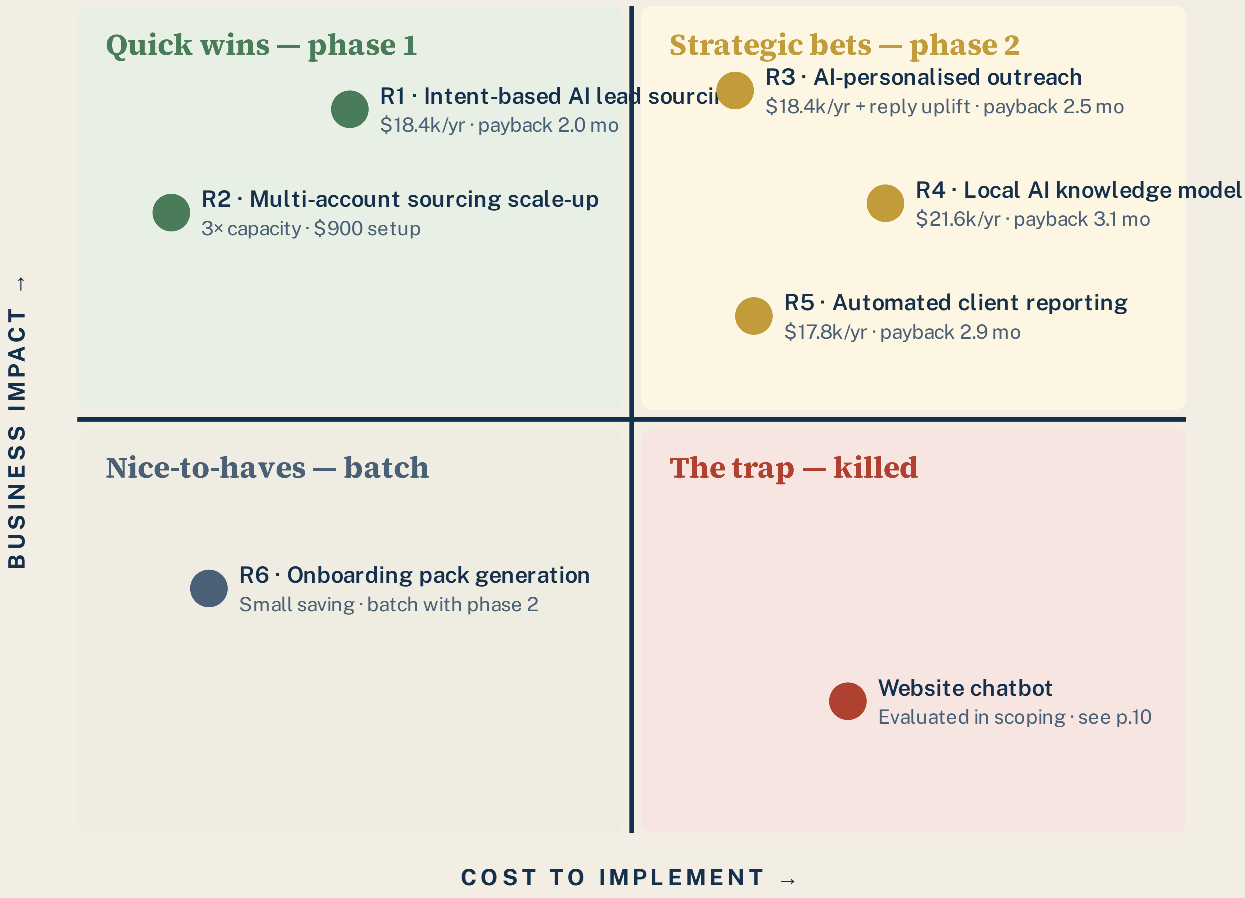

SECTION 5

Recommendations

R1 — Intent-based AI lead sourcing Quick win · fixes quality

Replace manual list building with AI scraping that targets **lead intent**, not just firmographics: companies hiring marketers, announcing funding, launching products, changing tech stacks, or engaging with relevant content. Each lead arrives pre-scored on intent, enriched, and deduplicated. AEs review a ready list instead of building one.

TOOLING	SAVING	BUILD	RUNNING	PAYBACK	RISK NOTES
Clay + Apollo intent data + Claude API scoring	34 hrs/mo list building · $1,530/mo · **$18,360/yr** — plus qualification pass rate up from ~70% to 85%+	$2,800	$150/mo	2.0 mo	Low–medium — B2B contact data only, sourced from licensed providers; AE review step retained

R2 — Multi-account sourcing scale-up Quick win · fixes quantity

The 240/month ceiling is seats and sending capacity, not demand. Add two further licensed sourcing seats and dedicated sending infrastructure — three warmed domains, six inboxes with SPF/DKIM/DMARC — lifting safe outreach volume to roughly 700 leads a month. Only viable because R1 and R3 remove the manual work per lead; without them, more volume just means more hours.

TOOLING	VALUE	BUILD	RUNNING	PAYBACK	RISK NOTES
Additional Sales Nav / Apollo seats + sending domains (Instantly-style infra)	~3× outreach capacity → est. +6 discovery calls/mo at improved reply rates	$900	$400/mo licences	<1 won client	Medium — domain warm-up takes 2–3 weeks; volume ramped gradually to protect deliverability

R3 — AI-personalised outreach Strategic bet · fixes response rates

For every qualified lead, an agent scrapes their **last six months of public activity** — LinkedIn posts, achievements, company news, job changes, published work — and generates a genuinely personal email and follow-up sequence referencing it. Not a template with a name swapped: a message that could only have been written to them. Sequences stop instantly on reply.

TOOLING	SAVING	BUILD	RUNNING	PAYBACK	RISK NOTES
Claude API + activity scraping/enrichment, sequenced via sending infra	34 hrs/mo writing time · $1,530/mo · **$18,360/yr** — reply rate est. 2.5% → 6–8%	$3,500	$120/mo	2.5 mo	Medium — every email AE-reviewed before send for the first month, then spot-checked; public data only

R4 — Local AI knowledge model in Slack Strategic bet · fixes communication

A **locally hosted AI model, tuned to Meridian's own data** — campaign archive, client docs, ICP definitions, SOPs — and integrated directly into Slack, where the questions are already being asked. "What did we run for [client] last quarter?" gets answered in seconds, in-channel, with the source linked. Local hosting keeps client data inside the building — nothing leaves for a third-party API.

TOOLING	SAVING	BUILD	RUNNING	PAYBACK	RISK NOTES
Open-weight local model (e.g. Llama via Ollama) + retrieval over the company drive + Slack app	~40 hrs/mo of searching and waiting · $1,800/mo · **$21,600/yr** — plus faster outreach, reporting, and onboarding	$5,000	$200/mo hosting	3.1 mo	Medium — answers cite sources so they can be verified; access mirrors existing drive permissions; client confidentiality preserved by local hosting

R5 — Automated monthly client reporting Strategic bet

Ad platform APIs feed a pipeline that fills the report template and drafts commentary against last month's numbers. AEs review and adjust; AMs approve. Human review retained deliberately — this is client-facing.

TOOLING	SAVING	BUILD	RUNNING	PAYBACK	RISK NOTES
Platform APIs + n8n + Claude, into Looker Studio	33 hrs/mo · $1,485/mo · **$17,820/yr**	$4,000	$100/mo	2.9 mo	Medium — commentary drafted, never auto-sent; AM approval mandatory

R6 — Onboarding pack generation Nice-to-have

Template-driven kick-off docs, folder structure, and Asana project from one intake form. Real but small — batch into phase two if capacity allows.

TOOLING	SAVING	BUILD	RUNNING	PAYBACK	RISK NOTES
Google Apps Script + Asana API	~4 hrs/mo · **$2,400/yr**	$900	$0	4.5 mo	Low

What we recommend you do NOT automate

KILLED — WEBSITE CHATBOT

Evaluated during scoping: $8–12k to build well, a handful of relevant website enquiries per month, and a relationship-led sales motion. High cost, low impact — the classic trap quadrant. Revisit only if inbound volume triples.

PARKED — FULL PROPOSAL AUTOMATION

Proposals carry pricing judgment and positioning decisions that lose deals when templated. Keep proposals human; a drafting assistant that pre-fills boilerplate is a sensible later experiment — full automation is not.

SECTION 6

Roadmap, timeline & cost

Estimates use PERT: best case (O), most likely (M), worst case (P), blended as (O + 4M + P) ÷ 6. The quote already contains the worst case — an overrun isn't a broken promise, it's inside the range.

Phase 1 — quantity & quality (weeks 1–3)

ITEM	BEST CASE	LIKELY	WORST CASE	PERT EST.	COST
R1 · Intent-based AI lead sourcing	4 days	6 days	11 days	**6.5 days**	$2,800
R2 · Multi-account sourcing scale-up	2 days	3 days	6 days	**3.3 days**	$900

Phase 2 — response rates, communication, reporting (weeks 4–8)

ITEM	BEST CASE	LIKELY	WORST CASE	PERT EST.	COST
R3 · AI-personalised outreach	5 days	8 days	15 days	**8.7 days**	$3,500
R4 · Local AI knowledge model in Slack	7 days	11 days	20 days	**11.8 days**	$5,000
R5 · Automated client reporting	5 days	8 days	16 days	**8.8 days**	$4,000

$16,200
TOTAL BUILD (BOTH PHASES)

$970/mo
RUNNING COSTS & LICENCES

$76,100
ANNUAL SAVING, YEAR ONE ONWARD

DELIVERY TERMS IN THIS EXAMPLE

Phase 1 quoted fixed at PERT estimates. Phase 2 scoped and re-quoted after phase 1 ships — by then the systems, data quality, and edge cases are known, and estimates tighten. Weekly progress note every Friday; no surprises.

SECTION 7

Next steps

Meridian now holds the complete roadmap — the processes, the priorities, the tools, the math. It can be handed to any capable builder. The consultant who wrote it is the obvious choice.

1. **Approve phase 1**
   Intent-based sourcing plus the capacity scale-up: $3,700, three weeks, and the quantity and quality constraints are gone.

2. **Ship, measure, report**
   Leads sourced, pass rate, and hours saved tracked against the estimates in this report. The numbers make the case for phase 2 on their own.

3. **Scope phase 2**
   Personalised outreach, the local knowledge model, and reporting automation — re-quoted with tightened estimates once phase 1 is live.

HOW TO USE THIS EXAMPLE

This report is the standard to hit. Notice what makes it land:

- The findings reduce to three numbers the owner already feels — quantity, quality, response rate.
- Every recommendation names a tool, a number, and a risk — nothing vague survives.
- Weak ideas were killed with evidence — including the tempting ones. That's the trust moment.
- Estimates are ranges, not promises — PERT protects both sides.
- Present it live. Walk them through it. Land on the quick wins. The audit does the selling.

SeanNoCode

The AI Audit Toolkit · seannocode.com

Meridian Digital is a fictional company. All figures are illustrative examples of audit-standard ROI math.