A FIELD MANUAL FOR OWNERS AND REVENUE LEADERS

The Revenue Loop

How B2B Companies Grow When AI Does the Research

Tim Doelger

The Revenue Loop: How B2B Companies Grow When AI Does the Research

First edition, July 2026

This book shares methods, frameworks, and research for general informational purposes. It is professional education, and it is neither legal nor financial advice for a specific situation. Statistics are cited to their sources in the back of the book and were verified as of July 2026. AI platforms, search behavior, and buyer research change fast; check the sources for the latest figures before quoting them in a board meeting.

Contents

About this book

This book was written for the person who owns the revenue number at a B2B company.

Maybe you founded the company and still carry the bag yourself. Maybe you run sales for an owner who wants growth without another layer of management. Maybe your title says CRO, VP of Sales, or Director of Revenue, and the board asks every quarter what the AI budget produced. The details differ. The situation rhymes: buyers changed how they buy, AI changed how work gets done, and the old playbook now produces activity where it used to produce revenue.

The observations in this book come from my consulting practice, Get 'er Done, and from a career spent inside B2B sales organizations, mostly companies between $2 million and $50 million in revenue. What you will find here is what I have watched work, what I have watched fail, and what the current research says about both. The book covers why revenue goes flat while pipelines look full, how AI assistants now shape which vendors make a buyer's shortlist, what a knowledge and data foundation changes about every AI tool a company runs, where recovered selling time comes from, and how a revenue operating system holds the pieces together over time.

Seven worksheets sit in the appendix. They are the instruments I use in real diagnostic work, printed here in full: the Revenue Loop Score, the AI Efficiency Self-Audit, the Knowledge Base Checklist, the Pre-Contact Visibility Check, the Trust Filter, the MEDDIC deal review, and the One-Workflow Rule worksheet. They are yours to run, adapt, or ignore.

Every statistic in this book was checked against current research in July 2026, and the sources are numbered in the back. AI platforms, buyer behavior, and search surfaces move fast; the sources are there so you can check the latest figures before quoting anything in a board meeting.

How to use this book

Read straight through and the material arrives in the order it builds: what changed, what a foundation looks like, how companies run on top of one.

The parts also stand on their own, so the book can be entered wherever your current question lives. If buyers cannot find you, Part 3 covers how AI systems decide which companies to name. If your team drowns in admin work, Part 4 covers where the hours go and how companies have recovered them. If your forecast gets negotiated every quarter, Part 5 covers the operating system that computes one instead. Each chapter closes on something you can examine in your own records this week, most of it free.

One observation worth carrying in before you start. The companies I have watched get real results from AI kept their tool count small. They picked one workflow, proved it in dollars against their own records, and widened from there. That pattern shows up in every part of this book, and in my experience it separates a revenue system from a software collection.

INTRODUCTION

The shrug backed by a usage dashboard

If you spent $25,000 on AI for your revenue team this year, which line in your P&L would prove it?

Most owners I put that question to go quiet. The tools got bought in 2024 and 2025. The budget reviews are happening now. And the honest answer, in most companies, is a shrug backed by a usage dashboard. Logins are up. Prompts are up. Revenue held flat.

The research on this pattern has hardened. MIT's Project NANDA studied 300 public AI deployments and found that 95 percent of generative AI pilots showed no measurable P&L impact. S&P Global Market Intelligence found that 42 percent of companies abandoned most of their AI initiatives in 2025, more than double the rate a year earlier. Deloitte's 2026 enterprise survey of 3,000 leaders found that 74 percent of organizations want AI to grow revenue, while 20 percent report actually seeing it happen. Gartner forecasts that more than 40 percent of agentic AI projects will be cancelled by 2027 for lack of clear governance and return on investment.

The MIT finding rewards a second read, because it holds the whole problem in one place. More than half of generative AI budgets flow to sales and marketing. The measurable returns concentrate in operational work with clean before-and-after numbers. Revenue AI can pay. It is also the category where companies most often skip the measurement discipline that would prove it, because activity in a sales pipeline is easy to generate and hard to attribute. The spend goes where the excitement is. The proof lives where the discipline is.

While the tool purchases piled up, something larger moved underneath them. Your buyers changed. They do more of their research without you, and they do it with AI assistants. Gartner's 2026 surveys show a solid majority of B2B buyers now prefer a rep-free experience, nearly half used AI during a recent purchase, and AI-assisted buyers increasingly finalize their shortlists before any vendor hears about the deal; Chapter 2 carries the numbers. The shortlist forms before your team learns the deal exists. When the assistant never names you, the shortlist forms without you.

So the revenue question of 2026 has two sides. Inside the company, AI spend without measurement produces activity. Outside the company, AI-assisted buyers form opinions before the first meeting. This book takes both sides in the order I have found they need attention.

Part 1 lays out what changed, with current numbers. Part 2 covers the foundation every AI tool draws on: company knowledge, organized and governed. Part 3 covers being findable when AI does the looking. Part 4 covers selling time and execution built around judgment. Part 5 covers the operating system that connects it. Part 6 covers two topics worth separate treatment: government revenue for manufacturers, and a 90-day sequence that has held up when companies install all of it.

The owners who run this material well share one trait, and it is a willingness to measure. Technical depth shows up far less often than people expect. What follows starts with the problem most readers already feel.

PART 1

The Ground Shifted

CHAPTER 1

Pipeline Full, Revenue Flat

Leads are up. Deal values are up. Meetings are on the calendar. And revenue stayed flat.

If that sentence describes your last four quarters, you have company. It describes the majority of B2B businesses I diagnose, and it points at the most misdiagnosed condition in small and mid-market B2B. A full pipeline measures activity. Revenue measures whether the activity lands anywhere. The two came apart quietly, and the industry data says the separation is structural.

Optifai's 2026 Pipeline Study, drawn from CRM stage-level data across 939 B2B SaaS companies, found win rates holding the same range they have held for years: roughly 28 to 35 percent on deals under $10,000, dropping to 12 to 18 percent above $100,000. The median sales cycle stretched to 84 days, 22 percent longer than in 2022. More stakeholders, more security review, more rounds of internal approval. Very little of that registers as a stalled deal in a CRM. It registers as a pipeline that looks full and a quarter that closes soft anyway.

Martal Group's 2026 B2B Data Industry Report supplies the other half of the pattern. Spending on sales data and tooling kept climbing, and the growth stopped translating into pipeline gains. Access to data stopped being the bottleneck years ago. Acting on it while it is fresh became the bottleneck instead.

Put the two together and you get the defining pattern of this era: investment climbs, cycles stretch, win rate holds still.

The five places revenue leaks

The common response to a flat quarter is a push for more activity. More outbound, more tools, more meetings. With a pipeline already full, that lever tends to produce more of what the company already has. In the companies I diagnose, the leak sits in one of five places, and the fix starts with naming which one.

1. The front door. The pipeline fills with the wrong deals: companies that will never buy at your price point, never clear procurement, never feel the pain you solve. Marketing celebrates lead volume. Sales inherits leads that were never qualified to enter.

2. The middle. Deals enter well and stall in the quiet middle. No champion inside the account, no quantified pain, no agreed next step. The CRM says "proposal sent." The buyer moved on and was polite enough to skip saying so.

3. The close. Winnable deals die late. Price pressure in the eleventh hour, a competitor named in the final week, a CFO who was never briefed. These losses feel new each time. In the records, they are usually the same three causes repeating.

4. The leak after the signature. Customers who bought once and never expanded. Renewals that churn quietly. A base of happy customers whom nobody asks for referrals. Existing revenue is the cheapest growth a company owns, and in most companies it gets the least attention.

5. The mirror. The forecast gets negotiated. Reps commit deals that were never qualified, leaders discount the number by feel, and planning, hiring, and spending all run on a figure nobody fully trusts.

Effort stayed off that list on purpose. In the flat-revenue companies I meet, the team is almost always working hard. The effort points at the wrong leak, or spreads thin across all five at once.

What a diagnosis looks like before a dollar moves

Here is a pattern I have learned to trust: money spent before the leak has a name tends to buy speed for whatever already exists. A company leaking in the middle benefits from deal discipline and qualification, which Part 5 covers. A company leaking at the front door benefits from message and visibility work, which Parts 3 and 4 cover. A company with a mirror problem benefits from the operating system in Part 5. The companies I have seen end up with twelve AI subscriptions and the same flat line usually bought the tools first and named the leak later, if ever.

The owners who find their leak fastest run a version of the same exercise. They pull their last twenty closed-lost deals and their last twenty stalled opportunities, then sit with five questions:

1. Where in the pipeline did most of them die? Stage by stage, from the records.
2. For the deals that stalled, what was the stated next step, and who owned it?
3. How many had a named champion inside the account, someone with a personal stake in the solution winning?
4. How many had the buyer's pain quantified in dollars?
5. Which single cause, if fixed, would have changed the most outcomes?

Most owners who run this find their answer inside an hour. The pattern reads as obvious once it sits on paper. It stayed invisible before because it was spread across forty CRM records and four people's memories.

Where the rest of the book meets your answer

Each leak maps to a part of this book. Wrong deals at the front door: Part 3 on visibility and Part 4 on outreach quality. Stalls in the middle: Chapters 14 and 17 on trust and qualification. Late losses: Chapter 17 on MEDDIC and Chapter 18 on deal coaching. Post-signature leakage: Chapter 15 on the Learning Loop. The mirror: Chapters 15 and 16 on the operating system and the audit.

One observation before moving on. It is tempting to read the five leaks and conclude the problem is all five. In fifteen years of this work I have yet to diagnose a company where all five weighed the same. One leak always costs the most this quarter. The companies that moved fastest found that one, named it, and let the rest of the material wait its turn.

CHAPTER 2

Your Buyers Changed How They Buy

While the leaks in Chapter 1 were forming, your buyers rebuilt their entire purchasing process. Most of it now happens without you.

Gartner's 2026 buyer research, drawn from surveys of roughly 650 B2B buyers, puts hard numbers on what sales leaders have been feeling. Sixty-seven percent of buyers prefer a rep-free experience, up from 61 percent the year before. Forty-five percent used AI during a recent purchase, mostly to research vendors and products. Buyers consult an average of seven information sources before deciding. And buyers spend about 17 percent of their total purchase time in contact with all potential suppliers combined. For one of four vendors under evaluation, that works out to roughly 4 percent of the buyer's attention.

The rest of the buying process happens where sales teams cannot see it: AI assistants, peer networks, review platforms, internal meetings. Forrester's 2025 research found generative AI tools were the single most cited research method among B2B buyers. The first conversation a buyer has about your category increasingly happens with ChatGPT, Copilot, Perplexity, or Gemini.

The shortlist forms early

This is the change with the most revenue attached to it. Gartner found that buyers who use generative AI in their evaluation are 2.3 times more likely to finalize a shortlist before contacting any vendor. Read that as an owner. The list of companies that will be considered gets written with an AI assistant, in a chat window no vendor will ever see, before marketing knows the account is in market.

A company the AI names gets a chance to compete. A company it skips was excluded from a deal it never knew existed. No loss record lands in the CRM for that deal. No reason code. The revenue simply never appears. This is the reason Part 3 of this book exists: visibility inside AI answers has become a revenue function, and three chapters there cover how it works and what has moved it.

Buyers doubt the AI too

Here is the part of the research that I think should shape a company's whole response. Buyers use AI heavily and distrust it at the same time.

Forrester found that 20 percent of buyers, and 28 percent of procurement professionals, walked away from AI research less confident in their decision because the information was unreliable or inaccurate. Gartner's May 2026 survey found 69 percent of buyers turn to a sales rep to validate what the AI told them. When Gartner asked buyers where they expect to encounter misleading information, 51 percent said generative AI and 49 percent said a sales rep. That reads as nearly a coin flip, and it carries a warning for both software platforms and sales teams.

The pattern has a name worth remembering: confident misunderstanding. Buyers do extensive research, form firm conclusions, and arrive at the first meeting holding views that feel authoritative and are partly wrong. Forrester's broader data shows what that costs everyone: 81 percent of buyers end up dissatisfied with the provider they chose, and 86 percent of purchases stall somewhere in the process. Gartner adds that 74 percent of buying teams experience unhealthy conflict during the decision, and teams that reach genuine consensus are 2.5 times more likely to report a high-quality outcome.

What I have watched work in this environment is a specific posture: the seller as the buyer's most trusted validator. The first-source job went to the machines, and it shows no sign of coming back. When a buyer arrives holding an AI-generated comparison of the category, the rep who can say "here is what that gets right, here is what it misses, and here is the question worth asking instead" earns the trust that closes the deal. Gartner's own guidance points the same direction: buyers who reach genuine value clarity are twice as likely to report a high-quality deal.

The committee grew and went quiet

Two more structural shifts belong in any revenue plan. Buying committees now average around eleven stakeholders, up from seven a decade ago, and each added stakeholder lowers the probability of a purchase. Much of the committee's research happens in private channels: Slack groups, peer texts, communities no analytics tool will ever see.

Eleven people will rarely sit in one meeting with you. One or two will, and they will carry the message back to the rest. That turns the materials covered in Part 2, the documents that explain a company clearly with nobody in the room, into a sales asset. And it makes the trust behaviors in Chapter 14 the difference, because a champion's willingness to repeat your story inside their own company decides most of what happens next.

CHAPTER 3

The Adoption Paradox

Here is the strangest pair of numbers in B2B sales right now.

Salesforce's 2026 State of Sales report, which surveys more than 4,000 sales professionals, found that 87 percent of sales organizations now use AI. The same report found that 73 percent of B2B buyers actively avoid sellers who send irrelevant outreach. Nearly everyone holds the tools. Nearly three quarters of buyers tune out what the tools produce.

That gap is the adoption paradox, and I would trade any single tactic in this book for a reader who understands it.

Everyone bought the same engine

The adoption numbers describe a settled market. AI ranks as the number one growth tactic named by sales teams for 2026. Ninety-four percent of sales leaders running AI agents call them critical to meeting business demands. Top performers are 1.7 times more likely to use AI agents for prospecting than underperformers. Fifty-four percent of sellers have used an agent, and nearly nine in ten plan to by 2027. Salesforce itself pointed agents at 130,000 leads no rep had touched and created 3,200 opportunities in four months from leads that would have aged out in the CRM.

So AI in sales works, and it has become table stakes. The advantage moved. It now lives in what a company feeds the tools and where it points them. The same Salesforce report shows the raw material: reps using AI agents expect research time per prospect to drop 34 percent and email drafting time to drop 36 percent. Forty-eight percent of reps say they lack the bandwidth for adequate cold outreach. The tools genuinely give time back. The open question in every company is what the hours become.

What the 73 percent are avoiding

Read the buyer data closely and the target of the avoidance comes into focus: carelessness. AI itself barely registers with buyers. A buyer has no way to tell that AI helped write the email that references their actual tech stack, their actual compliance constraints, and their actual timeline. The same buyer can tell within seconds when an email found one public fact about their company and bolted a generic pitch onto it. The phrases repeat, and buyers in 2026 have seen them hundreds of times.

Practitioner data puts a size on the penalty. Analysis of more than a billion cold emails shows fully AI-generated copy producing roughly a 0.3 percent reply rate, against 4 percent or better for human-written copy to comparable lists. Industry-wide, cold email reply rates slid from around 6.8 percent in 2023 to roughly 4 to 5 percent in 2025 and 2026, partly because Gmail and Microsoft tuned their filters against AI-generated patterns. The volume play keeps getting weaker at exactly

the moment everyone can run it.

The 27 percent who still respond are responding to evidence that a real person did real thinking about their specific situation. I watched this up close with a client evaluating an enterprise purchase. Thirty-four outbound emails arrived in two weeks. Two got responses. The thirty-two that were ignored followed one pattern: a company name, a recent funding round, then a generic pitch. The two that landed demonstrated specific understanding of the buyer's environment and timeline. One of those two was drafted with AI assistance. The difference lived in the judgment applied to the output.

The reply-rate data confirms the same split at scale. Signal-personalized outreach, built on real triggers like a job change, a funding event, or a hiring pattern, produces 15 to 25 percent reply rates against the 3 to 5 percent industry average. The same tool category produces either 0.3 percent or 20 percent, and the judgment wrapped around it decides which.

The test every message faces

Buyers now run a simple credibility check on everything that reaches them: does this prove the sender understands my situation, or does it prove they own good automation?

Most AI-generated outreach fails the check because it was built for speed. It finds a detail and mentions it, with no way to know whether the detail matters, because it holds none of the account strategy, the competitive picture, or the politics of the buying committee. That context lives in a team, or it lives nowhere.

The reps I have watched win with AI use it to become more prepared. They take the research time savings and dig deeper. They let the machine draft, then add the judgment that makes the message credible. They send fewer emails, and each one lands with more precision.

What owners say they want

When B2B owners talk plainly about what they want from AI in sales, the list runs short and stays practical. Time back from admin work. Outreach that sounds like their best rep wrote it. Knowing which deals are real without a meeting to find out. New hires ramping in weeks instead of quarters. And someone willing to say what to skip buying.

The list carries no mention of more agents, more seats, or more features. Every item on it is an outcome, and every outcome on it is reachable with the material in this book. The tools are already good enough. The scarce ingredient is the operating discipline around them, which costs nothing and which few vendors have an incentive to sell.

Two messages, one prospect

Same prospect, same product. The first message is the kind the 73 percent delete: "Hi Sarah, I noticed your company is growing fast and thought our solution could help you scale even faster. Do you have fifteen minutes this week?" Nothing in it proves the sender knows anything. The name and

the growth mention read as merge fields, and the buyer feels it.

The second is the kind that earns replies: "Sarah, your team posted four openings for field service techs in the Midwest last month, which usually means the dispatch board is getting stretched. We helped a similar distributor cut scheduling time by about six hours a week per coordinator. Worth comparing notes?" A real signal, a specific inference, a verifiable claim, a low-pressure ask. AI can research the signal and draft the skeleton. A rep's judgment picks the signal that matters and writes the sentence that proves understanding. That division of labor recurs through every chapter that follows.

Quality is the strategy

One frame carries the rest of the book, and it is worth holding from here forward. AI handles research and synthesis. Humans handle judgment and relationships. It shapes the knowledge base in Part 2, the visibility work in Part 3, the workflow choices in Part 4, and the governance in Part 5. The companies I have watched get revenue from AI in 2026 pointed a small number of tools at well-defined work, with a human accountable for the output. Agent count had nothing to do with it.

The 73 percent remain reachable. They respond to a seller who did the work. The rest of this book covers how a team becomes that seller at scale, with AI doing the preparation.

PART 2

Foundations Before Tools

CHAPTER 4

The Knowledge Base Is the Asset

Here is an experiment I run with owners, and it takes five minutes. We ask whatever AI tool they already pay for to write a sales email for their company. Then we read the result closely. Does it know who the best customers are, what those customers worry about before they buy, what the company actually sells, and what proof it can stand behind? Or could the email belong to any business in the category?

For most B2B companies, the second answer wins. The tool writes something. It summarizes, it drafts, it produces a plan. Sounding like the company sits out of its reach, because nobody has told it what the company knows.

That gap explains most of what feels generic about AI output. The tool works without the company's offers, its customer questions, its best examples, its objection answers, its follow-up standards, its pricing logic, and its way of helping someone make a good decision. Handed none of that, it fills the space with the average of everyone else.

An AI tool produces useful, on-brand work when it can draw on the company's own knowledge. Organized once, that knowledge improves every tool the company touches. Skipped, it guarantees that every new subscription repeats the same generic output with a different logo on the invoice. This is why the knowledge base opens Part 2, ahead of visibility, workflows, and agents. It is the asset everything else draws from.

The cost gets paid every single time

When knowledge lives only in an owner's head, in old emails, in past proposals, and across a handful of people, the company pays for it repeatedly. It pays when a new rep ramps for six months because so little is written down. It pays when a proposal goes out with an old price or the wrong proof point. It pays when a good customer asks a question and gets three different answers from three different people. It pays when an AI agent, working from an empty or stale base, tells a prospect about something the company stopped offering a year ago.

Boston Consulting Group's 2026 research on AI in B2B sales framed the economics as the 70/20/10 rule: 70 percent of AI value comes from people and process, 20 percent from data and technology, and 10 percent from algorithms. Most of the budgets I see run in reverse order. The knowledge base funds the 20 and the 70 at once: the data layer, and the human ownership that keeps it true.

The shape of a base that works

The knowledge bases I have watched hold up share a specific shape: a governed collection of documents built so a new hire and an AI system can both find the right answer fast and trust what they find. A wiki gets partway there. A shared drive full of decks gets less far. Four dimensions

decide the outcome.

Accuracy. Every document reflects how the company sells today, as opposed to how a deck described it two years ago. **Coverage.** The base holds what sellers and buyers actually ask about, which differs from what was easy to write down. **Retrievability.** A person or a machine can pull the right passage in seconds. **Governance.** Every document carries one named owner and a review date. When any of the four slips, the base quietly loses trust, first with the team and then with the AI tools reading it.

One architecture decision comes with the territory. Knowledge can live in documents that AI tools read directly, or it can be indexed into a vector database for programmatic retrieval. For most B2B teams of five to fifty revenue-facing staff, the document approach inside an AI workspace the company already pays for has proven the right starting point: cheaper, easier to govern, readable by humans. The vector database earns its cost later, when retrieval volume and automation outgrow the simple setup.

Writing for the machine and the human at once

The single most useful design rule I know is the atomic principle: every document answers one question, addresses one scenario, or describes one entity. A document that tries to cover a whole manufacturing vertical, mixing ideal customer profile, case studies, pricing tiers, and competitor positioning, retrieves as a blob and frustrates everything that reads it. Four documents, one per topic, retrieve precisely.

The operational test fits in one sentence: what does this document answer? An answer that requires the word "and" points at a document ready to be split.

Structure carries the rest. Independent retrieval benchmarks published in 2026 by FloTorch and Firecrawl found that the plain approaches beat the clever ones: documents with clear titles, a summary at the top, consistent headings, and short sections outperformed exotic chunking schemes. The sales knowledge documents I see retrieve best all share the same skeleton:

- A clear title naming the entity or question. "Pricing: Professional tier (US, 2026)" beats "Pricing Overview."
- A one-sentence summary at the top. Retrieval systems and busy humans both read the first lines most.
- A named human owner. A person, since a department name behaves like no owner at all.
- A review date: when the document was last verified, and when it comes due again.
- The body, in plain language. Short sections, specific numbers, zero marketing filler.

What belongs in the base

The builds that stick start where retrieval demand runs highest and the habit change runs smallest. Four categories cover most of what a revenue team needs.

Offer and product. One document per thing the company sells: what it is, who it is for, what it costs, what it includes, what it excludes. A document that can state neither the price nor the reason there is no fixed price still has a section missing.

Proof. Case studies, references, results, and evidence a buyer can verify. Every claim a team makes in a deal traces to a proof document the AI can also find.

Objections and risk. The twenty questions buyers actually ask, with the answers the best rep gives. Security, implementation, contract terms, switching costs, "what happens if it does not work." This is where a seller earns the right to ask for the close, and it is the content AI tools lean on hardest when drafting responses.

Process and operations. How the deal actually moves: stages, qualification standard, handoff rules, follow-up expectations. This content makes a new rep productive and keeps an AI agent inside the company's rules.

Drafts, duplicates, outdated decks, and anything ownerless stay out. A smaller base of current documents beats a large one nobody trusts.

Maintenance carries most of the value

The build turns out to be the easy part. Keeping the base accurate is where most teams quit, and where the value compounds for the teams that keep going. Four habits show up in every base that has stayed alive.

One named owner per document. When ownership transfers, the new owner reviews within 30 days.

A review cadence by document type. Pricing and contract language every 30 days. Competitive comparisons every 90 days. Case studies every 180 days. Methodology annually. Trigger events force early review: a pricing change, a competitor launch, a lost deal naming that competitor.

Drift detection. When the base says one thing and the team says another, the base loses. The teams that catch this early assign someone to sample answers monthly: ask the AI ten real buyer questions and check the answers against the base. Wrong or stale answers trace to the source document and get fixed there, at the source, since patches applied in the chat evaporate.

Deprecation. Old documents get archived out of the active set. An AI that retrieves a 2024 price list in 2026 does more damage than no AI at all.

What a finished document looks like

An offer document, in full, might run 150 words: "Conveyor maintenance contract (Mid-Atlantic). For distribution centers running 40 or more hours of weekly throughput. Quarterly inspection, belt and roller service, four-hour emergency response, all parts included. $3,400 per month per facility, twelve-month term. Excludes motor replacement and controls work, quoted separately. Proof: three case studies in the proof library, including the Newark facility that cut unplanned downtime 31

percent. Owner: Dana. Reviewed: June 2026. Next review: September 2026."

That one document answers a buyer's question, feeds an AI writing a proposal, onboards a rep, and satisfies a retrieval system. Multiplied across offers, proof, objections, and process, the base exists. The work rewards precision far more than volume.

The honest cost model

Building the base costs owner and team time, roughly forty to sixty focused hours across a quarter for a typical SMB. The alternative carries its own price, paid forever: every AI tool producing generic output, every new hire ramping on folklore, every proposal rebuilt from the last proposal. One cost is finite and produces an asset the company owns. The other renews monthly.

A 90-day build that has held up

The builds I have watched succeed inside one quarter, without new headcount, ran in three phases. Days 1 to 30: inventory and triage. Collect what exists, retire what is stale, list what is missing, assign owners. Days 31 to 60: the canonical content. Write the atomic documents for offers, proof, objections, and process, with metadata on every one. Days 61 to 90: integrate and measure. Point the AI tools at the base, run retrieval tests with real buyer questions, fix what comes back wrong, and put the review cadence on the calendar.

The full checklist, including the metadata fields and the retrieval test protocol, sits in Appendix C.

CHAPTER 5

Keep the Keys in Your Hands

A knowledge base your AI reads from is an attack surface. That sentence sounds dramatic until the research sits next to it, and then it reads as planning.

In 2025, researchers at Penn State published work at the USENIX Security Symposium showing that an attacker who injects as few as five malicious texts into a knowledge base of millions of documents can make the AI system return an attacker-chosen answer, with a 90 percent success rate. The attack is called PoisonedRAG. The researchers tested existing defenses and found them insufficient. Related work at the same conference showed attackers can also jam retrieval with blocker documents, and 2026 follow-on research demonstrated the same poisoning pattern against AI security agents reading public write-ups.

Brought home, the stakes look like this. A revenue knowledge base holds pricing logic, objection answers, customer proof, and competitive positioning. An AI agent reads from it to answer buyers, draft proposals, and brief reps. Anyone who can write to that base without review can change what the company says. And any team member who pastes sensitive material into the wrong tool can move the company's positioning into a competitor's AI session. Both risks argue for governance, and governance runs cheap next to either outcome.

Three controls and a person

The security model I have seen work for a B2B revenue knowledge base has three controls and one human habit.

Narrow write access. Read access can stay broad. Write access stays a short, named list. Every document has an owner, and only owners and approvers modify their sections. This single control closes most of the poisoning risk, because poisoning requires write access.

Recorded provenance. Every document carries its origin: who wrote it, when, from what material, and when it was last verified. When an AI answer looks wrong, the trail leads to the exact source document in seconds. Without provenance, a wrong answer stays a mystery. With it, a wrong answer becomes a one-line fix.

Approved updates. Changes to pricing, contract language, competitive claims, and customer proof pass a second person before going live. Finance teams apply this discipline to payments as a matter of course. The base shapes what the company says to buyers, and it deserves the same treatment as what the company pays.

Then the habit: a person stays above the system. Once a week, someone owns a short review. They sample recent AI outputs that drew on the base, check them against the source documents, and log what they find. Fifteen to thirty minutes. This is the human-above-the-loop pattern from Chapter 3 applied to the company's own data: the machine does the work, a person owns the

judgment.

The leak that runs the other direction

The second risk flows outward: team members feeding sensitive material into tools that should never hold it. This one is already common. An owner pastes a customer list into a chatbot to draft an email. A rep uploads a proposal to summarize it. A manager drops pipeline notes into a tool nobody vetted. Each action reads as convenient in the moment. Together they hand the company's competitive position to systems outside its control.

The fix I have watched hold is a one-page policy. Three rules cover most of the risk. First, classify before pasting: customer data, pricing, contracts, and deal strategy go only into company-approved, access-controlled AI workspaces, and never personal accounts. Second, the knowledge base serves as the buffer: reps prompt against the governed base, which already holds the approved version, and skip pasting raw files. Third, one named approver decides which tools are approved, and the list stays short enough for an index card. The card gets posted where the team can see it.

A format choice reduces a third risk, lock-in. Canonical knowledge kept in open, portable formats, Markdown or plain documents with standard metadata, moves between AI platforms as the market shifts, with no rebuild. The knowledge is the asset. The tool is a reader. Companies that let the reader own the asset have discovered what that costs at migration time.

A six-point safety check

A base a company already runs can be checked in under an hour. Six questions:

1. Can you name every person who can write to the base, from memory?
2. Does every document show its owner and last review date?
3. Do pricing, contract, and competitive documents require a second person's approval before changes go live?
4. When an AI answer was last wrong, could you trace it to the exact source document?
5. Is there a written, one-page rule for what can and cannot be pasted into AI tools?
6. Is the canonical knowledge stored in a format the company could move to a different AI platform within a week?

Every "no" marks a specific, fixable gap. None of them requires new software. All of them require an owner, which by now reads as the recurring theme of this book: the technology arrived ready. The governance is the job.

CHAPTER 6

The Data Foundation

Here is a sentence I hear some version of every month: "We bought the AI tool, pointed it at the CRM, and the output was garbage." The tool takes the blame. In the audits I run, the tool earns it about one time in ten. The data underneath earns it the rest.

AI magnifies the quality of whatever revenue system already exists. Pointed at clean, connected data, it produces sharp research, accurate briefs, and useful drafts. Pointed at the average B2B CRM, where close dates read as wishes, fields sit empty, and the same account exists three times, it produces confident nonsense at scale. Salesforce's 2026 data makes the stakes concrete: 51 percent of sales leaders say disconnected systems are slowing their AI initiatives, and 74 percent of sales professionals now prioritize data cleansing and integration. Among high performers, that figure runs 79 percent. Among underperformers, 54. The teams getting value from AI fixed the data first, and the gap between those two numbers looks like more than coincidence to me.

A diagnostic that fits in two mornings

Before any foundation work, it helps to know how bad it is, and two mornings answer the question.

Morning one, the 90-minute data trust test. Pull twenty open opportunities at random. For each one, five checks: Is the close date real or a placeholder? Is the next step filled in, specific, and dated? Is the stage accurate based on what actually happened last? Is there a named contact with a role? When was the record last meaningfully updated? Each deal scores out of five. An average under four means the AI tools are reading fiction.

Morning two, the 60-minute knowledge layer test. Ask the AI tool ten questions a new rep would ask: What do we sell? What does it cost? Why do customers choose us? What are the top three objections and our answers? Count the answers that came back accurate, current, and specific to the company. Fewer than seven right points at the knowledge layer from Chapter 4 as the first project, ahead of any workflow automation.

The scores, written down, become the baseline. Everything in Part 4 gets measured against them.

The seven layers, in order

A revenue data foundation has seven layers, and they build in sequence. The companies I have seen skip ahead ended up with a dashboard nobody trusts.

1. **People and accounts.** One record per company, one per contact, deduplicated, with roles. Everything else hangs off this layer.
2. **Activity truth.** Calls, meetings, emails, and next steps logged where they happened, close to when they happened. A team that logs activity on Friday from memory runs its data on a week of guesses.

3. **Pipeline state.** Stages with entry and exit rules written down, close dates that mean something, and a next step on every open deal.
4. **Outcome data.** Win, loss, stall, and the reason, recorded consistently. This layer feeds the Learning Loop in Chapter 15, and most CRMs I open lack it entirely.
5. **Knowledge assets.** The governed documents from Chapter 4, linked to the deals and accounts they support.
6. **Signals.** The external events that change timing: job changes, funding, hiring patterns, contract renewals. This layer powers the outreach quality in Chapter 13.
7. **Metrics and forecasts.** Reports and projections computed from the layers below, untouched by hand. A number in a board deck that was edited in the spreadsheet reveals an unfinished foundation underneath it.

Most companies hold layers one through three in partial shape and almost nothing above them. That reads as normal, and the fix runs sequential: layer seven waits until layers one through four earn trust.

Deciding backward from the decision

The design principle that keeps all of this practical is decision-backward data design. It starts from the five decisions leadership actually makes: Where does the next dollar go? Which deals are real? Which reps need coaching? What will close this quarter? What is broken? Each decision gets a list of the fields and records its answer depends on. Those fields earn required-field discipline, review, and ownership. Everything else stays optional, and optional fields stay empty, so the honest move is a short list.

The same principle draws the boundary around the CRM. The CRM holds structured, decision-driving data: accounts, contacts, stages, amounts, dates, next steps, outcomes. Documents, call recordings, proposals, and the knowledge base live outside it, linked in. Teams that stuff everything into the CRM end up with a cluttered system nobody maintains. Teams that keep the boundary clean end up with a CRM that answers questions and a knowledge layer that explains them.

One last standard makes every later chapter work, and I have come to treat it as the foundation's keystone: the revenue truth model. When two numbers disagree, everyone in the company knows which system wins. The CRM beats the spreadsheet. The spreadsheet beats the anecdote. Written down, the hierarchy ends a whole category of argument. The forecast chapter shows why: a forecast holds only as well as the truth model underneath it, and most forecast arguments turn out to be two people citing two different sources of truth.

PART 3

Being Found When Ai Does The Looking

CHAPTER 7

How AI Agents Find Businesses

For the last twenty years, being found meant ranking in a search engine. Show up in Google results and you had a chance to win the conversation. That model still matters, and a second discovery system now runs alongside it, on different rules.

When a buyer asks ChatGPT, Copilot, Perplexity, or Gemini for companies that do what you do, the AI reads, synthesizes, and names a small set of businesses, often with a short description of each. The buyer sees an answer where the list of links used to be. Your company is either inside that answer, described accurately, or it sits outside the conversation entirely.

The scale of this channel has moved past speculation. BrightEdge measured Google AI Overviews on roughly half of US searches by February 2026, up from about 6 percent in early 2025, and Advanced Web Ranking measured the figure near 60 percent by April 2026. AI Overviews reach about 2.5 billion users monthly. ChatGPT reported 900 million weekly users in February 2026, and Sensor Tower data reported by Reuters showed its app crossing 1 billion monthly users in June 2026, the fastest consumer app to reach that mark. Perplexity serves around 100 million monthly users, per the Financial Times. Google's AI Mode crossed a billion monthly users on its own.

Meanwhile, the traffic math for websites inverted. Seer Interactive, tracking billions of impressions, found organic click-through on queries with AI Overviews fell by roughly two-thirds before partially recovering, and Pew Research measured users clicking a traditional result only 8 percent of the time when an AI summary is present, versus 15 percent without one.

One number in that research deserves an owner's full attention. On queries where an AI Overview appears, Seer found brands cited inside the answer earn roughly twice the clicks per impression of brands left out. The gap between cited and uncited has become the new page one.

How the machine reads you

Humans read pages. Agents parse structure. When an AI system evaluates a business for an answer, it works through three quiet questions: who this organization is, what problem it solves, and when it should be recommended. Unclear signals send the system on to a source where the signals are clear.

That evaluation draws on more than the company's website. AI systems assemble a picture from the site, its schema markup, business profiles, review platforms, directory listings, press mentions, and third-party articles. They cross-check. When the site says one thing, the LinkedIn page says another, and an old directory listing says a third, the machine loses confidence, and confidence is the currency that decides who gets named.

This is why strong AI visibility rests on structured clarity ahead of clever marketing language. A page that says "we provide commercial HVAC installation for buildings up to 100,000 square feet across

the Chicago metro, with four-hour emergency response" hands the machine everything it needs. A page that says "we deliver excellence through decades of combined experience" hands it nothing. Both read fine to a human skimming. Only one survives extraction.

The four signals that decide whether you get named

Across the sites I audit, four signals separate the companies AI recommends from the ones it skips.

Clear entity definition. The machine can state, in one sentence, who you are, what you do, where you do it, and for whom. That sentence stays consistent across the homepage, the profiles, and the schema. A test I offer owners: cover the logo and ask a stranger to describe the business after thirty seconds on the homepage. Whatever they say is roughly what the AI says.

Structured service descriptions. Each service gets its own page with an answer-first opening, specifics, and schema markup identifying it as a service offered by the organization. One page listing twelve services in bullet points gives the machine twelve weak signals. Twelve pages give it twelve strong ones.

Verifiable proof. Statistics, named sources, case studies, and credentials. The peer-reviewed research in the field known as GEO, the study of how content earns visibility inside AI-generated answers, found that adding citations, statistics, and quotations produced the largest visibility gains of any content change tested. Proof functions here as raw material: it is what the machine uses to justify naming you.

Consistent cross-references. Profiles, listings, and third-party mentions tell the same story the site tells. The machine treats consistency as a trust signal, because inconsistency is what fabricated businesses look like.

A worked example

Same company, two homepage openings. Version one: "Welcome to Apex Industrial. For over 30 years, our family-owned business has been committed to excellence, integrity, and customer satisfaction across a wide range of services." Every word is true and none of it is usable. An AI asked "who does conveyor maintenance in Ohio" can extract zero relevant facts from it.

Version two: "Apex Industrial Services designs, installs, and maintains conveyor and material handling systems for distribution centers and manufacturers across Ohio, Indiana, and Kentucky. Family-owned since 1994. 140 active maintenance contracts. Four-hour emergency response, guaranteed in writing." Now the machine knows what the company does, where, for whom, and why it is credible. So does every human who lands on the page. Clarity for machines and clarity for people turn out to be the same writing.

Human choice became a machine signal

One 2026 development points at where all of this heads. Google's Preferred Sources feature lets a person hand-pick the sites they want to see more of, and it expanded into AI Overviews and AI

Mode in May and June of 2026, alongside a Highly Cited badge for original, widely referenced reporting. Google reports that users click through to their preferred sources about twice as often.

Every other lever in search is something a company tunes until an algorithm rewards it. This one has no setting. The only input is a human deciding a company is worth choosing. Machine visibility has started to compound human trust. A business earns AI recommendations the same way it earns human ones: by being clear, useful, and verifiably good at what it does. The companies working from that premise are building assets. The ones chasing tricks are renting rankings.

CHAPTER 8

The Visibility Playbook

This chapter lays out the working sequence I have watched move a company's AI visibility: finding out where it stands, understanding why it is invisible where it matters, and closing the gaps across ninety days. It is written in enough detail to run without hiring anyone.

Phase 1: Baseline, one afternoon

AI visibility is measurable, and a baseline takes an afternoon. It starts with a prompt list: ten to twelve questions your buyers actually put to an AI assistant. Three categories cover most of it. Category prompts: "best commercial HVAC companies in Chicago" or "who installs conveyor systems in New Jersey." Problem prompts: "how do I reduce downtime on a packaging line." Comparison prompts: "[your company] vs [your top competitor]" and "is [your company] good."

Every prompt runs on the platforms your buyers use: ChatGPT, Google AI Mode and AI Overviews, Perplexity, and Claude. Copilot joins the list when buyers live in Microsoft environments. Three things get recorded per answer: whether you were named, whether the description was accurate, and who was named instead.

The results score on four dimensions. Discovery: how often you appear in category and problem prompts. Shortlist: how often you appear when the buyer asks for options. Credibility: whether the description is accurate and current. Risk: whether the AI says anything wrong, stale, or damaging. Most owners finish this exercise with a clear picture inside two hours, and most find at least one surprise: a competitor they beat in person winning the machine conversation, or the AI describing a service the company retired two years ago.

The full structure for this baseline, including the twelve-prompt template and the scoring sheet, sits in Appendix D.

Phase 2: Diagnose why

Invisibility has a short list of causes, and they rank by how often I find them. **Thin entity signals:** the machine cannot tell what the company does, because the pages never say it plainly. **No extractable answers:** the content was written for persuasion, so nothing clean exists to quote. **Weak third-party footprint:** nothing outside the company's own site corroborates it. **Inconsistency:** different profiles tell different stories. **Technical gaps:** missing schema, blocked crawlers, or pages that bury the answer under scripts and pop-ups.

The baseline results point at the cause. Named nowhere: entity and third-party problems. Named with wrong descriptions: extractability and consistency. Named on some platforms and skipped on others: usually technical, because platforms weight sources differently.

Phase 3: The fixes, ordered by impact

Weeks 1 to 2, claim the digital real estate. Free and high-impact. A complete Google Business Profile, LinkedIn company page, and the two or three directories the buyers actually use, with identical descriptions everywhere. This is the consistency layer, and it costs an afternoon.

Weeks 2 to 6, rewrite the core pages for extraction. Answer-first structure: every page and every section opens with a direct statement of fact. Evidence leads: real numbers and named sources on the best pages. A real FAQ page carries fifteen to twenty questions buyers actually ask, with answers in the forty-to-sixty-word range that extraction favors. One category-definition guide, the "what is this service and how do you choose a provider" piece, becomes the most durable citation asset a company can own. Credentialed authorship and visible last-updated dates round it out, because freshness now multiplies citations: ConvertMate's tracking found pages updated within the past month earning citations at roughly triple the rate of stale ones.

Weeks 4 to 8, add the machine-readable layer. Schema markup for the organization, services, FAQs, and articles. This is the explicit version of everything on the page, written in the vocabulary machines parse natively. It is tedious, it matters, and Chapter 9 covers why it matters more every quarter.

Ongoing, build third-party authority. The hardest signal to fake, and the strongest. Customer reviews on the platforms buyers trust, press and trade coverage, partner pages, association listings. Brands with active profiles on review platforms earn citations at multiples of those without. This work compounds slowly, and no shortcut I have seen survives a cross-check.

The two content assets that earn the most citations

Two of the fixes carry enough citation weight to deserve detail.

The FAQ page. A working FAQ differs from the decorative kind most sites carry. It holds fifteen to twenty questions buyers actually ask, pulled from sales calls and emails, each answered in forty to sixty words of plain, self-contained fact. "How much does conveyor maintenance cost?" gets "Most single-facility contracts run $2,800 to $4,200 per month depending on line count and response requirements. Emergency coverage is included, parts are included, motors are quoted separately." That answer is extractable: an AI can lift it whole and cite the source. The pages that work get written from the inbox, marked up with FAQ schema.

The category guide. One substantial page that defines the category and explains how to choose a provider, written to genuinely help a buyer decide, including the criteria where the company is a poor fit. A "how to choose a conveyor maintenance provider" guide that honestly discusses when an in-house tech makes more sense will out-earn any page that pretends otherwise. AI systems reward the same thing buyers do: the source that teaches without flinching. Published once and kept current, it works every day.

Phase 4: Track it like revenue

AI answers shift as models, sources, and competitors move, so the baseline lives as a loop. The lightweight version costs thirty minutes a month: rerun the prompt list, log the four scores, watch the trend. In June 2026, Google added a dedicated generative AI report to Search Console showing impressions inside AI Overviews and AI Mode, which gives every site free first-party data for the Google surface. Citation-tracking tools cover the rest whenever paying for them starts to make sense.

One habit costs nothing and catches more than any dashboard: be your own customer. Once a month, ask an AI assistant to help you buy what you sell, and follow the answer where it leads. A path that skips your company is the path your buyer is on.

A note on honesty

AI visibility moves. Answers vary by platform, prompt wording, location, and model updates, so any promise of a permanent position amounts to selling weather insurance. What a company can own is the baseline, the fixes, and the retest. That is the loop: measure, close the highest-value gaps, measure again. The companies running that loop quarterly are pulling away from the companies still debating whether AI search is real.

CHAPTER 9

The Machine-Readable Web Is Next

Everything in Chapter 8 makes a business readable to AI systems that visit and parse pages. The next layer is already in motion: a web built for machines to act, beyond just reading.

Two standards define it, and both crossed important lines in the past year.

MCP, the Model Context Protocol. Anthropic introduced MCP in late 2024 as a universal way for AI systems to connect to tools and data. In place of every AI needing a custom integration for every service, a service exposes its capabilities once through an MCP server, and any compatible AI can use them. OpenAI adopted it in March 2025, Google followed, and in December 2025 the protocol moved to the Linux Foundation under a new Agentic AI Foundation co-founded by Anthropic, OpenAI, and Block, with AWS, Google, Microsoft, and Cloudflare among the supporting members. The registry now lists thousands of servers, and the SDKs see tens of millions of downloads a month. When the companies competing hardest in AI agree on one plumbing standard, the plumbing is settled.

WebMCP. In February 2026, Google's Chrome team and Microsoft's Edge team published WebMCP as a W3C draft. It lets a website declare what it can do, directly to an AI agent, as callable tools. In place of an agent scraping a page and guessing what the "Book a consultation" button does, the site tells it: here is a booking action, here are the inputs it needs, here is what it returns. The agent calls the action. The site runs it. As of June 2026, WebMCP sits in a public origin trial in Chrome 149, which means real websites can test it on live traffic while the specification keeps evolving. Analysts project broader default-on availability in late 2026, and the Chrome team has committed to no public date, so the honest description is: real, moving, and early.

Alongside the standards, a simple convention keeps spreading: the llms.txt file, a plain-text index at the root of a website telling AI systems what the site is and where its important content lives. Rankability's adoption tracker put it at 8.7 percent of the web's top thousand sites by June 2026, up more than fivefold in twelve months, with Cloudflare, GitHub, and WordPress among the adopters. It costs an hour to create, and it adds one more structured signal to the stack from Chapter 7.

Why an owner might care now

Right now, AI agents mostly read and recommend. WebMCP and MCP form the plumbing for the next step, where agents act: checking availability, requesting quotes, booking meetings, initiating orders, on behalf of a buyer who never opened the website. A company whose services exist as structured, callable actions will be transactable by agents. A company that exists as a pile of pages will wait.

The response I have seen make sense runs at three speeds, and all three are deliberately boring.

Now, cheap: the Chapter 8 work continues. Structured content, schema, consistent profiles. Every one of those investments carries forward into the machine-readable layer, because the standards all consume the same clarity.

This year, small: publish an llms.txt. For a company with a booking, quoting, or ordering flow, one conversation with the web developer answers what it would take to expose that flow as a clean, structured action. In most cases the answer runs smaller than expected, and the work doubles as an accessibility and integration improvement.

Watch, without chasing: WebMCP tooling, agent directories, and the agentic commerce protocols the platforms are racing to define. The standards will keep shifting through 2026. The companies with a reason to experiment early are the ones whose buyers already transact through agents, which today means software and digital services more than manufacturing or professional services. For everyone else, the strong position is structural readiness when the buyers' agents arrive, and readiness is exactly the work in Chapters 7 and 8.

The pattern underneath all of this echoes Part 2: the companies that win each wave had their foundations in order before it arrived, and the ones that chased the wave itself mostly paid for the lesson. The machine-readable web rewards the same clarity that human-readable selling always did. Build once, benefit repeatedly.

PART 4

Selling Time And Execution

CHAPTER 10

The Selling Time Recovery

A question worth asking your reps, privately: how much of the week actually goes to selling? Their answer usually lands close to the research. Salesforce's data has put actual selling time at roughly a third of a rep's week for years, with the rest consumed by data entry, internal meetings, research, scheduling, and system maintenance. The bandwidth finding from Chapter 3 pairs with a stranger one in the same report: reps devote almost a full day a week to prospecting and still report coming up short.

The gap has a plain name: the most expensive employees spend most of their time on the cheapest work.

For a five-person sales team at a fully loaded cost of $500,000, recovering ten hours per rep per week equals adding a full seller without adding headcount. That is the reclamation math, and it explains why this chapter comes ahead of any talk about agents or automation strategy. The largest AI opportunity in most B2B companies looks less like a moonshot and more like getting back the hours already being paid for.

The four-step reclamation method

The method I have watched return the most hours runs in four steps.

Step 1: Identify the drains. For one week, each rep logs their work in thirty-minute blocks. One week, and done. The log surfaces the recurring tasks that fill the week without producing revenue: CRM updates, proposal formatting, meeting notes, scheduling, research rabbit holes, internal status pings.

Step 2: Sort by judgment. The list splits into two columns. Work that requires judgment, relationships, and accountability stays human. Work that runs repetitive, rules-based, and low-risk becomes the automation candidate list. Meeting notes, CRM hygiene, first-draft research, follow-up scheduling, and proposal assembly show up as the usual top five.

Step 3: Deploy against one drain at a time. The single biggest time drain gets fixed before the second gets touched. The failure pattern I keep finding is five half-finished automations. The winning pattern is one workflow, fully working, measured, then the next.

Step 4: Measure the recovery. Hours returned per rep per week, and what the hours became. The second number matters more than the first. Recovered time that becomes selling time shows up in pipeline. Recovered time that becomes longer lunches shows up nowhere. The teams that keep the gain put the reclaimed hours against a specific activity before the automation goes live.

The five recovery plays that pay first

Across the teams I work with, five plays consistently return the most hours for the least risk.

Meeting recovery. AI handles scheduling, reminders, notes, and CRM entry after calls. A rep with fifteen meetings a week typically gets back three to five hours, and CRM record quality improves at the same time.

Prospect intelligence. Research briefs generated before first meetings: the account, the people, the triggers, the likely objections. The State of Sales data shows reps expect AI to cut research time by about a third. The deeper gain: every rep walks in prepared, beyond just the diligent ones.

CRM hygiene automation. Field updates, deduplication, stale-record flagging, and next-step nudges running in the background. This is the unglamorous play that keeps Chapter 6's foundation real week after week.

Proposal acceleration. Drafts assembled from the knowledge base: the right boilerplate, the current pricing, the relevant proof. The rep edits for judgment, with the blank page already gone. Teams typically cut proposal time in half, and the documents get more consistent at the same time.

Forecast preparation. Pipeline summaries, week-over-week changes, and stall lists prepared automatically before the weekly review. The meeting turns from a readout into a decision session.

The two habits that multiply everything else

Two daily habits, borrowed from the owners who run this well, cost nothing and compound with every play above.

Habit one: give the machine context before asking for output. Who you are, who the audience is, what good looks like, and one example to imitate. Thirty extra seconds of input routinely doubles the usefulness of the output. The owners who describe AI as producing generic work are almost always giving it generic input.

Habit two: iterate. The first output serves as a draft. One round of "shorter, plainer, drop the adjectives, keep the number" gets most outputs to usable. Two rounds gets them to good. The skill turns out to be editing, which every owner already knows how to do.

When the calendar opens up

There is a moment coming for most teams that few owners have planned for. The automation works, the admin load drops, and suddenly the reps have open hours. What happens next decides whether any of this chapter mattered.

Unplanned, three things tend to happen. Work expands to fill the time. Reps spend the hours on more low-value activity, because busy feels safe. Or the best reps, the ones who carried the admin load as a tax on their selling, start interviewing at companies that will let them sell.

Planned, the recovered hours become the largest capacity increase a company will ever get for free: ten hours per rep per week, redirected into first meetings, follow-up discipline, and the

customer visits nobody had time for. The owners I have watched get this right decide the redirection before the automation ships, say it out loud to the team, and measure it after. They treat the reclaimed calendar as inventory to allocate.

CHAPTER 11

The One-Workflow Rule

Here is the rule that would have saved most of the AI money spent in the last two years: pick one workflow, price the leak in dollars, baseline it, fix it, and measure the result in your own records before funding anything else.

Recall the introduction's numbers: MIT's finding that 95 percent of generative AI pilots produced no measurable P&L impact, and Deloitte's gap between the three quarters of organizations that want AI revenue growth and the one fifth that report seeing it. The companies in the successful minority run fewer pilots, with the measurement discipline the rest skipped. A pilot with no baseline has no way to succeed, because even when the technology works, nothing exists to compare the result against.

The selection framework

Choosing the first workflow is the decision most teams get wrong, usually by starting with the tool. The companies that get it right start with the work, scoring candidate workflows on four tests and taking the one that passes all four.

The dollar test. Can the workflow's current cost be named, in hours or dollars? "Research takes my reps six hours a week each" passes. "We should be doing more with AI" fails. A leak that has no price offers no way to prove the fix.

The repetition test. Does the work happen every week, in roughly the same shape? Weekly recurrence means the savings compound and the measurement stays clean. One-off projects prove nothing.

The judgment test. Can the workflow split into machine work and human judgment? Good first workflows carry a large mechanical core with a thin layer of human review on top. Work that runs mostly on judgment belongs in coaching. Work that runs mechanical from start to finish scores even better.

The blast radius test. When the workflow produces a bad output, what happens? First workflows should fail cheaply: an internal brief that reads wrong, a draft that needs a heavier edit. Customer-facing, money-moving, and compliance-touching workflows come later, after the measurement habit exists.

Scored honestly, one workflow usually separates itself. In most B2B companies it comes from the five recovery plays in Chapter 10: meeting administration, prospect research, CRM hygiene, proposals, or forecast prep.

Baseline before building

Before anything changes, the workflow gets measured as it runs today. Hours per week, cost per run, quality as judged by a human, and the downstream metric it feeds, written down and dated. This takes a week of lightweight logging, and in my read of the Deloitte data, it is the week that separates the 20 percent from the 80.

One metric defines success, chosen in advance, in writing. A single metric keeps the verdict clean. "Research brief time drops from ninety minutes to twenty, with rep-rated quality unchanged" works as a success definition. "Reps feel more productive" leaves nothing to measure.

One person owns it. A single person, since every failed AI deployment I have audited shares the same org chart: everyone involved, nobody accountable. The owner runs the baseline, supervises the build, reviews the outputs weekly, and reports the metric.

Then the workflow runs for a defined period, usually thirty to sixty days, and gets judged only against the written metric. A hit produces a documented dollar value and clears the way for workflow two. A miss produces a cheap lesson and a baseline that makes the next attempt smarter. Either outcome moves the company forward. The only losing configuration I know is the unmeasured one.

Why the discipline is the strategy

The One-Workflow Rule feels slow. In practice it has been the fastest path I have watched, because it compounds. Workflow one returns hours and proof. Workflow two comes easier because the knowledge base, the data habits, and the owner's skill all carry over. By workflow three, the company holds something no vendor sells: an internal capability to evaluate, deploy, and measure AI against its own P&L. That capability is the actual asset. The workflows are where it practices.

CHAPTER 12

The AI SDR Reality Check

Somewhere in the last eighteen months, someone pitched you an AI SDR. The demo was impressive. The agent built lists, wrote sequences, sent them, handled replies, and booked meetings while you slept. Here is what the demo left out.

Jason Lemkin at SaaStr has watched more than twenty B2B companies deploy AI SDRs, and his summary runs blunt: 90 percent get absolutely nothing. Zero pipeline, zero meetings. The difference he identifies sits in ownership: the successful minority treat the agent like a $100,000 hire, trained daily, reviewed daily, managed by a human who owns the output, while the rest treat it like a $29-a-month SaaS tool. SaaStr's own deployment ran on the hire model, more than twenty agents under daily human management, and generated millions in pipeline across its first year.

Chapter 3 carried the failure mechanism: buyers pattern-match AI-written emails within seconds, and fully AI-generated copy pulls replies at a tiny fraction of human-written rates. The aggregate stick rate for fully automated deployments after ninety days runs around 2 percent. The pattern repeats: a few novelty replies in week one, then decay toward zero as the market learns to recognize the voice.

And the category's history carries a due-diligence lesson. In March 2025, TechCrunch reported that 11x.ai, one of the most funded AI SDR vendors, had claimed roughly $10 million in recurring revenue when revenue from contracts past their three-month break clause was closer to $3 million, and had displayed customer logos without permission. The company disputed the retention characterization. Either way, the story hands every buyer its first due-diligence question, and it leads the checklist later in this chapter. A vendor proud of retention answers it happily. Discomfort at the question is itself the signal.

When the tools actually work

The deployments I have seen produce results share five conditions. All five present makes a pilot worth running. Three or fewer usually means a human rep with AI research support will return more.

1. **A simple product with a one-sentence value proposition.** A sale that requires explanation makes AI copy read as shallow.
2. **A non-technical buyer.** Technical buyers detect and dismiss AI outreach fastest.
3. **A small, well-defined list, under roughly two thousand contacts.** Past that, personalization degrades into mail merge with better grammar.
4. **A human reviewing copy before every send.** Teams with human approval on drafts get three to five times the results of teams sending autonomously.
5. **A short sales cycle.** Under thirty days, the exposure from an AI-sounding message stays limited. Across a six-month enterprise cycle, the artificiality compounds through a dozen

touches.

The five conditions describe a high-volume, lower-ACV, transactional motion, and that is where autonomous agents earn their keep. The further a sale moves toward consultative, high-ACV, committee-driven, the more the economics favor humans doing the talking and machines doing the preparation.

Where AI fits in a complex sale

For the consultative, committee-driven sale, the fit map looks different from the SDR playbook, and it deserves plain statement because most vendor marketing skips it.

AI earns its place in a complex sale at the edges of the conversation: account research and trigger monitoring, meeting preparation, call summarization, CRM hygiene, proposal assembly, and follow-up drafting for review. The middle of the conversation still belongs to people: the discovery call where trust forms, the negotiation where judgment prices risk, the moment a champion decides whether to stake their reputation on you. Gartner's validation finding from Chapter 2 says the same thing from the other side: most buyers take what the AI told them to a rep to check. The more complex the sale, the more the human is the product.

So the complex-sale pattern reads simply: automation before and after the conversation, humans inside it.

The hybrid model that wins

The configuration producing the best economics in 2026 splits the work by what each side does well. AI handles the 80 percent of outbound that never needs to sound like a specific human: account research, enrichment, list building, lead scoring, reply classification, send-time selection, and first drafts. Humans own the 20 percent that touches the prospect: the final message, the reply that needs judgment, the call, the relationship.

The numbers keep pointing the same direction. One widely cited controlled test booked 847 meetings with an AI-only configuration at an 11 percent conversion rate, against 312 meetings in a hybrid configuration at 38 percent conversion. The hybrid produced roughly 2.3 times the revenue from fewer meetings. The exact multiple reads as directional, and the direction matches everything else in this chapter: volume behaves as a vanity metric, and revenue per meeting behaves as the real one.

The vendor due-diligence questions

Before signing with any AI sales vendor, these questions, with the answers written down, have saved my clients real money. What is your actual post-trial recurring revenue, and can I speak to three customers who renewed past the break clause? What reply and meeting rates do your median customers get, beyond the best case study? Who owns the sending domains, and what happens to my domain reputation if the campaign underperforms? What data of mine trains your models, and can I opt out? What does the human approval workflow look like in the product, and what does it

cost to use? What happens to my data and my sequences at cancellation?

A vendor who answers cleanly has nothing to hide. A vendor who pivots back to the demo has just described the next twelve months.

A 90-day pilot that has held up

For a company whose five-condition check says pilot, the sequence I have watched work runs in three gated phases, with widening earned at each gate.

Days 1 to 30, pilot on a narrow slice. One segment, one motion, a human approving every send. The test covers data quality, routing, and deliverability, with meetings as a later concern. Sends per mailbox stay capped, around twenty-five per day, because Gmail and Microsoft now actively filter AI-pattern bulk outreach, and a burned sending domain takes months to recover. Success in this phase means clean data.

Days 31 to 60, validate revenue. The approval gate loosens where the agent earned trust. Conversion and revenue per meeting get tracked against the human baseline. Multiple analyses report account executives closing AI-sourced meetings at a lower rate than human-sourced ones, so the gap gets measured before scaling, with qualification tightened if it shows.

Days 61 to 90, scale the hybrid pod. Only now does the funnel widen, at roughly one human per two agent seats, with a named person owning the agents daily. That ownership line is the single strongest predictor in every case study I have read, positive and negative. Every failure shares set-and-forget. Every success shares daily human management.

Run this way, an AI SDR becomes pressure on a motion the company already understands, owned by a team that keeps the judgment. Run without the discipline, it becomes a faster way to send the wrong message to the wrong people, with a monthly invoice attached.

CHAPTER 13

Ammunition, Not Automation

Gartner projects that by 2027, 95 percent of B2B seller research workflows will begin with AI, up from under 20 percent in 2024. In the best organizations, the cycle already runs. A trigger event happens: a funding round, a leadership change, a new facility, a contract expiration. Within seconds, AI has synthesized the signal, enriched the record, and drafted a relevant message. The research-to-outreach cycle that took a diligent rep most of a day now runs in under a minute.

That moves where the advantage lives. It used to belong to the rep willing to spend hours researching an account. That diligence became free. The advantage moved entirely to what the rep does with the research: the judgment about whether to reach out, what to say, and what to ask. I watched an early version of this at High Reliability Group, where I built outbound systems for consulting services and we talked with global industrial firms running some of the largest research operations in the world. The teams that won moved from signal to conversation fastest while keeping the credibility that comes from genuine understanding. Speed mattered, and it only paid when the message still landed as thoughtful.

This is why I teach ammunition over automation. Automation aims at sending more. Ammunition aims at sending better. The collapse of the research cycle hands a team more time to think before hitting send, more time to review the draft before it ships, and more time to prepare for the conversation after the reply arrives. The teams that spend the saved hour on judgment capture the entire gain.

Speed is cheap, accuracy is expensive

When every competitor can research and draft in seconds, the volume of noise rises and buyers raise their shields. Chapter 3 carried the numbers: 73 percent of buyers avoid irrelevant outreach, and reply rates on obvious AI copy run under half a percent. The barrier to sending a message has never sat lower. The barrier to sending one that gets a response has never sat higher.

The accuracy problem carries a second layer most teams miss: the research itself can be wrong. AI synthesizes from whatever data it can reach, and a CRM that says the contact left the company eight months ago sends the perfectly drafted message to a dead address with your name on it. Speed magnifies whatever it touches. It magnifies good data into fast, relevant outreach. It magnifies stale data into fast, confident embarrassment.

Three habits protect the cycle, and all three cost discipline, and none of them cost money.

Data hygiene as operations. The trigger signals feeding outreach decay weekly: job changes, bounced addresses, dead accounts. The Chapter 6 foundation work keeps the 60-second cycle pointed at real people, on a schedule, with automation where it fits and review where it does the most good.

A human in the verification loop. Before the first send to any account, a person checks three things: Is this person still there? Is the trigger real and recent? Does the message make a claim we can defend? Thirty seconds per account, and it separates a research cycle that produces ammunition from one that produces noise with the company's domain reputation attached.

Sawdust sorted from garbage. The Salesforce untouched-leads result from Chapter 3 worked because untouched differs from worthless; it means unworked. The teams that repeat the result sort the lead pile first. Sawdust: real fits nobody had time for. Garbage: bad fits that were never going to buy. AI sweeps sawdust beautifully. It will sweep garbage into the sending reputation just as fast.

The pre-flight check, and the apprenticeship problem

One more thing changed while research got fast: the teaching disappeared. Senior reps learned research the slow way, through thousands of hours of it, and junior reps learned judgment by watching them. With the machine doing the research in seconds, the apprenticeship quietly vanished. New reps get the output without the scar tissue that taught the seniors what to do with it.

The leaders I watch handle this treat it as part of the job now. They review AI-drafted outreach with reps before it ships, out loud, as a demonstration of judgment applied to real work. Why this trigger matters and that one carries nothing. Why this sentence stays and that one goes. Ten minutes a week, one message. That is how the next generation of a team learns the part the machine cannot do.

And before any campaign launches, the pre-flight check runs as a team: Is the data current? Is the trigger real? Is the message defensible? Is a human accountable for what ships? Four questions, ninety seconds. The teams that ask them are the ones whose 60-second research cycle produces revenue.

CHAPTER 14

Trust Is Still the Product

Everything in this book so far has covered machines: what they read, what they write, what they speed up. This chapter covers the thing they cannot produce, and why its value keeps rising as everything else gets automated.

The best sales lesson I ever picked up started at a fishing conference. Tim Sanders stood on stage at the American Sportfishing Association summit and told three hundred people that the biggest catches come from what happens below the surface. He was talking about trust. Trust built from knowledge, network, and compassion, shared generously.

That was years before buyers started deleting dozens of AI-generated emails a day. The water looks calm above the surface now. Below it, credibility still carries everything. Many leaders point at AI as the trust problem. What I keep finding underneath is a team using AI to look prepared, where the buyers wanted a team that was prepared. AI can tell a rep what to say. Making the buyer believe it stays out of its reach. Credibility resists automation; it can only be architected for.

Gartner's projection points the same direction: by 2030, three quarters of B2B buyers will prefer sales experiences that prioritize human interaction. Buyers keep the technology and price the difference between information and judgment.

The seven trust truths

These seven behaviors are the ones no tool performs on a person's behalf. They are also, and this reads as no coincidence, the behaviors that move serious deals.

1. Build your character first. Trust rests on character and competence, and character comes first. Buyers tolerate a learning curve far faster than they tolerate manipulation or half-truths. AI makes it easier to sound polished. Polish without grounded judgment creates suspicion.

2. Admit mistakes immediately. Every relationship includes a moment where something goes wrong. The seller who names the mistake before the buyer finds it converts a liability into a deposit. The one who lets the buyer discover it converts a small problem into a character reference nobody wanted.

3. Make consistent deposits. Trust accumulates through repeated, small, no-ask interactions: the useful article sent with zero pitch attached, the introduction that helps the buyer, the question remembered from three months ago. Relationships get built by a line of ordinary, generous interactions. A single brilliant meeting stays a dot.

4. Check your self-orientation. Buyers can feel whose interests a meeting serves. Every interaction answers a silent question for them: is this person here to help me decide well, or to close me? Performing helpfulness while tracking the commission fails the test fastest. Genuine orientation

toward the buyer's outcome passes it durably.

5. Listen to understand. Most discovery calls run as two people waiting for their turn to talk. The rep who listens for what is actually at stake, and plays it back accurately enough that the buyer says "exactly," earns access no sequence can buy.

6. Create safety to say no. Counterintuitive and powerful: the seller who makes it safe for the buyer to walk away gets told the truth. The truth about budget, about the competing offer, about the internal politics. Sellers who punish honesty with pressure get polite fiction, and polite fiction is what stalls deals in the quiet middle.

7. Dare to be vulnerable. The scripted seller stays replaceable. The one who says "I do not know, and I will find out today" or "we are a poor fit for this part, and here is who fits" buys credibility that outlasts any single deal.

The automation trap

There is a pattern worth recognizing before it costs a relationship built over years. A team adopts AI for follow-up and nurturing. Response times improve, touch frequency rises, every metric the dashboard tracks looks better. Then a long-time customer gets three automated touches in a week, each one cheerful, each one slightly wrong about their situation, and they call to ask if anyone still works there.

Speed spent the trust the relationship was built on, and nobody noticed because the activity metrics stayed green. That is the automation trap: the same AI that makes a careful team more present makes a careless team more absent. The guardrail states simply and runs on discipline: automate the mechanics, keep the meaning human. Anything the buyer will read as "a person from my vendor thought about me" gets a person behind it, every time.

The credibility scorecard

Where a team stands is measurable across six dimensions, and the reviews I have seen work best run on live deals.

Message honesty. Does every claim in the outreach and proposals survive a buyer's fact-check? **Discovery depth.** Do the reps understand the buyer's situation before prescribing, or pitch on the first call? **Follow-through reliability.** Do commitments made on Monday happen by Friday, every time? **Proof quality.** Can every cited result be verified by the buyer independently? **Human judgment in AI-assisted selling.** Does a person review what ships, and would the buyer agree it shows? **Handoff integrity.** Does the customer get what the deal promised?

Scored honestly across the last ten deals, the pattern that emerges is the company's trust position, and in my experience it predicts win rate better than any activity metric on the dashboard.

The irreplaceable revenue organization

Pulled together, the chapter yields a design principle for the whole revenue organization: automate everything mechanical, humanize everything relational, and keep the two lists separate. The companies building this way are becoming genuinely hard to compete with, because their efficiency funds more human attention. The best deal I closed last year started with a handwritten note. Real ink, ninety seconds, in the mail. The prospect called three days later. He had deleted forty-seven AI-generated sales emails that same morning. He kept the one envelope with handwriting on it.

That is the revenue organization this book keeps describing: AI doing the preparation, humans doing the moments that matter, and a team that knows which is which.

PART 5

The Revenue Operating System

CHAPTER 15

The Four Loops

Every fix in this book so far has been a component. This chapter covers the chassis they bolt onto. Without it, the components decay: the knowledge base goes stale, the data foundation silts up, the visibility work fades, the recovered hours drift back into admin. With it, the components feed each other and the system compounds.

The chassis is four loops. Every B2B revenue organization runs them, consciously or otherwise. The live question is whether they are closed.

The Signal Loop. Buyer signals flow in: website behavior, trigger events, conversation notes, intent data, market movement. In a closed loop, signals get captured, routed, and acted on while they are fresh. In an open loop, they die in inboxes, in notebooks, and in the memory of whoever heard them. The outreach quality from Chapter 13 lives here.

The Execution Loop. Work flows out: the right next step on the right deal, at the right time, by the right person. In a closed loop, reps act from a shared standard and the standard stays visible. In an open loop, everyone improvises and the week decides what happens. The workflow discipline from Chapters 10 through 12 lives here.

The Learning Loop. Outcomes flow back: wins, losses, stalls, and the reasons, captured honestly and reused. In a closed loop, the tenth deal is smarter than the first because the system remembers. In an open loop, the same three loss reasons repeat for years and every new rep re-learns them personally. This is the loop I find fully open most often, and it is the most expensive one.

The Forecast Loop. Reality flows up: leadership knows what is real, in time to act. In a closed loop, the forecast is computed from evidence and misses get examined. In an open loop, the number is negotiated, sandbagged, or wished for, and every plan built on it wobbles.

Seeing your own loops

The loops can be evaluated straight from records, and records settle what opinions argue about. The full version of this evaluation scores sixteen fields across the four loops, computed from a company's own CRM records, calendar records, and documented artifacts. No surveys, no invented benchmarks. A lighter version runs on four questions, one per loop:

Signal: Pick last week's three best buying signals. Can you show where each was captured, who it routed to, and what happened next?

Execution: Pick five open deals at random. Does each carry a next step that is specific, dated, and agreed with the buyer?

Learning: Pull the last ten closed-lost. Is the loss reason recorded the same way every time, and has anyone changed a play because of it this quarter?

Forecast: Compare last quarter's commit to last quarter's close. What was the gap, and what was done about the three deals that caused it?

Four questions, one honest afternoon, and the weakest loop shows itself. In the companies I have watched, that is where the next dollar goes. The sixteen-field worksheet sits in Appendix A.

The recursion underneath

A deeper pattern runs inside the loops, and it explains why some teams' AI spend compounds while other teams' spend evaporates. Learning itself is a loop with four stages, and AI accelerates whichever state it finds.

Data: deal outcomes flow back into the system cleanly. **Detection:** the system flags deals while they drift, ahead of the slip. **Decision:** when a flag fires, a human decides, and the decision gets logged. **Distribution:** when a tactic wins on one rep or one team, the pattern reaches the others without manual rewriting.

A company that closes those four stages gets a little smarter every week, and its AI tools get smarter with it, because the tools read the same records. A company that leaves them open buys the same lessons repeatedly. This is the quiet reason two companies can buy identical tools and get opposite results. The tools match. The loops differ.

The mistakes that kill the system

Five mistakes account for most failed installs I have examined, and every one is avoidable. Automating without verification: outputs ship unchecked, and one visible error costs more trust than a hundred good outputs earn back. Adding tools without subtracting tasks: the AI arrives, nothing leaves, and the team carries both. Measuring activity while outcomes go unmeasured: dashboards fill with sends and logins while revenue stays flat. Skipping the owner: everyone involved, nobody accountable, and the system quietly dies by week six. And the original sin, buying before diagnosing: the tool chosen before the leak was named speeds up whatever it touches, leaving whatever hurts untouched.

The install sequence below gives each mistake a gate that catches it, and the gates hold only when kept.

The install sequence

For a company installing this itself, the order carries most of the outcome, and the installs I have watched succeed all ran the same way. Diagnosis first: two weeks of leak mapping, data review, and deal-stall analysis before anything changes. Then the rules: stage definitions, qualification standards, next-step discipline, the knowledge base starter, and the governance rules for AI use. Then operation: weekly deal reviews running on the new standard, coaching on live deals, the first

clean leadership brief. Then transfer: an internal owner trained, the rhythm documented, the before-and-after measured against the company's own baseline.

Ninety days in that order has proven enough to close the loops in a typical five-to-fifty-person revenue team. The pieces built earlier in this book slot in along the way: the knowledge base feeds Signal and Execution, the data foundation feeds all four, the trust behaviors are what Execution is made of, and the audit in the next chapter keeps the whole thing honest.

CHAPTER 16

The Revenue Leak Audit

Every company in Chapter 1 had a leak. Most had a guess about where it sat. The Revenue Leak Audit replaces the guess with evidence, and the core of it fits in ten working days.

The method carries one rule that separates it from a typical pipeline review: every finding traces to a record. "Reps feel like deals stall after demos" stays a feeling. Records mean the last ninety days of stage durations, the actual emails, the real call notes. Feelings point at leaks. Records prove them.

Days 1 and 2: the leak map. The audit pulls every deal that touched the pipeline in the last ninety days. Open, closed-won, closed-lost, and the silent majority that just stopped moving. Charted by stage: how many entered, how many exited forward, how long each stage actually took. The shape of the funnel tells you where to dig. A bulge at one stage means deals arrive and wait. A cliff means they die at a specific moment. A long tail means nobody says no, to the deals or to the reps carrying them.

Days 3 and 4: the stall analysis. Every deal that sat longer than twice the normal cycle gets its record read in full. The reading looks for the five middle-leak patterns from Chapter 1: no champion, no quantified pain, no agreed next step, no economic buyer contact, no decision process mapped. Counted, the most frequent pattern becomes the first fix, and in my experience it is usually embarrassingly consistent. I audited one company where eleven of fourteen stalled deals shared the same cause: proposals sent to contacts who had never once mentioned budget. Nobody there needed new leads. They needed to stop writing proposals for tourists.

Day 5: the data truth check. The morning-one test from Chapter 6, run across the full open pipeline. Close dates, next steps, stage accuracy, contact roles, update recency. The score tells you how much of the pipeline reporting is real. Below an average of four out of five, every downstream analysis, including the forecast, is running on partial fiction.

Days 6 and 7: the tool and spend inventory. Every sales and marketing tool the company pays for, listed with its cost, its users, and what it produced that anyone can point to. Most owners find at least one tool nobody has opened in ninety days, and one duplicate doing another's job. Per S&P Global's research, the average company scraps nearly half its AI proofs of concept before production, which means the average stack carries the debris of abandoned experiments.

Days 8 and 9: the conversation. The two or three best revenue people, individually, with the same three questions. Where do deals actually die? What do you stop doing when you get busy? What would you fix first if you owned it? Their answers get compared to the records. Where the records and the people agree, the leak has been found. Where they disagree, the records win, and a coaching topic has been found.

Day 10: the readout. One page, since the deck adds nothing. The leak, the evidence, the dollar estimate of what it costs per quarter, the fix, the owner, the ninety-day target. That page is the audit.

Everything else was excavation.

The audit's output feeds everything after it. The leak map feeds the loops in Chapter 15. The stall analysis feeds qualification in Chapter 17. The data check feeds the foundation work in Chapter 6. The spend inventory feeds the One-Workflow decision in Chapter 11. Ten days of evidence, and the rest of the book knows exactly where to aim.

CHAPTER 17

Qualification That Holds Up

Ask a rep how a deal is going and the answer arrives as an adjective. Ask where the deal stands against a qualification standard and the answer arrives as evidence. MEDDIC is the standard I have seen hold up best for B2B deals with real stakes, and this chapter is the working version.

MEDDIC was built at PTC in the 1990s, where it helped take a $300 million revenue base to a billion in four years, and it has survived every sales fashion since because it runs as six questions a deal must answer before anyone believes in it. In 2026 it matters more: with the eleven-stakeholder buying committees from Chapter 2 and most of the buying process happening without the seller present, the deals that close are the ones where somebody mapped the machine early.

Metrics. The quantified outcome the buyer expects. "They want to improve efficiency" stays a wish. A number works: what it costs them today, what it will cost after, how they will measure it. A deal with no metrics carries no business case, and the first CFO review kills it.

Economic Buyer. The person who can say yes when everyone else has said maybe. In SMB deals this is often the owner. In larger ones, whoever owns the budget line. A contact who cannot or will not open that door marks the deal as a conversation with a recommender, and recommenders decline to sign.

Decision Criteria. The written and unwritten requirements the buyer will judge against. The direct question works: "What will you evaluate, and how will you score it?" Buyers respect it, because their own procurement team asks the same thing.

Decision Process. The steps, people, and timeline from here to signature, including the paper process: legal, security review, procurement. Deals rarely stall in demos. They stall in approval gates nobody mapped. A champion walking through the last purchase of similar size supplies the best predictor available.

Identify Pain. The specific, current, expensive problem driving the evaluation, and the cost of doing nothing. Pain is what moves the decision from someday to this quarter. No pain, no urgency, no deal.

Champion. A person inside the account with influence, access to the economic buyer, and a personal stake in the solution winning. A champion sells while you are absent, which, per Chapter 2, covers most of the buying process. The test: will they share the internal politics, the competition, the real timeline? A contact who will do none of that is a supporter. Useful, and different from a champion.

Two extensions round out the modern version. **MEDDPICC** adds Paper Process as its own element, reflecting how often deals now die in procurement and security review, and Competition, because buyers who generate AI comparisons before the first call have already priced the field, and a seller

who pretends otherwise loses ground with every meeting.

Making it stick

MEDDIC fails as a form and works as a habit, and the teams where it sticks share two practices.

First, the six elements live in the CRM as fields on every opportunity, with stage advancement depending on them. No economic buyer and no identified pain means no Stage 3. The fields mark the moment the standard becomes real.

Second, the deal review question changes. "How is the deal going" draws an optimistic adjective every time. The MEDDIC questions draw evidence. Who is the economic buyer, and when was the last conversation with them? Which decision criteria do we win on? What is the paper process? What is the champion saying about the competition? Deals with empty answers read as unforecastable deals, and surfacing that is exactly what the review exists to do.

Where AI helps, and where it stops

AI earns real value in qualification, in the preparation layer: researching the account, drafting the economic buyer hypothesis, summarizing the call, flagging the deal with no next step, scoring completeness across the pipeline. The qualifying itself stays with people. Building the champion relationship, reading the room, hearing the hesitation before "we should be able to get approval." Those are trust behaviors from Chapter 14, and they decide which deals are real.

The confidence metrics belong here too. Activity metrics count what reps did: calls made, emails sent. Confidence metrics measure what the buyer did: champion introduced the economic buyer, criteria shared, next meeting on the calendar with an agenda. Buyer actions predict outcomes. Rep actions predict busyness. The pipelines I trust get graded on the former. The full deal review checklist sits in Appendix F.

CHAPTER 18

Building Reps in the AI Era

The tools changed. The way a rep gets good held steady. A rep gets good the same way they always have: repetitions on real deals, with honest feedback, from someone whose judgment they trust. What changed is what the repetitions should be spent on, and how much of the old grind still deserves a place in the week.

The coaching model I have watched work in 2026 has four pieces, and they map directly to the earlier chapters.

Preparation coaching. Every rep walks into every meeting holding the AI-assisted brief from Chapter 10. The coaching lives in the review of that preparation. Beyond "did you read the brief," the questions run: what did you take from it, what are you going to ask, and what would make this meeting a success? The machine makes preparation universal. Coaching makes it useful.

Conversation quality. Real calls and real messages get reviewed weekly, against the trust behaviors from Chapter 14. Did the rep listen or wait to talk? Did they make it safe for the buyer to say no? Did the follow-up happen when promised? One call, one message, specific feedback, every week. This is the repetition that builds the judgment the machine cannot supply, and it also closes the apprenticeship gap from Chapter 13: seniors thinking out loud about real work.

Deal execution discipline. Live deal reviews run on the MEDDIC standard from Chapter 17. Where is the champion? What is the paper process? What did the buyer do this week that signals confidence? The rep learns qualification on their actual pipeline, which is the only place I have seen it stick.

Workflow ownership. Each rep owns their AI workflows with a verification habit: the draft reviewed, the facts checked, the judgment added before anything ships. The verification checklist runs short. Is this accurate? Is this current? Does this sound like us? Would I say this in the room? Four questions, applied every time, until reflex.

The weekly rhythm

Good coaching runs as a rhythm, and the calendar version is simple enough to hold without a system. Monday, thirty minutes: pipeline review on the MEDDIC questions, one deal per rep in depth. Wednesday, fifteen minutes: one call or one message reviewed together, with feedback specific enough to act on. Friday, ten minutes, async: each rep sends the three buyer confidence signals they earned this week; the activity log stays home. The rhythm totals under an hour of meeting time per rep per week, and it compounds because the standards get reinforced every single week.

Two rules keep the rhythm honest in the teams I have watched. The rhythm coaches live deals; hypothetical exercises produce hypothetical improvement. And sessions end with one action the rep

chose. Self-selected actions get done.

The ramp problem, and the teams that halved it

New rep ramp time sits among the largest hidden costs in B2B sales. Six months is common. Nine is unremarkable. Most of that time goes to learning the company: what we sell, what it costs, why customers buy, what to say when they push back. The selling skill itself takes a far smaller share. That is knowledge base content, and it is exactly what Chapter 4 built.

A new rep with a governed knowledge base, AI-assisted research, and weekly coaching on live deals ramps in a fraction of the usual time. The base answers the company questions on demand. The AI handles the research grunt work. The coaching hours go to judgment, with orientation handled by the base. Teams running this model report ramp compression of 40 to 60 percent, which for a team hiring two reps a year outweighs most tool subscriptions.

A certification habit makes it concrete. Before running a solo cycle, the new rep passes three gates: they can run the research workflow and defend the output, they can pass a mock discovery against the trust standard, and they can walk a manager through a MEDDIC-complete deal plan. Gates beat calendars. A rep who passes in week five starts in week five. A rep who has not passed by week ten gets coaching ahead of quota.

The setup that makes AI-booked meetings convert

One tactical piece deserves its own paragraph because so many teams stumble on it. AI can now book discovery calls: the agent finds the signal, sends the outreach, gets the yes, and puts it on the calendar. Then the rep walks in cold, the buyer realizes the rep knows less than the email implied, and the meeting dies in the first five minutes. The teams that convert these meetings run a handoff standard: no AI-booked meeting happens without a brief the rep has actually read, a hypothesis about why the buyer agreed, and one question the rep plans to ask that proves a human reviewed the file. The machine can open the door. A prepared human walks through it credibly.

CHAPTER 19

Leadership Options

Somewhere between five reps and twenty, most B2B companies hit the same wall. The owner can no longer run sales personally, the team needs real leadership, and every option looks expensive. A full-time VP of Sales at a growth company runs $250,000 to $400,000 fully loaded, takes months to hire, and fails at a rate that sobers anyone who has made the hire twice. Fractional leadership reads as a compromise. Doing nothing reads as safe, until the quarter says otherwise.

The math underneath the choice gets discussed less than it deserves.

The fractional math

A senior fractional sales leader typically runs $8,000 to $15,000 a month depending on scope, against a full-time hire's all-in cost of $25,000 to $35,000 a month once salary, bonus, equity, benefits, and tools are counted. Done right, the fractional leader arrives as a more senior operator than the company could afford full-time, working a defined scope: the operating system from Chapter 15, the deal reviews, the coaching cadence, the forecast discipline, the hiring profile when the time comes to scale.

The commitment objection deserves a straight answer, and the answer I have watched play out surprises boards. The worry runs that a fractional leader will care less than a full-time hire. In practice the accountability runs the other way. A fractional leader holds no political capital to spend, no empire to build, and no option to coast. Their entire business model rests on producing visible results at clients who can end the engagement on thirty days' notice. What boards need, underneath the presence question, is pipeline truth, forecast accuracy, and a plan that survives the next hire, and a good fractional leader delivers those faster because they have installed the same system before.

When fractional makes no sense

Honesty here matters more than any pitch, and three situations point away from fractional. When the team runs large and complex enough to need daily, full-time leadership: above roughly fifteen to twenty reps, the math and the attention both favor the hire. When the real problem lives in the product or the market: no leader, fractional or otherwise, fixes that from the sales chair. And when the owner wants a savior: leadership installed on top of chaos produces better-organized chaos, and only a system prevents that. The foundation chapters come first either way.

The hiring half of the equation

Leadership models fail when the seats hold the wrong people, and the cost of a bad sales hire keeps climbing: the salary, the ramp time, the deals lost while the seat underperforms, the second search. Most owners hire the way they were taught, on interviews and gut feel, then discover at

month six what a structured process would have shown in week two.

The hiring standards I have watched work carry the same shape as everything else in this book. The seat gets defined in writing before the search: deal size, cycle length, motion, and the ninety-day expectations. Candidates score against that definition, with charm scored at zero. Verification runs on proof: past numbers, references who managed them, a working session on a real deal scenario. Then the investment gets protected with structured onboarding and coaching from day one, because even the right hire fails in a seat with no system around it. Some recruiting arrangements now share this risk directly, pairing the placement with structured ramp support, which closes the gap most SMBs cannot staff alone.

Protecting pipeline when the team is in flux

Whatever leadership model a company chooses, transitions arrive: a rep quits, a leader changes, a restructure lands. The companies that lose pipeline in transitions share one trait: the relationships and the knowledge lived in the departing person's head. The defense I have watched hold has three layers, and most of it is this book.

Layer one: be findable. Inbound interest that comes from visibility work stays when an employee goes. The pipeline most exposed to turnover is the pipeline that was one person's network.

Layer two: knowledge above the individual. Accounts, histories, decision maps, and champion relationships documented in the CRM and knowledge base, current, and owned. When a rep leaves, the next rep inherits a file, complete and current. The human-above-the-loop structure means judgment transfers with documentation.

Layer three: leadership continuity. A documented operating rhythm, the weekly reviews, the forecast method, the coaching cadence, written down and owned by the company. When leadership changes, the system keeps running while the new leader ramps into something real.

With the three layers running, a departure becomes a staffing problem, which is solvable. Without them, it becomes a revenue crisis, which compounds. Transitions arrive on their own schedule, and the defense has to exist before it is needed.

PART 6

Special Plays And The Plan

CHAPTER 20

The Government Play

For a manufacturing company or an industrial services firm, there is a buyer with a legal obligation to pay on time, a budget set by law, and a public record of what it plans to buy. Most competitors will never sell to it, because the paperwork scares them off.

Ask owners why they consider government work and the honest answer is the invoice. The biggest commercial customers set the terms. They stretch payment to Net 60, 90, or 120 because they can, and the supplier floats the cost to keep the account. The federal government works the other way. Under the Prompt Payment Act, agencies must pay a proper invoice within 30 days, and late payment accrues interest automatically. The rate is published twice a year; the Federal Register set it at 4.75 percent for July 1 through December 31, 2026. An additional penalty, up to 100 percent of the interest owed, applies when the interest itself goes unpaid. Money earned from this buyer arrives on a clock, with statutory interest as the late fee.

The pool is real. Federal prime contracts worth $183.5 billion went to small businesses in fiscal 2024, a record per the US Small Business Administration, and contracts renew on a rhythm a company can plan around. One agency account can replace several flaky commercial ones. For a manufacturer watching private-sector orders soften, that looks less like a side bet and more like a second engine.

The honest warnings first

Anyone who describes this channel as fast is selling something. Three truths belong up front. It runs on documentation ahead of charm: wins come from answering the requirement exactly, with proof, in the format requested. Most failures are preventable: registration errors, wrong codes, missed deadlines, and generic capability statements kill most first attempts before competition does. And the costs come first: mobilization, materials, and compliance precede the first invoice, so the cash math needs planning even though payment itself is reliable.

The entry path, in order

The demand check. The government publishes what it buys and what it paid. An afternoon of research answers whether your products or services have real demand, at what volumes, and through which contract vehicles. Everything after this step depends on it.

Registration, once and correctly. SAM.gov registration, the right NAICS codes, and set-aside eligibility: small business, HUBZone, woman-owned, veteran-owned, service-disabled. The set-asides matter because entire categories of contracts are reserved for them, and eligibility a company already holds but never claimed is the cheapest advantage available.

A readable business. Contracting officers research vendors the way every other 2026 buyer does: online, increasingly with AI assistance. The capability statement, the website, and the past-performance record say the same thing in plain language, with codes, certifications, and differentiators machine-readable. Everything in Part 3 applies here, with a contracting officer as the buyer.

Known before the solicitation. The awards that look like upsets rarely are. They went to the company the buyer already knew. APEX Accelerators, the free government-funded counseling program in every state, exists to make those introductions and teach the process. The free help comes first because it is genuinely free and genuinely good.

Pursuit filtered to the winnable. The pattern among the winners I have studied: fewer bids, answered exactly. Opportunities filtered to the ones matching the company's codes, capacity, and geography, with the requirement answered to the letter. Ten precise responses beat fifty generic ones, in this channel more than anywhere.

Six questions before the first dollar

A readiness check worth running before investing in the channel. Do you have real capacity to deliver beyond the current book of business? Can you carry mobilization and material costs for sixty to ninety days before the first invoice pays? Is the SAM.gov registration accurate, or does it need to happen first? Do you hold, or can you document eligibility for, at least one set-aside category? Can you name the NAICS codes the work falls under, or find them this week? And would a stranger reading the website and capability statement understand the business, per the Chapter 7 standard?

Six yeses reads as ready now. Four or five points at a month of preparation. Fewer than four means the channel is real and belongs on next quarter's list, with the preparation work sitting right there in the questions.

The capability statement and the compliance gate

Two artifacts carry more weight than most first-timers expect.

The capability statement works as the company's one-page resume for contracting officers: what you do, the NAICS codes, the differentiators with proof, the certifications and set-asides, past performance, and contact details, in plain language on one page. It is the document that gets forwarded inside agencies, so the strong ones are written for the stranger who reads them with nobody present. Everything in Chapter 4 about precise, verifiable writing applies double here.

The compliance gate is newer and catches manufacturers off guard: cybersecurity requirements now gate many Department of Defense awards. A self-assessment against the required controls, done in month one, answers eligibility before months get spent pursuing a contract the company cannot yet accept. A pass becomes a selling point. A miss becomes a remediation list with a self-set deadline.

The part nobody expects

The assets built for government work pay twice. The capability statement sharpens how the company describes itself to every commercial prospect. The knowledge base that speeds a federal proposal speeds every quote. The machine-readable site that helps a contracting officer find the company helps every buyer's AI find it. The pipeline discipline that tracks a twelve-month procurement cycle makes commercial follow-up look casual, then fixes it. The companies I have watched build this channel properly ended up with a better commercial engine as a side effect.

CHAPTER 21

The 90-Day Owner's Plan

Every piece now sits on the table. This chapter offers the sequence, because in the installs I have watched, the order is what turned twenty chapters into one working system. The plan assumes a normal B2B company: an owner or revenue leader, a handful of sellers, a CRM in imperfect shape, and zero new headcount. Treat it as a map. The pace and the route stay yours.

Days 1 to 10: Diagnosis. The Revenue Leak Audit from Chapter 16 runs first: the leak map, the stall analysis, the data truth score, the tool and spend inventory, the one-page readout. The Pre-Contact Visibility Check runs the same week, so the AI visibility baseline exists too. End of week two, two documents exist: where revenue leaks, and what AI says about the company. Both came from evidence.

Days 11 to 30: Foundation. The knowledge base build from Chapter 4 starts: inventory, triage, owners assigned, the first canonical documents written for offers and objections. The two data tests from Chapter 6 run, and the worst twenty records get fixed. The one-page AI data policy from Chapter 5 gets published. None of this requires buying anything.

Days 31 to 60: First workflow and first visibility fixes. One workflow gets chosen by the rule in Chapter 11, baselined, and deployed. In most companies it turns out to be prospect research or meeting administration. In parallel, weeks one through six of the visibility roadmap run: profiles claimed, core pages rewritten answer-first, the FAQ built, the category guide published. The weekly deal review starts on the MEDDIC standard. Recovered hours start getting logged.

Days 61 to 90: Operate and measure. The workflow gets its thirty-day measurement against the written metric. The visibility scores get their first retest. The deal reviews produce the first honest forecast. The knowledge base hits its first review cadence. On day 90, the four loops from Chapter 15 get scored against the day-one baseline, and the next quarter's plan writes itself from the gaps.

What exists at the end of the quarter: a named leak with a fix in progress, a knowledge base with owners, a data foundation worth trusting, one measured AI workflow returning real hours, a visibility baseline with fixes underway, a qualification standard the deals actually get graded against, and a forecast computed from evidence. That reads less like a transformation program and more like a quarter of disciplined work, and it compounds.

Four starting points for this week

For a reader who wants motion before the full plan, four moves I have watched owners take first, any one of which stands alone.

Inspect the process where trust is built or lost. Three recent deals, read for where the buyer was deciding versus where the team was performing. That gap tends to be the first fix.

Protect the high-judgment work. The one workflow this quarter where a human decision moves revenue, with the hours for it protected ahead of any tool talk.

Treat the CRM as an operating system. The fields leadership actually uses get cleaned, the fields nobody touches get retired, and next steps become required. A CRM earns its keep by answering questions; storing intentions comes free.

Read the homepage as a buyer's AI would. If the homepage and the best sales page skip saying plainly what the company does, for whom, where, and with what proof, that becomes this month's writing project, and Chapter 8 holds the method.

The habits that keep it running

The system survives on four habits, none longer than an hour a week. The weekly deal review, thirty minutes, on the MEDDIC questions. The weekly AI output sample, fifteen minutes, someone above the loop checking what shipped. The monthly visibility retest, thirty minutes, the prompt list and the scores. The monthly knowledge review, one document type per cycle, owner verifying accuracy. A quarter without them and the system drifts. A quarter with them and it compounds, because every week the records get cleaner, the answers get sharper, and the next dollar of spend lands on a foundation that can hold it.

A note before the first ten days

The temptation most owners report is running three tracks at double speed and finishing in six weeks. The companies I have watched succeed with this material ran it at the pace written here, because the pace is what lets each layer hold. The companies that struggled shared a different sequence: tools first, foundation skipped, technology declared overrated when the flat line held.

The technology earned its reputation. The discipline remains underrated. Both now sit in your hands, and the first ten days are there whenever you decide to run them.

APPENDICES

The Worksheets

The seven instruments below are the working tools from my diagnostic practice. Print them, copy them, run them with your team, adapt them freely.

APPENDIX A

The Revenue Loop Score

Score each field 0 to 2: 0 = not happening, 1 = happening sometimes or only with a hero, 2 = happening every week without a hero. Compute from records, not impressions. Maximum score: 32.

Signal Loop

1. Buyer signals (site visits, trigger events, intent) are captured in the CRM, not in inboxes.
2. Every signal routes to a named owner within one business day.
3. Outreach references a specific, current signal in the majority of first touches.
4. Stale signals are purged or refreshed on a documented cadence.

Execution Loop

1. Every open opportunity has a next step that is specific, dated, and buyer-agreed.
2. Stage entry and exit rules are written and followed.
3. Reps walk into first meetings with a documented brief.
4. AI-drafted buyer-facing work passes human review before it ships, every time.

Learning Loop

1. Win, loss, and stall reasons are recorded the same way on every closed deal.
2. Someone reviews loss reasons monthly and changes a play because of it.
3. Winning tactics are documented and shared across reps within a month.
4. The knowledge base has an owner and a current review date on every document.

Forecast Loop

1. The forecast is computed from pipeline evidence, never adjusted by feel.
2. Commit accuracy (committed deals that actually closed) is measured every quarter.
3. Forecast misses get a named cause and a corrective action.
4. Leadership can state this quarter's confidence number without a meeting to prepare it.

Reading the score: 26 to 32, the loops are closed; protect the rhythm. 18 to 25, one loop is open; find it and spend the next dollar there. Under 18, start with the data foundation before anything else.

APPENDIX B

The AI Efficiency Self-Audit

Fifteen minutes. Five sections.

Section 1: Workflow audit. List the ten tasks that fill your week. For each, mark J (requires judgment) or M (mechanical). The M column is your automation candidate list.

Section 2: Knowledge base health. Check whether a new hire or an AI tool could answer these from your documents today: what we sell, what it costs, why customers choose us, the top objections and our answers, how a deal moves. Each "no" is a document to write.

Section 3: Master the inputs. Rate yourself 1 to 5 on the three habits that beat prompt tricks: giving context (your role, the audience, the definition of good), giving examples, and iterating past the first output.

Section 4: Declare your stack. Write down every AI tool you pay for, its monthly cost, and the workflow it serves. One tool should run most of your work. Anything without a workflow is a cancellation candidate.

Section 5: Time loss and priorities. Rate these eight drains 1 to 5 for your team: meeting scheduling and notes, CRM updates, prospect research, proposal writing, follow-up tracking, internal status reporting, data cleanup, forecast prep. Your top three scores are this quarter's targets, in order.

APPENDIX C

The AI-Ready Knowledge Base Checklist

Structure

- Every document answers one question or describes one entity (the atomic principle).
- Every document has a clear title naming the entity or question.
- Every document opens with a one-sentence summary.
- Every document names one human owner and a review date.

Coverage

- One document per offer, covering definition, audience, price, and exclusions.
- Proof library: case studies, results, references a buyer can verify.
- The twenty real buyer objections, with your best answers.
- Process documents: stages, qualification standard, handoff rules, follow-up expectations.

Governance

- Write access restricted to a named list.
- Provenance recorded: author, date, source material.
- Second-person approval on pricing, contract, and competitive documents.
- Review cadence by type: pricing 30 days, competitive 90 days, case studies 180 days, methodology 365 days.
- Monthly retrieval test: ten real buyer questions asked of the AI, answers checked against the base.
- Deprecated documents archived out of the active set.

Portability

- Canonical documents stored in open formats (Markdown or standard documents).
- The base lives in a workspace your company controls, never a personal account.

APPENDIX D

The Pre-Contact Visibility Check

Thirty minutes. The only tools required are the AI platforms your buyers use.

1. **Build the prompt list.** Four category prompts ("best [your service] companies in [your market]"), four problem prompts ("how do I [problem you solve]"), four comparison prompts ("[you] vs [competitor]", "is [your company] good", "who offers [your service] for [your niche]").
2. **Run every prompt** on ChatGPT, Google (AI Mode and AI Overviews), Perplexity, and Claude. Add Copilot if your buyers are Microsoft shops.
3. **Score four dimensions, 0 to 10 each.** Discovery: how often you are named. Shortlist: how often you appear when buyers ask for options. Credibility: whether descriptions are accurate and current. Risk: whether anything wrong, stale, or damaging appears (score in reverse: 10 means nothing wrong found).
4. **Record who gets named** in your place, and what the AI says about them.
5. **Write the headline finding in one sentence.** Example: "We are invisible in category prompts on three of four platforms, and our top competitor is named in all of them."
6. **Retest monthly.** Same prompts, same platforms, dated log.

APPENDIX E

The Trust Filter

Run any deal, message, proposal, AI workflow, or piece of content through these six questions before it ships. Three or more "no" answers means the work goes back to the desk.

1. **Diagnose.** Is the buyer's actual problem visible, specific, and quantified before any solution shows up?
2. **Quantify.** Is the cost of inaction on the page, in dollars?
3. **Teach.** Does the non-buyer leave smarter? (Only a small fraction of your market is buying right now. The rest still need to learn something from you.)
4. **Tie.** Is this connected to something the buyer cares about now: a current event, a recent change, a number they read this month?
5. **Deposit.** Does this strengthen the relationship before drawing on it? Deposits must outnumber withdrawals.
6. **Repeat.** Is this a sustainable cadence? Trust builds along a line of touches; a single touch stays a dot.

APPENDIX F

The MEDDIC Deal Review

For every deal past your qualification stage, answer in writing. Empty fields are coaching topics, never forecast entries.

- **Metrics:** What is the quantified outcome the buyer expects, and how will they measure it?
- **Economic Buyer:** Who can say yes? When did we last speak with them?
- **Decision Criteria:** What will they evaluate, and which criteria do we win on?
- **Decision Process:** What are the steps, stakeholders, approval gates, and timeline to signature? What does the paper process look like?
- **Identify Pain:** What specific problem is driving this now, and what does doing nothing cost?
- **Champion:** Who sells for us inside the account? What have they done in the last two weeks that proves it?
- **Competition:** What else are they evaluating, including doing nothing?
- **Buyer confidence signals:** What did the buyer do this week (not what did we do)?

Deal rule: below four of six MEDDIC elements complete at mid-pipeline, the deal stays out of the commit forecast.

APPENDIX G

The One-Workflow Rule Worksheet

The candidate: Name the workflow in one sentence. ("Reps research accounts before first meetings.")

The dollar test: What does this workflow cost today, in hours per week and dollars per month?

The repetition test: How many times per week does it run, in roughly the same shape?

The judgment test: Which parts are mechanical (machine) and which require judgment (human)? List both.

The blast radius test: If the output is wrong, what happens? (Internal-only consequences for a first workflow.)

The baseline: Hours per week now ____ · Cost per run now ____ · Quality score now ____ · Downstream metric it feeds ____

The success metric (one, in writing, in advance):

The owner (one person):

The review date (30 to 60 days out):

The verdict: Hit / missed, the dollar value proven, and the next workflow in line.

Free Tools and Further Reading

Interactive versions of the worksheets in this book, along with the full library of research and field guides behind these chapters, are free at geterdone.ai/tools/. No email gate on the tools. Use whatever helps.

Sources

Statistics and research cited in this book, verified as of July 2026.

[1] MIT Project NANDA, "The GenAI Divide: State of AI in Business 2025." 95 percent of generative AI pilots showed no measurable P&L impact.

[2] S&P Global Market Intelligence, 2025. 42 percent of companies abandoned most AI initiatives, up from 17 percent the prior year; roughly half of AI proofs of concept scrapped before production.

[3] Deloitte, "State of AI in the Enterprise," 2026 (survey of 3,000 leaders). 74 percent of organizations want AI to grow revenue; 20 percent report seeing it.

[4] Gartner, 2025 to 2026 press releases and surveys. 67 percent of B2B buyers prefer a rep-free experience, up from 61 percent (March 2026, survey of 646 buyers); 45 percent used AI during a recent purchase; 69 percent validate AI-generated insights with sales reps (May 2026); 51 percent expect misleading information from generative AI versus 49 percent from a sales rep (May 2026); buyers using generative AI are 2.3 times more likely to finalize shortlists before vendor contact (2025); buyers consult an average of seven information sources; buyers spend 17 percent of purchase time with all suppliers combined; buying committees average roughly eleven stakeholders; 74 percent of buying teams show unhealthy conflict, and consensus teams are 2.5 times more likely to report a high-quality outcome; by 2030, 75 percent of B2B buyers will prefer human-led sales experiences (August 2025); by 2027, 95 percent of seller research workflows will begin with AI, up from under 20 percent in 2024; more than 40 percent of agentic AI projects forecast cancelled by 2027.

[5] Forrester, State of Business Buying and 2025 buyer research. 81 percent of buyers dissatisfied with the chosen provider; 86 percent of purchases stall; generative AI the most-cited research method; 20 percent of buyers and 28 percent of procurement professionals less confident after AI research.

[6] Salesforce, "State of Sales Report 2026" (survey of more than 4,000 sales professionals). 87 percent of sales organizations use AI; 73 percent of B2B buyers avoid sellers who send irrelevant outreach; 94 percent of leaders with agents call them critical; top performers 1.7 times more likely to use AI agents for prospecting; agents expected to cut research time 34 percent and drafting time 36 percent; 48 percent of reps lack bandwidth for cold outreach; 54 percent of sellers have used an agent; 130,000 untouched leads yielding 3,200 opportunities in four months; 51 percent of leaders cite disconnected systems; 74 percent prioritize data cleansing (79 percent of high performers versus 54 percent of underperformers); actual selling time at roughly one third of a rep's week.

[7] Seer Interactive, "AIO Impact on Google CTR: 2026 Update" (2.43 billion impressions). Organic click-through on AI Overview queries fell roughly two-thirds before partially recovering; brands cited inside the answer earn roughly double the clicks per impression of uncited brands.

[8] Pew Research Center, July 2025 (68,879 real searches). Users click a traditional result 8 percent of the time when an AI Overview is present, versus 15 percent without one.

[9] BrightEdge, February 2026. AI Overviews on roughly half of US queries, up from about 6 percent in early 2025; reach of about 2.5 billion users monthly.

[10] Advanced Web Ranking, April 2026. AI Overviews measured on roughly 60 percent of US searches.

[11] OpenAI, February 2026 (900 million weekly ChatGPT users); Sensor Tower data reported by Reuters, June 2026 (ChatGPT app crossing 1 billion monthly users). Perplexity at roughly 100 million monthly users per

the Financial Times, April 2026. Google AI Mode above 1 billion monthly users per Google.

[12] Aggarwal et al., "GEO" (KDD 2024), the peer-reviewed study of content visibility in AI-generated answers. Adding citations, statistics, and quotations produced the largest visibility gains of tested content changes.

[13] ConvertMate citation tracking, 2026. Pages updated within the past month earn citations at roughly three times the rate of stale pages.

[14] Google, May and June 2026. Preferred Sources expanded to AI Overviews and AI Mode, with the Highly Cited badge; users click through to preferred sources about twice as often; Search Console generative AI performance report announced June 3, 2026.

[15] Model Context Protocol documentation and Linux Foundation announcements. MCP launched November 2024 by Anthropic; OpenAI adoption March 2025; donated to the Linux Foundation's Agentic AI Foundation December 2025, co-founded by Anthropic, OpenAI, and Block, with AWS, Google, Microsoft, and Cloudflare supporting; thousands of registry servers; tens of millions of monthly SDK downloads.

[16] W3C Web Machine Learning Community Group, "WebMCP," Draft Community Group Report, February 10, 2026, co-developed by the Google Chrome and Microsoft Edge teams; public origin trial in Chrome 149, June 2026.

[17] Rankability llms.txt adoption tracker, June 2026. 8.7 percent of the top 1,000 websites publish llms.txt, up more than fivefold in twelve months.

[18] Zou, Geng, Wang, Jia, "PoisonedRAG: Knowledge Corruption Attacks to Retrieval-Augmented Generation of Large Language Models," USENIX Security Symposium 2025. 90 percent attack success with five injected texts; evaluated defenses insufficient. Related retrieval-jamming work at the same conference, and 2026 follow-on research on poisoning against AI security agents.

[19] Jason Lemkin, SaaStr, "The AI SDR Reality Check" and subsequent deployment reports, 2025 to 2026. 90 percent of AI SDR deployments produce nothing across more than twenty companies observed; winners manage agents like $100,000 hires; SaaStr's own multi-agent deployment generated millions in pipeline across its first year under daily human management.

[20] Astra GTM, "The AI SDR Reality Check," May 2026. Roughly 0.3 percent reply rates for fully AI-generated cold email versus 4 percent or better for human-written copy, from analysis of more than 1 billion emails; roughly 2 percent stick rate for full-automation deployments after 90 days.

[21] TechCrunch, March 2025. 11x.ai claimed roughly $10 million in recurring revenue against roughly $3 million past the three-month break clause, and displayed customer logos without permission; the company disputed the retention characterization.

[22] Digital Applied, "AI SDR Agents in 2026: The Realistic Buyer's Guide." Hybrid configuration booked 312 meetings at 38 percent conversion versus AI-only at 847 meetings and 11 percent, roughly 2.3 times the revenue (directional); cold email reply rates sliding from 6.8 percent in 2023 to 4 to 5 percent in 2025 and 2026; roughly 25 sends per mailbox per day as a deliverability ceiling.

[23] Autobound, "State of AI Sales Prospecting 2026." Signal-personalized outreach at 15 to 25 percent reply rates versus a 3 to 5 percent average.

[24] Optifai, "2026 Pipeline Study" (939 B2B SaaS companies). Win rates of 28 to 35 percent under $10,000 in deal value and 12 to 18 percent above $100,000; median sales cycle of 84 days, 22 percent longer than 2022.

[25] Martal Group, "2026 B2B Data Industry Report," July 2026. Data and tooling spend no longer translating to pipeline gains.

[26] Boston Consulting Group, 2026 B2B sales AI research. The 70/20/10 rule: 70 percent of AI value from people and process, 20 percent from data and technology, 10 percent from algorithms.

[27] FloTorch and Firecrawl retrieval benchmarks, 2026. Plainly structured documents with clear titles, summaries, and short sections outperformed exotic chunking approaches.

[28] Prompt Payment Act, 31 U.S.C. 3901 et seq., implemented at 5 CFR Part 1315. 30-day payment requirement; automatic interest on late payment; additional penalty up to 100 percent of interest owed. Federal Register, July 2026: prompt payment interest rate of 4.75 percent per annum for July 1 through December 31, 2026.

[29] US Small Business Administration. $183.5 billion in federal prime contracts to small businesses in fiscal year 2024, a record.

About The Author

Tim Doelger is the founder of Get 'er Done, a revenue consultancy for B2B companies between $2 million and $50 million in revenue, based in New Jersey. He has spent his career inside B2B sales organizations, including building outbound systems at High Reliability Group for Fortune 500 industrial clients, and now works as a fractional revenue leader, coach, and AI governance advisor for owner-led companies. His practice runs on one standard: every framework must be provable in the client's own records, and everything installed belongs to the client.