

THE LEAK AUDIT

Nine checks. A score out of 18 and out of 100. One workflow you run on your own records.

THE ECHOCODY OS LOOP · 01-05 BY DEPLOYMENT · 06 COMING

CAPTURE → IDENTITY → BRAIN → ACTION → LEDGER → INTELLIGENCE

We read the market reports. Most recycle each other. Two studies survive. Everything else in this file is the instrument — so you can price the leak and decide whether the fix is AI, a process, or a person.

INSIDE · PAGE 8 IS THE INSTRUMENT

- | | | | |
|----|-------------------------------|----|-----------------------------------|
| 01 | How to use this | 07 | The scored audit → /100 |
| 02 | AI in plain English | 08 | Operating contract + receipts |
| 03 | Prompts, practice, tools | 09 | Before AI touches a phone |
| 04 | Seven leaks + the two studies | 10 | What still matters as AI improves |
| 05 | Put a number on it | 11 | Seven-day challenge |
| 06 | Worked example | 12 | Book a review · sources |

PAGE 8 IS THE AUDIT.

Pages 2–7 teach you to run it. Pages 9–12 are depth, compliance, and the next step. Start with one useful task. Then one broken workflow. Then evidence.

You are not late to AI.

You are probably leaking money right now.

Most owners do not need to know how a model is trained. They need to know what AI is good at, where it goes wrong, and where it is worth pointing. Most businesses that adopt AI adopt it somewhere harmless — a better draft, a tidier email. Meanwhile the actual money is escaping through a handoff nobody has looked at in two years.

This guide runs the other way. Understand the tool. Find the leak. Measure it. Then decide whether AI, a process change, or a person is the fix.

01 UNDERSTAND

Know what AI is good at, where it can fail, and which terms actually matter.

02 PRACTICE

Use AI on low-risk work first so confidence comes from experience, not hype.

03 IMPROVE

Find one recurring handoff where better speed, ownership, or follow-up can be measured.

THREE IDEAS TO REMEMBER

- **AI is a tool, not an authority.**
You remain responsible for the decision.
- **Good results come from context.**
Tell it what you want, what to use, and what it must not invent.
- **Value comes from the workflow, not the answer.**
A clever response is nice. A reliable process with a measurable outcome is the business.

We will tell you when the answer is not AI.

Most free guides in this category are published by companies selling AI, and every one of them concludes that you should buy AI. On page 8 you will find a business whose largest leak was fixed by changing a lunch schedule. If the honest answer for you is that you do not need AI, you should be able to reach it using this guide.

02 · AI IN PLAIN ENGLISH

What AI is — and what it is not.

A useful mental model beats a hundred acronyms. AI feels sudden because the public experience changed quickly. The field itself has been building for decades. The important shift is not that AI suddenly became “smart.” It is that AI became easier for ordinary people to direct, and easier to place inside real workflows.

WHEN	SHIFT	WHAT CHANGED
1950s	The question	Can machines reason, solve problems, use language?
1980s	Expert systems	Specialist rules in software. Useful, rigid, hard to maintain.
1990s–2010s	Machine learning	Systems learn patterns from data instead of only hand-written rules.
2010s	Deep learning	Speech, vision, translation and prediction jump.
2020s	Generative AI	Ordinary language, files, images, audio and code.
Now	Workflow AI	Inside processes: tools, rules, human review, system updates.

AI

A very fast pattern-and-language engine that can help you think, draft, compare, summarize, create, and sometimes act through connected tools.

NOT THIS

Not an authority.

It can sound certain and still be wrong.

Not a database of truth.

Important facts still need verification.

Not a replacement for judgment.

Humans remain responsible for decisions and exceptions.

Not magic.

Better context, rules, examples and review usually produce better results.

SEVEN TERMS WORTH KNOWING

TERM	MEANING
Generative AI	Creates new text, images, audio, code, or structured output.
LLM	A language model trained to predict and generate language.
Prompt	The goal, context, inputs, rules, and output you ask for.
Agent	AI that can work through multiple steps or use tools.
Hallucination	A confident-sounding answer that is unsupported or incorrect.
Automation	A repeatable action triggered by rules or events.
Human-in-the-loop	A person reviews, approves, handles exceptions, or owns the final decision.

03 · PROMPTS, PRACTICE, TOOLS

Your first 30 minutes with AI.

Do not begin with your most sensitive process. Build confidence on ordinary work you already know well enough to judge. Do not paste passwords, private customer records, confidential contracts, or protected health data into an AI service until you understand that account's privacy, retention, and data controls.

- 1 Explain** "Explain this like I am a business owner, not a programmer. Define every technical term."
- 2 Rewrite** Paste a short email. Ask for clearer, shorter language that preserves meaning.
- 3 Compare** Give 2–3 choices. Ask for a table of differences, tradeoffs, and open questions.
- 4 Organize** Paste meeting notes. Ask for decisions, actions, owners, due dates, unanswered questions.
- 5 Challenge** "What am I assuming? What evidence would prove me wrong? What should I verify first?"

A PROMPT IS A MINI OPERATING BRIEF

FIELD	THE QUESTION
GOAL	What do you want to accomplish?
CONTEXT	What should the AI know about the situation?
INPUT	What information should it use?
RULES	What must it do or avoid?
OUTPUT	What should the final answer look like?
BOUNDARY	What should be flagged for human review?

COPY + PASTE STARTER

I own a service business. I want to review the last 30 days of inbound leads and find where follow-up broke down. Use only the records I provide. Do not invent missing facts. Separate what the records prove from what you infer. Return: (1) the three most repeated gaps, (2) evidence for each, (3) one simple process fix, and (4) anything that requires human review.

FIVE PROMPTS THAT MAKE AI SAFER

- 01 Translator** Explain this like I am an experienced owner new to AI. Define jargon the first time.
- 02 Evidence** Separate what this proves from what you infer. Do not invent missing facts.
- 03 Process map** For each step: trigger, owner, action, due time, system, what proves completion.
- 04 Leak finder** Review these records for handoff failures. Evidence, frequency, smallest fix to test.
- 05 Challenger** Strongest argument against your recommendation, assumptions, what evidence would change it.

LEARN IT PROPERLY, THEN PICK ONE TOOL

RESOURCE	WHERE	BEST FOR
Anthropic — AI Fluency	anthropic.com/ai-fluency	4Ds: Delegation, Description, Discernment, Diligence.
OpenAI Academy	academy.openai.com	Foundations → applied → agents and workflows.
Google AI Essentials	grow.google/ai-essentials	Workplace fundamentals. Certificate option.
ChatGPT / Claude	chatgpt.com · claude.ai	One general assistant. Start here.
NotebookLM	notebooklm.google.com	Questions across your own documents.
Zapier / Make	zapier.com · make.com	Connect apps only after the process is stable.

04 · THE SEVEN REVENUE LEAKS

Where the money actually escapes.

Revenue leakage is money your business had a reasonable chance to earn and lost because an opportunity was missed, delayed, mishandled or never finished. The seven tells below are the market. Trade blogs that publish a single annual-loss figure for a named trade are selling a fix. Read your records, not their average.

#	LEAK	THE TELL IN YOUR RECORDS
1	Missed calls	Inbound calls with no outbound callback the same day
2	Slow lead response	Form timestamp more than an hour before first outbound touch
3	No next action	A logged conversation with no task, no owner, no due date
4	Estimate follow-up	Quotes sent 14+ days ago still in the status they were sent in
5	No-show recovery	Appointments marked no-show with no rebook attempt attached
6	Old leads	Contacts with 90+ days of no activity, never formally closed
7	CRM handoff	Phone, calendar and CRM disagree about the same customer

Most businesses have all seven. Normal, and not the point. The point is which one costs you most — and you cannot know that until you price it.

TWO STUDIES SURVIVE. WE KEEP THEM SEPARATE.

2007 · LEAD RESPONSE MANAGEMENT

Oldroyd, MIT Sloan / InsideSales.com. Three years, six companies, 15,000+ web leads, 100,000+ call attempts. Measured contact and qualification — not close rates.

**5 min vs 30 min: ~100× contact odds, ~21× qualify odds.
5 min vs 10 min: contact odds drop ~5×.**

2011 · HARVARD BUSINESS REVIEW

Oldroyd, McElheran, Elkington. Audit of 2,241 US firms, one web test lead each — a stopwatch, not a survey. Separate qualification analysis.

37% responded in under an hour. 23% never. 42-hour average among those who did. Qualify: ~7× inside one hour vs waiting another hour; ~60× vs waiting 24h+.

WHAT WE REFUSE TO CITE

Trade blogs in 2025–2026 publish “annual leak” totals (\$45k, \$124k, \$167k) without a public method, then sell the answering product in the next paragraph. The 21× qualify figure is still credited to Harvard almost everywhere. It is the 2007 study. Both papers are real. They are different. A guide that cannot cite its own evidence has no business asking you to trust page 6. Price the leak from your records.

05 • PUT A NUMBER ON IT

A leak without a dollar figure is a complaint.

A leak with a dollar figure is a decision. The common formula — missed opportunities × close rate × job value — overstates the loss, because it assumes every miss is worth zero. Some call back. The honest figure is the difference between two close rates: what you close when you catch it, and what you close when you miss it.

THE THREE-MINUTE VERSION — FROM MEMORY

Roughly how many opportunities slip past you in a month? → _____

What is a customer worth to you the first time? → \$ _____

Multiply the two, then take 10% and 25%. Rough annual range: \$ _____ to \$ _____

THE FULL VERSION — RUN IT TWICE, REPORT THE RANGE

INPUT	CONSERVATIVE	REALISTIC
A Missed opportunities / month	_____	_____
B Share who were genuine buyers	_____ %	_____ %
C Close rate when you catch it	_____ %	_____ %
D Close rate when you miss it	_____ %	_____ %
E Average customer value	\$ _____	\$ _____
$A \times B \times (C - D) \times E \times 12 = \$ \text{_____} \text{ to } \$ \text{_____} \text{ per year}$		

The output is driven almost entirely by the gap between C and D — and D, your close rate on the ones you miss, is the number almost nobody actually knows. “Somewhere between \$40,000 and \$90,000 a year, and here is the arithmetic” survives a skeptical partner. “\$63,412” does not.

IF YOU ARE GUESSING, START HERE — THEN REPLACE THESE WITHIN A MONTH

INPUT	STARTING POINT
B — genuine buyers	40%
C — close rate when answered	50–70% for urgent work
D — close rate when missed	10–20%

These are borrowed numbers and we are telling you so. They come from published vendor calculator methodology, not peer-reviewed research, and the vendors publishing them sell the fix. They are a starting point for an owner with no data and nothing more. Then check the annual range against what a fix costs. If it sits below the cost of fixing it, document the leak and ignore it. That is a legitimate answer. If it is several times the cost, you have found this quarter’s project.

06 • WHAT THIS LOOKS LIKE ON A REAL WORKFLOW

Illustrative. Every assumption stated so you can swap in your own.

A seven-truck service company believes it is “pretty good” at answering the phone. The owner pulls one month of call logs instead of trusting that belief.

412

inbound calls

71

rang out to voicemail

9

got a callback the same day

62

got nothing, ever

	CONSERVATIVE	REALISTIC
A Missed opportunities	62	62
B Genuine buyers	30%	40%
C Close rate when answered	45%	55%
D Close rate when missed	20%	15%
Recoverable jobs	4.7	9.9
E Average first job	\$840	\$840
Annual leak	≈ \$47,000	≈ \$100,000

The owner reports the range, not the midpoint. Both ends justify acting.

The fix was not an AI receptionist.

The records showed 51 of the 62 came in between 11:30 and 1:30 — the office was at lunch and nobody was covering the line. The first fix was a staggered lunch schedule. It recovered most of the leak for the cost of a conversation. AI came second, for the after-hours remainder, where a human answer genuinely was not available.

This is why you audit before you shop. Sometimes the AI answer is right. It was not the right first answer here, and finding that out cost nothing.

07 · THE 30-MINUTE LEAK AUDIT

Nine checks. Score out of 18, then out of 100.

Pull the last 20–50 examples from the workflow you priced — last, not best. Use real records, not memory. If you cannot find the records, you have already found a leak: score line 1 zero and keep going. 2 = true on nearly every record. 1 = inconsistent. 0 = rarely or never. Index /100 = score × 100 ÷ 18.

#	CHECK	2	1	0	NOTES
1	Was the event captured?	☐	☐	☐	_____
2	Was the first response recorded, with a timestamp?	☐	☐	☐	_____
3	Was the outcome documented?	☐	☐	☐	_____
4	Was a next action assigned?	☐	☐	☐	_____
5	Was there a named owner?	☐	☐	☐	_____
6	Was there a due time?	☐	☐	☐	_____
7	Was the follow-up actually completed?	☐	☐	☐	_____
8	Was the system of record updated?	☐	☐	☐	_____
9	Can you verify the final outcome today?	☐	☐	☐	_____

YOUR SCORE

____ / 18

FIRST LEAK

____ / 100

ANNUAL \$ _____ to \$ _____

READING THE NUMBER · WORKING BANDS, NOT INDUSTRY BENCHMARKS

15–18 · 83–100

Instrumented

Visible process, records you can trust.
Ready to automate the highest-volume step.

9–14 · 50–78

Partially blind

Work is happening. Evidence is not. Fix lines 4, 5 and 6: owner, next action, due time.

0–8 · 0–44

Flying on memory

Do not automate yet. Get one workflow writing a timestamp and an owner.
Nothing else until that works.

08 · OPERATING CONTRACT + RECEIPTS

Build the workflow before you automate it.

If you cannot describe the process without AI, you are not ready to add AI. Vagueness in any answer is the thing to fix first. Capture → Connect → Act → Prove. Miss any one of the four and you have a hope, not a system.

#	FIELD	THE QUESTION
1	Trigger	What specific event starts this workflow?
2	Required information	What must be known before the next action can happen?
3	Allowed action	What can the person or system actually do?
4	Human boundary	What requires judgment, escalation, approval, or a live person?
5	Owner	Who owns the next action after the system acts?
6	Due time	When must the next step occur?
7	System update	Where does the outcome get written?
8	Completion test	What evidence proves the business result happened?

Field 4 deserves more attention than it gets. Write down before you build: an angry customer, a legal question, a price exception, a safety issue, and anyone who has asked you to stop. Those go to a human. Every time. Most bad AI deployments are a missing boundary, not bad technology.

EVERY IMPORTANT ACTION SHOULD LEAVE A RECEIPT

RECEIPT	WHAT YOU CAN SEE AFTERWARD
Capture	The call, form, message, appointment, estimate or task exists, with a timestamp.
Decision	The classification or routing decision is recorded, with the reason.
Action	The actual message, booking, task, transfer or system update is visible.
Ownership	A specific person or queue owns the next action and knows when it is due.
Outcome	The business records whether the desired result happened.

THE QUESTION TO ASK ANY AI VENDOR

"Show me the exact business event you will improve, the rule the system will follow, where the result gets written, who owns the exceptions, and how I verify the outcome myself." A vendor who answers that in one sitting is worth your time. One who redirects to capabilities, models or a demo is selling you software and leaving the workflow design to you.

This is the page most AI guides leave out.

It can cost more than the leak you are fixing. This page describes general principles and is not legal advice. Requirements vary by state and change frequently. Confirm obligations with qualified counsel before running any outbound program.

1 Federal rules already apply

Telephone consumer protection rules attach statutory penalties to each individual call or text, not to the campaign. Regulators have confirmed that AI-generated voices fall inside existing artificial-voice rules. One mistake is cheap. The same mistake automated across ten thousand contacts is not.

2 States do not agree

Several states now impose AI disclosure duties. They differ on what triggers them, who they cover, and whether disclosure must be proactive or only on request. Treat every state-law claim you read — including this page — as a pointer to a primary source, not as an answer. Statutes and effective dates are public.

3 Vendor blogs go stale

Statutes in this area have been amended, delayed, and in at least one case proposed for repeal inside eighteen months. There are live pages today stating requirements that were amended over a year ago, including pages published by companies selling AI voice agents.

FIVE QUESTIONS BEFORE YOU GO LIVE

CHECK	THE QUESTION
Consent	Can you produce evidence that each contact asked to hear from you?
Disclosure	Does the system identify itself as AI, in the opening, on every channel?
Opt-out	Has every stop-route been tested end to end, including one spoken aloud?
Quiet hours	Is the check against the contact's local time, not your office's?
Record	When someone opts out, is it written to the system of record, or only said out loud?

Question 3 fails most often, and it fails silently. An assistant can confirm out loud that someone has been removed and never write it anywhere. Test it with a real number before you trust it. The cheapest posture: saying plainly that the caller is speaking with an AI assistant costs one sentence and satisfies the strictest current standard. EchoCody is inbound only. Spoken name: Cody. First sentence: "This is an AI assistant, not a person."

10 · WHAT STILL MATTERS AS AI GETS BETTER

No predictions. A better question.

Anyone confident about what AI looks like in 2060 is selling something. If AI becomes as capable as its most optimistic advocates claim, which parts of running a business get easier, and which parts get harder? Whatever gets harder is where your durable advantage lives.

RUNG	WHAT THE SYSTEM DOES	WHAT YOU DO
1. Assistance	Drafts, summarises, suggests	You decide everything
2. Delegation	Completes bounded tasks end to end	You verify the output
3. Operation	Runs standing processes continuously	You handle the exceptions
4. Negotiation	Transacts with systems on the customer's side	You set the terms it may agree to
5. Institution	Holds durable responsibility over long horizons	You answer for it

Most small businesses are on rung 1 and believe they are on rung 2. Every rung still requires the eight fields on page 9 — not fewer as you climb. A system that operates unsupervised needs a sharper human boundary than one that only drafts.

THREE THINGS GET SCARCER

Proof

When generating a flawless proposal costs nothing, none of them persuade. A timestamp, a signed change order, a completed job with a customer who will pick up the phone — that holds value.

Judgment

Automation takes the repeatable middle. What is left on your desk is the tails. Plan for a business where your people make fewer decisions and each one matters more.

Accountability

A system can decide. It cannot be answerable. The human boundary never disappears. It migrates. For a service business, that is the product.

THE FUTURE-PROOF TEST · USE IT ON MONDAY

1. If the AI in this became ten times better overnight, does my advantage grow or vanish? 2. If it stopped improving tomorrow, was this still worth doing? 3. When it fails, who is answerable, and can they prove what happened? Something that survives all three is an investment. Something that needs the technology to keep improving is a bet. And if capability plateaus tomorrow, every recommendation in this guide still holds. Knowing where your money leaks, proving what your process does, and naming who owns the exceptions were good ideas in 1985.

11 · FROM AI-CURIOUS TO AI-CAPABLE IN ONE WEEK

Seven days. Two numbers most shops have never calculated.

DAY	MOVE	DO THIS
1	Learn the basics	First section of AI Fluency, AI Foundations, or AI Essentials.
2	Low-risk task	Something you know well enough to judge.
3	Practice verification	Prompt 02 — separate fact from inference.
4	Map one process	Trigger → Action → Owner → Due time → Evidence.
5	Score 20 records	Run the audit. Get your number out of 18 and out of 100.
6	Price the leak	Run it twice — conservative and realistic. Report the range.
7	Test one change	Measure before and after. Keep what works.

By Sunday you have a process score and an annual leak range. Confidence comes from practice, not from memorising terminology.

YOUR NEXT STEP

You do not need more AI. You need a better next action.

Free 30-minute Revenue Capture Review. Bring one workflow and 20–50 real records. We score it against the 18-point audit, price the largest leak with your numbers, and write a one-page operating contract for the fix — whether or not that fix involves AI, and whether or not you hire us. If you have no records yet, spend two weeks capturing timestamps and owners, then come back.

echocody.ai/book · (801) 980-0308



BOOK

WHY THIS GUIDE READS DIFFERENTLY

EchoCody is not an AI reseller. The method is evidence-first: inspect the records, compare what the workflow was supposed to do with what actually happened, then fix the highest-value gap before buying more technology. That is why page 7 ends with a lunch schedule and page 10 tells you not to fully trust page 10. A guide that only ever concludes “buy the thing I sell” is an advertisement with a table of contents.

Research: Lead Response Management Study, Dr. James Oldroyd, MIT Sloan with InsideSales.com (2007); Oldroyd, McElheran & Elkington, The Short Life of Online Sales Leads, Harvard Business Review (2011). Learning: Anthropic Academy, OpenAI Academy, Google AI Essentials (details verified against official pages, August 2026). This guide is educational. Results depend on your process, data, market, team and implementation. The compliance page is not legal advice. Benchmark defaults on page 6 are starting points from published vendor methodology — replace them with your own numbers. Review current privacy terms before sharing sensitive information with any AI service. EchoCody AI is a trade name of North Star Intelligence LLC, of Utah. North Logan, UT. Inbound only. echocody.ai · admin@echocody.ai · (801) 980-0308 · © 2026 EchoCody AI