

THE AI BUSINESS OPERATING SYSTEM

*The field manual for teams
that run on AI*



ROGER WONG WON

FOUNDER & CEO, CAPTURE THAT MEDIA

CAPTURE THAT MEDIA

The AI Business Operating System

The field manual for teams that run on AI

Roger Wong Won

Founder & CEO, Capture That Media

First edition, 2026. Written and published by Capture That Media, LLC,
12015 Radium Street, San Antonio, Texas 78216.

All content original. Copyright 2026 Capture That Media, LLC.
Share it freely with attribution. Do not resell it.

The prompts in this manual are model agnostic. They have been used against
current frontier assistants and are written to survive model updates.

Nothing here is legal, financial or tax advice.

AAF GOLD ADDY · BEST OF SAN ANTONIO 2023

www.capturethatmedia.com · roger@capturethatmedia.com

VISIBILITY, EVERYWHERE.

CONTENTS

What is inside

INTRODUCTION	
Ten Percent Of A Tool You Are Paying Full Price For	3
PART ONE	
The Five Moves Almost Nobody Makes	6
PART TWO	
Five Operating Systems By Role	10
PART THREE	
The Visibility Layer	15
PART FOUR	
Twenty Five Workflows That Move Real Numbers	19
PART FIVE	
Working With Your Actual Data	26
PART SIX	
Fifteen Automations You Can Build This Month	30
PART SEVEN	
Connectors Worth The Setup Time	34
PART EIGHT	
The Ninety Day Rollout	37
PART NINE	
The Prompt Library	40

How to use this

Every section runs Explain, then Show, then Do. Open your AI tool alongside this document and run each section as you read it. Reading it without touching a keyboard is the least valuable thing you could do with it.

INTRODUCTION



Ten Percent Of A Tool You Are Paying Full Price For

Most teams are not bad at AI. They are stuck at the first move.

This manual is not a survey of the field. It is an operating system: the specific habits, structures and prompts that separate a team getting a 4 percent lift from a team getting back a day a week. Everything here is built to be used the same afternoon you read it.



There is a gap between what AI can do and what your team actually gets out of it, and that gap is almost never about the model. It is about operating discipline.

I run a marketing agency, Capture That Media, that has built AI into client delivery since our founding week, and I spend the rest of my time on applied AI work with teams trying to do the same thing. In both, I watch the same thing happen: a company buys seats, sends around a link to a prompt library, and six months later the honest internal answer to whether it changed anything is no.

What went wrong is rarely the tool. It is that the team never moved past asking questions. They typed into a box, got a decent paragraph back, pasted it somewhere, and called that adoption. That is the ten percent. This manual is about the other ninety.

The Four Levels

Every person using AI sits at one of four levels. The level is not about how smart they are or how much they know about machine learning. It is about the **shape of the interaction**. Diagnose the level first, because the fix is different at each one.

LEVEL	WHAT IT LOOKS LIKE	CEILING
Prompter	One question, one answer, close the tab. Treats the model as a faster search engine.	Saves minutes. Never compounds.
Operator	Gives a precise instruction with format, length and audience. Gets usable output first try.	Saves hours. Still one task at a time.
Architect	Chains steps into a pipeline. Output of one becomes the input of the next. Reuses the chain.	Saves days. Now a system, not a session.
Owner	Builds context once, delegates whole workflows, reviews exceptions rather than every output.	Changes headcount math.

Most professionals plateau at Operator and assume they have arrived. They write good prompts. They get good answers. And they still do every step by hand. The jump from Operator to Architect is the single highest return move in this manual, and it is not a skill. It is a decision to stop treating each task as new.

FIND YOUR LEVEL IN NINETY SECONDS

1. Open your AI tool and look at your last ten conversations.
2. Count how many were a single question with a single answer. That is your Prompter rate.
3. Count how many produced something you shipped without rewriting it. That is your Operator rate.
4. Count how many you would ever open again and reuse. That number is almost always zero, and that is exactly the problem.



THE DIAGNOSTIC

If you cannot point to a single conversation you would reuse, you do not have an AI workflow. You have an AI habit. Habits do not scale across a team and they do not survive the person leaving.



How To Use This Manual

Every section follows the same three beats. **Explain** gives you the principle in a paragraph. **Show** gives you the thing itself, usually a prompt you can copy. **Do** gives you the specific action to take before you close the file. If you only ever read the Do boxes, you will still be ahead of most of your competitors.

Open your AI tool alongside this document. Read a section, run it immediately, then move on. Reading this cover to cover without touching a keyboard is the least valuable thing you could do with it.

The teams that win with AI are not the ones with the best prompts. They are the ones who stopped starting from scratch.



The Five Moves Almost Nobody Makes

Underused capabilities that change the economics of the tool.

These are not tricks. Each one removes a specific tax you are currently paying: re-explaining context, accepting a mediocre first answer, doing work the model could verify, or asking for prose when you needed a decision.



Move 1: Build Context Once, Not Every Time

The most expensive habit in AI use is re-explaining who you are at the start of every conversation. Your industry, your customer, your voice, your constraints, your definition of good. Typed again and again, badly, in a hurry, slightly differently each time.

Write it once. A context brief is a standing document, roughly 400 to 800 words, that you paste or attach at the start of any serious conversation. It should cover: what the business does in one paragraph, who the customer is and what they actually worry about, the three outcomes that matter this quarter, how you sound, and what you never say.

CONTEXT BRIEF GENERATOR

I want to build a reusable context brief I can paste at the start of future conversations so you understand my business without me re-explaining it. Interview me. Ask one question at a time, maximum twelve questions, covering: what we sell and to whom, the customer's real problem in their own words, our differentiator, our tone of voice, our three priorities this quarter, and the things we never say. After the last answer, output the finished brief as a single block under 700 words, written in the second person addressed to an AI assistant.

WHY IT WORKS

A good context brief typically improves first draft usability more than any prompt engineering technique. You are not making the model smarter. You are removing its need to guess, and guessing is where generic output comes from.

Move 2: Ask For The Draft You Did Not Ask For

The first answer is the model's safest answer. It is the average of everything reasonable. If you accept it, you ship average. The fix costs one sentence.

THE THIRD OPTION

Give me three distinct approaches to this, not three variations of one approach. Approach one should be the obvious, safe version. Approach two should be what a sharp competitor would do. Approach three should be the one most people would reject as too bold. For each, give me the strongest argument for it and the specific risk it carries. Then tell me which you would pick and why.

The value is not the third option. It is that seeing the range recalibrates your judgement on the first one. You will often ship a sharpened version of option one, but you will ship it knowing what you traded away.

Move 3: Make It Argue With Itself

Models are agreeable by default. Left alone they will polish your idea rather than test it. You have to explicitly commission the opposition, and you have to do it in a separate step, because asking for a draft and a critique at the same time gets you a soft critique.

ADVERSARIAL PASS

Forget that you wrote this. You are now a skeptical reviewer whose reputation depends on catching weak reasoning. Go through the material above and list: every claim that is asserted but not supported, every place the logic skips a step, every assumption that would break the argument if it turned out to be false, and anything a well informed reader would push back on immediately. Be specific and quote the line. Do not soften anything. Do not offer fixes yet.

Run the fixes as a third message, not the second. Separating find and fix keeps the critique honest, because a model that knows it has to repair the damage will find less of it.

Move 4: Ask For The Decision, Not The Essay

A great deal of AI output is prose nobody needed. You asked what to consider, so you got eight considerations, and you are no closer to a decision than before. Specify the artifact.

INSTEAD OF

ASK FOR

What should I consider when pricing this?	A recommended price, the two closest alternatives, and what has to be true for each to be right.
Analyze this feedback.	The three themes ranked by how often they appear, with one representative quote each and the decision each theme forces.
Help me think about this hire.	A go or no-go with the single strongest argument against your own answer.
Summarize this meeting.	Decisions made, owners, dates, and the open questions nobody answered.

THE RULE

Name the artifact and the shape before you name the topic. The difference between a useful assistant and an expensive autocomplete is almost entirely in this one habit.



Move 5: Give It Your Actual Material

Teams describe their data to the model instead of handing it over. They summarize the customer complaints rather than pasting them. They characterize the spreadsheet rather than uploading it. Every layer of summary you add is a layer of your own bias baked in before the analysis starts.

Give it the raw export. The full transcript. The whole thread. The messy CSV with the inconsistent column names. Modern models handle volume far better than they handle your compression of that volume, and the surprises, which are the entire point, live in the parts you would have edited out.

DO THIS NOW

1. Build your context brief with the generator above. Fifteen minutes.
2. Save it somewhere you can paste from in two seconds. Not a folder you will never open.
3. Take the next real task you were going to hand to AI and run it twice: once cold, once with the brief attached. Compare.
4. Keep the version that wins and delete the other. That comparison is the whole argument for everything in this manual.

Five Operating Systems By Role

Ready to deploy setups for Founder, Sales, Content, Operations and Client Delivery.

An operating system here means a standing context, a set of recurring jobs, and the prompts that run them. Pick the one closest to your seat, install it this week, and adapt the others as you grow into them.



The Founder OS

The founder's problem is not lack of information. It is that every decision arrives without the analysis attached, and the analysis is what takes the week. The Founder OS front loads structure onto decisions so you spend your judgement, not your time.

STANDING JOBS

- **Monday decision queue.** Dump every open decision, get them sorted by reversibility and cost of delay.
- **Pre-mortem on anything expensive.** Before committing, force a written account of how it failed.
- **Weekly numbers narrative.** Turn the dashboard into three sentences you could say out loud to an investor.
- **Inbox to commitments.** Extract everything you actually promised someone this week.

THE DECISION QUEUE

Here is every open decision on my plate this week. For each one, classify it as reversible or one-way, estimate the cost of delaying it by two weeks, and identify the single piece of information that would most reduce my uncertainty. Then order them: one-way and expensive to delay first, reversible and cheap last. For the bottom third, tell me which I should simply decide in the next sixty seconds without further analysis, and give me your recommendation for each so I can just confirm or overrule.

THE PRE-MORTEM

It is eighteen months from now and the decision described below has gone badly wrong. Write the internal post mortem. Be concrete about what specifically failed, in what order, and what the early warning signs were that we ignored at the time. Then list the three cheapest things I could do today that would have prevented the top failure mode.

WHY IT WORKS

The pre-mortem works because prospective hindsight beats forecasting. Asking what could go wrong produces vague risk lists. Asserting that it already went wrong and asking for the story produces specific, testable failure modes.

The Sales OS

Sales teams misuse AI by asking it to write emails. That is the least valuable thing it can do for you. The high value work is preparation and pattern reading: knowing what this specific buyer cares about before

the call, and knowing what your last forty calls have been trying to tell you.

PRE-CALL BRIEF

I have a call in an hour with the company described below. Using their site, their recent announcements and anything else you can find, give me: their likely top three priorities this quarter in their own language, the specific pain our offer maps to, two questions that will make them feel understood rather than sold to, the objection I am most likely to hear from someone in this role, and one thing I should absolutely not say. Keep the whole thing under 300 words. I am reading it in the car.

LOSS PATTERN ANALYSIS

Attached are the notes from every deal we lost in the last two quarters. Do not summarize them. Find the patterns: at which stage do we most often lose, what objection appears most frequently in the deals we lost but rarely in the ones we won, and is there a segment or company profile where we systematically underperform? Quote the specific notes that support each pattern. Then tell me the one change to our process that would address the largest cluster.

OBJECTION HANDLING THAT DOES NOT SOUND LIKE A SCRIPT

OBJECTION DRILL

You are the buyer described below and you have just raised this objection. I am going to respond. After each of my responses, stay in character, react the way that buyer realistically would, and escalate if my answer was weak. Do not break character to coach me until I type COACH. When I do, tell me exactly which of my responses lost ground and what a stronger version would have been.

The Content OS

The failure mode in content is volume without a spine. Teams generate forty posts that all say nothing in particular. The fix is to make the model work from a fixed point of view rather than from a topic.

A point of view is a claim you would defend in an argument. Not a subject area. Not a keyword cluster. The Content OS starts by extracting yours, then produces everything downstream from it.

POINT OF VIEW EXTRACTION

Below are transcripts, notes and past writing from me. I want you to find my actual point of view, not my topics. Specifically: what do I believe about this industry that most people in it would argue with? What am I consistently annoyed by? What advice do I give that contradicts the conventional advice? Give me five candidate points of view, each phrased as a single sentence I could defend on stage, ranked by how strongly the evidence in my own material supports it.

ONE IDEA, FIVE FORMATS

Take the point of view below and produce: a 900 word article that argues it properly with a concrete example, a LinkedIn post under 200 words that opens with the sharpest line and does not bury the claim, a 45 second video script written for speaking not reading, three short posts that each attack one narrow slice of the idea, and a subject line plus 150 word email that gets someone to read the article. Keep my voice consistent across all of them. Do not soften the claim to make it more agreeable in the shorter formats.

WHY IT WORKS

Repurposing fails when it becomes reformatting. The same paragraph chopped into a carousel is not five pieces of content, it is one piece wearing hats. Attacking different slices of one idea is what makes a body of work feel deep instead of repetitive.

The Operations OS

Operations is where AI pays back fastest and gets deployed slowest, because the work is unglamorous and nobody owns it. The three jobs worth automating first are documentation, exception triage, and status reporting.

SOP FROM A RECORDING

Attached is a screen recording transcript of someone doing a task. Turn it into a standard operating procedure that a new hire could follow without asking questions. Number every step. Where the person made a judgement call, say so explicitly and give the rule they seemed to be applying. At the end, list the three points in this process most likely to go wrong and what to do when they do. Flag anything they did that looked like a workaround for a broken system, because that is what we should fix rather than document.

EXCEPTION TRIAGE

Here is this week's queue of support tickets and internal escalations. Group them by root cause, not by symptom. For each group, tell me how many tickets it accounts for, roughly how much time it consumed, and whether the fix is a product change, a process change, or a documentation change. Put the group with the worst ratio of effort to ticket count at the top.

The Client Delivery OS

For any agency or services business, the margin is eaten by the same three things: scoping badly, reporting slowly, and rewriting the same deliverable structure for every client. All three are structural, and structure is what AI is best at.

SCOPE STRESS TEST

Here is a scope of work we are about to send. Find every place where the language is loose enough that a reasonable client could expect more than we intend to deliver. For each, quote the phrase, explain the gap, and give me tighter language. Then list what is missing entirely: the assumptions, dependencies and client obligations that always cause friction later and are not written down here.

MONTHLY CLIENT REPORT

Here is the raw data for this client's month plus my rough notes. Write the report as a narrative, not a data dump. Open with the single most important thing that happened. Explain what changed and why, including anything that went backwards, said plainly. Connect the activity to the outcome they actually care about, which is stated in the brief. Close with what we are doing next month and what we need from them. Under 600 words. No filler about being excited to partner with them.

DO THIS NOW

1. Pick the OS closest to your seat. Just one.
2. Take its two highest frequency prompts and run them on live work today, not a test case.
3. Save the versions that worked into a shared document your team can reach.
4. Put one of them on a recurring calendar block so it becomes a habit rather than an experiment.

The Visibility Layer

Being good at AI and being found by AI are two different problems.

Your buyers are asking assistants who they should hire before they ever reach a search results page. Most companies have no idea what those assistants say about them. This part is about the half of the AI shift that has nothing to do with your internal productivity.



Every other part of this manual makes your team faster. This one decides whether anyone finds you at all.

For twenty years the discovery question was where you ranked on a page of ten blue links. That page is no longer the first thing a buyer sees. They ask an assistant, they get one answer with two or three companies named in it, and they act on that. If you are not one of the named companies, you were never in the running, and unlike a search ranking, there is no page two to be on.

Three Questions, In Order

AI visibility breaks into three separate problems, and they have to be solved in sequence. Most companies jump to the third and wonder why nothing moves.

LAYER	THE QUESTION	WHAT FAILS
Findable	Can the assistant retrieve accurate facts about you at all?	Inconsistent business details across the web. No structured data. Key pages not crawlable. The model has nothing solid to work from.
Recommendable	When someone asks for a provider like you, are you named?	You exist but you are generic. Nothing in your public material makes you the obvious answer to a specific question.
Actionable	Can the buyer take the next step without leaving the answer?	No clear pricing, no obvious contact path, no proof. The assistant names you and the buyer stalls.

LAYER ONE: FINDABLE

Assistants assemble answers from what they can retrieve and verify. Ambiguity is fatal. If your company name, address, service list and category are stated three different ways across your site, your listings and your profiles, the model resolves the conflict by trusting none of it and naming someone else.

- **One canonical description.** The same sentence about what you do, in the same words, everywhere.
- **Structured data that matches your copy.** Organization and service markup that says the same thing your page says, not aspirational language.
- **Entity consistency.** Name, address, phone, founding date and category identical across every listing you control.
- **Crawlable answers.** Your best explanations should be in text on public pages, not locked in a PDF, a video, or a gated form.

LAYER TWO: RECOMMENDABLE

Findable gets you into the index. Recommendable gets you into the answer. The mechanism is specificity. Assistants name the company that most cleanly matches the shape of the question, and most company websites are written to match no question in particular.

Write the answer to the exact question your buyer asks, using their words, on a page that does nothing else. Not a services page listing nine things. A page that answers one question completely and better than anyone.

QUESTION SURFACE MAPPING

You are a buyer with the problem described below. You have decided to ask an AI assistant for help finding a provider. Write out the twenty questions you would actually type, in the order you would think of them, from the earliest vague version to the final comparison question right before you contact someone. Use realistic phrasing including the imprecise early ones. Then mark which of the twenty a provider could plausibly own with one excellent page.

LAYER THREE: ACTIONABLE

The last layer is the one nobody checks. An assistant names you, the buyer asks a follow up, and the assistant cannot say what you charge, who you serve, or how to start. The recommendation dies there. Publish the boring specifics: pricing or at least ranges, service area, who you are not a fit for, and a single obvious next step.

Run The Audit On Yourself

This takes twenty minutes and is the most uncomfortable exercise in this manual. Do it before you spend anything on visibility.

1. Open three different assistants in fresh sessions with no history.
2. Ask each the five questions your best customer would have asked before they found you. Do not name your company.
3. Record every company named, in order, across all three.
4. Now ask each one directly what it knows about your company. Record what is wrong, what is missing, and what is out of date.
5. Ask what it would need to see to recommend you for those five questions. The answers are unusually specific and unusually actionable.

WHAT YOU WILL FIND

In most audits the same three things surface: a competitor you do not consider a real competitor is being named consistently, your own description is two years out of date, and at least one factual detail is simply wrong. All three are fixable in a week and none of them are visible from inside your own website.

This is not a one time exercise. The answers change as models update and as the web around you changes. Re-run it monthly and track which companies get named. That list, over time, is the most honest competitive intelligence available to you.

Twenty Five Workflows That Move Real Numbers

*Copy, adapt, run. Grouped by function, ordered by
payback.*

Each workflow states the job, the prompt, and the realistic time it saves.
Nothing here requires a developer, a subscription you do not have, or a
change to how your company is structured.



Strategy And Decisions

1. THE COMPETITIVE TEARDOWN

PROMPT

Here are three competitors' websites and public material. Do not summarize them. Tell me: what each one is claiming as their differentiator, where those claims overlap so much that they cancel each other out, and the specific position none of them are occupying. Then tell me whether that empty position is empty because it is valuable and hard, or because it is worthless.

2. THE PRICING SANITY CHECK

PROMPT

Here is our pricing and what is included at each tier. Argue that we are underpriced. Then argue that we are overpriced. Use only evidence from the material provided plus what you know about this market. Then tell me which argument is stronger and what single piece of data would settle it.

3. THE STRATEGY COMPRESSION

PROMPT

Compress the strategy document below into: one sentence a customer would understand, one sentence an investor would fund, and one sentence an employee could use to make a decision on Tuesday without asking anyone. If the document cannot support all three, tell me which one it fails and what is missing.

4. THE REVERSIBILITY FILTER

PROMPT

For each decision below, tell me the cost of being wrong and the cost of being slow. Mark which ones I should decide immediately with incomplete information because reversing them later is cheap, and which ones justify two more weeks of work.

5. THE ASSUMPTION AUDIT

PROMPT

List every assumption this plan depends on, including the ones so obvious nobody wrote them down. Rank them by how badly the plan breaks if the assumption is false. For the top three, tell me the cheapest test that would tell us within two weeks.

Sales And Revenue

6. TRANSCRIPT TO NEXT STEPS

PROMPT

Here is the transcript of a sales call. Extract: what the buyer said they needed in their exact words, what they said that suggests a concern they did not state directly, who else is involved in the decision, what I committed to, and the strongest reason this deal does not close. Then draft the follow up email in under 150 words using their language, not mine.

7. THE PROPOSAL DIFFERENTIATOR

PROMPT

Here is my proposal. A competitor is submitting one for the same job. Write the section of their proposal that would beat mine. Then tell me what to change in mine, ranked by impact.

8. SILENT PIPELINE DIAGNOSIS

PROMPT

These deals have gone quiet. For each, based on the notes, tell me the most likely actual reason for the silence, and write a re-engagement message that would work for that specific reason. No generic checking in. Each message should give them a reason to reply that serves them, not me.

9. ICP REFINEMENT FROM REAL DATA

PROMPT

Here are our thirty best customers and our fifteen worst, with revenue, tenure and support load. Find what actually separates them. Not what we say our ideal customer is, what the data says. Then write the disqualification criteria we should be using on discovery calls.

10. THE REFERENCE STORY**PROMPT**

Turn this customer interview transcript into a case study with a real narrative: what their situation actually was including what was going badly, what they tried first that did not work, what changed, and the specific numbers. Keep their voice in the quotes. No superlatives that they did not say themselves.

Content And Marketing

11. THE HOOK REWRITE**PROMPT**

Here is my opening paragraph. Write eight alternative first lines: two that open with a concrete number, two that open with a claim most readers would dispute, two that open mid-scene with a specific moment, and two that open with a question I actually answer. Then tell me which one earns the second line best and why.

12. VOICE MATCHING**PROMPT**

Here are five things I have written. Describe my voice as a specification another writer could follow: sentence rhythm, how I open, how I handle transitions, what I never do, the words I overuse. Then rewrite the attached draft to match that specification exactly.

13. THE COMMENT MINE

PROMPT

Here are 200 comments and replies from our audience. Find the questions being asked repeatedly, the misconceptions that keep surfacing, and the language they use for problems we describe differently. Rank by frequency. Each item should map to a piece of content we have not made yet.

14. EDITORIAL CALENDAR FROM A POINT OF VIEW**PROMPT**

Here is our core point of view. Build eight weeks of content where every piece attacks a different facet of that one idea rather than covering different topics. For each: the angle, the format, the specific claim it makes, and how it connects back to the central argument.

15. THE DE-FLUFFING PASS**PROMPT**

Cut this by 40 percent without losing a single point. Remove throat clearing, restatement, hedging and anything that could be deleted without the reader noticing. Show me the cut version, then list the three sentences you removed that I might want to argue about.

Operations And Admin

16. MEETING TO COMMITMENTS**PROMPT**

From this transcript, extract only: decisions actually made, who owns each, the date committed, and the questions that were raised and never resolved. If a decision was implied but never explicitly agreed, list it separately as needs confirmation.

17. THE POLICY PLAIN-LANGUAGE PASS

PROMPT

Rewrite this policy so a new employee understands it on one read. Keep every obligation intact. Where the original is ambiguous, do not smooth it over, flag it as unclear and tell me what it could be read to mean.

18. VENDOR COMPARISON**PROMPT**

Compare these vendor proposals on total cost over three years including the costs they have not itemized, what happens if we need to leave, and what each one is deliberately vague about. Then tell me which questions to ask before signing.

19. JOB DESCRIPTION FROM REALITY**PROMPT**

Here is what this role actually does day to day, from the person currently doing it. Write the job description honestly, including the parts that are tedious. Then write five interview questions that would reveal whether someone can do the hard part.

20. THE HANDOVER DOCUMENT**PROMPT**

Someone is leaving in two weeks. Based on the material below, build their handover document: what they own, what is in flight with current status, who the relationships are with, and the things they know that are not written down anywhere. Flag the gaps you cannot fill so I know what to ask them before they go.



Analysis And Reporting

21. THE DASHBOARD NARRATIVE

PROMPT

Here is this month's data. Write three sentences: what changed, why it most likely changed, and what we should do about it. Then tell me the one number in here that looks fine but should worry me.

22. COHORT BEHAVIOUR**PROMPT**

Break this customer data into cohorts by signup month. Tell me whether retention is improving or degrading over time, and whether any cohort behaves differently enough to be worth investigating. Show me the numbers you used.

23. SURVEY SYNTHESIS**PROMPT**

Here are 400 open ended survey responses. Cluster them by underlying theme, not keyword. Give me the size of each cluster, two verbatim quotes that best represent it, and the action each one implies. Flag any cluster where respondents contradict each other, because that is a segmentation problem.

24. ANOMALY HUNT**PROMPT**

Look at this data and tell me what is weird. Not what is expected, what is anomalous: outliers, sudden changes, relationships that should exist and do not, and anything that suggests a data quality problem rather than a real effect.

25. THE FORECAST CHALLENGE**PROMPT**

Here is our forecast and the assumptions behind it. Build the downside case using the same data: which assumptions are most likely optimistic, what the number looks like if the three weakest ones are wrong, and what the leading indicator would be that we are heading there.

Working With Your Actual Data

PDFs, exports, transcripts and the analysis you keep not doing.

The single biggest unclaimed value in most companies is the data they already have and never look at properly. Not because it is hard, but because it takes an afternoon nobody has.



The Rule Of Raw

Give the model the export, not your description of the export. Every summary you write before the analysis is a decision about what matters, made before you knew what mattered.

Practically this means: upload the full CSV including the columns you think are irrelevant. Paste the whole transcript including the small talk. Attach the whole PDF including the appendix. The interesting finding is almost never in the part you would have kept.

Ask For The Method Before The Answer

The failure mode with data analysis is a confident number you cannot check. Force the method into the open first, then let it run.

TWO STEP ANALYSIS

Before you analyze anything, tell me how you plan to approach this: what you will calculate, what you will exclude and why, what assumptions you are making about the data, and what this analysis will not be able to tell me. Wait for me to approve the approach before running it.

WHY IT WORKS

This one habit catches most analysis errors before they become slides. It also surfaces the data quality problems that would otherwise show up as a confidently wrong conclusion three weeks later.

Documents

CONTRACT REVIEW

Review the attached agreement from my side. Flag: anything that creates an obligation for me that is not matched by one for them, every deadline or notice period, how it terminates and what survives termination, and any clause a lawyer would tell me to negotiate. Quote the exact language for each. I am not asking for legal advice, I am asking what to bring to my lawyer so the call is short.

LONG DOCUMENT INTERROGATION

I need to understand this document but I do not have time to read it. Give me: the three things it wants me to believe, the evidence it offers for each, the strongest counter-argument it fails to address, and the two pages I should actually read myself. Then wait, I am going to ask follow up questions.

Spreadsheets And Exports

MESSY DATA FIRST PASS

Here is a raw export. Before analyzing anything, tell me what is wrong with it: missing values and where they cluster, duplicates, inconsistent formatting, columns that appear to mean the same thing, and any values that are impossible rather than just surprising. Then tell me what I can and cannot reliably conclude from this data in its current state.

THE QUESTION BEHIND THE QUESTION

I asked for this metric, but tell me first whether it is the right metric for what I am actually trying to understand, which is stated below. If a different cut of this data would answer my real question better, say so and run that instead.

Transcripts

Call transcripts are the most underused asset in most businesses. Every objection, every misunderstanding, every phrase your customer uses instead of yours is sitting in them. Analyze them in batches, not one at a time, because the value is in the pattern.

BATCH TRANSCRIPT ANALYSIS

Here are thirty call transcripts. Do not summarize each one. Across all of them find: the words customers use for our product category that differ from ours, the question that comes up most often and how well we answer it, the moment in the conversation where deals most often turn, and anything customers assume about us that is not true. Quote specific calls as evidence for each finding.

DO THIS NOW

1. Find one export you have never properly analyzed. Support tickets, churned accounts, call notes, anything.
2. Run the messy data first pass on it before anything else.
3. Ask it one real question you have wondered about but never had time to answer.
4. Write down what surprised you. That single line is usually worth more than the report you were going to build.

Fifteen Automations You Can Build This Month

Real builds with triggers, steps and the failure mode nobody warns you about.

Automation is where AI stops being a tool you visit and starts being infrastructure that runs without you. Start with the ones that fire often and fail safely.



Before You Automate Anything

Three rules save more money than any tool choice.

- 1. Do it manually ten times first.** If you cannot describe the rule you are applying, you cannot automate it, you can only automate your confusion.
- 2. Automate the frequent and boring before the impressive.** A thing that runs forty times a week and saves four minutes beats a clever thing that runs monthly.
- 3. Decide what happens when it is wrong.** Every automation needs a human review point or a defined blast radius. Output that goes straight to a customer needs a person in the loop. Output that goes to a draft folder does not.

Tier One: Fire And Forget

These produce drafts or internal artifacts. Low risk, high frequency, build these first.

AUTOMATION	TRIGGER	WHAT IT DOES	WATCH FOR
Meeting to tasks	Transcript lands in storage	Extracts decisions, owners and dates, creates tasks in your tracker	Implied commitments read as firm ones. Route to a review list, not straight to assignment.
Inbound lead brief	New form submission	Researches the company, drafts a pre-call brief, attaches it to the CRM record	Confidently wrong company matches on common names.
Content atomizer	New article published	Produces the social variants and email version as drafts	Voice drift over time. Refresh the voice spec quarterly.
Review responder	New public review	Drafts a reply matched to sentiment and specifics, queues for approval	Never auto-publish. Ever.
Weekly numbers narrative	Schedule, Monday morning	Pulls the dashboard, writes the three sentence narrative, posts to the team channel	Explaining noise as signal. Ask it to say when a change is not meaningful.

Tier Two: Human In The Loop

AUTOMATION	TRIGGER	WHAT IT DOES	WATCH FOR
Ticket triage	New support ticket	Classifies by root cause, sets priority, drafts a first response for an agent to send	Miscategorising the rare urgent one. Keep an escalation keyword bypass.
Proposal assembly	Deal reaches a stage	Builds the draft from the discovery notes and your library of standard sections	Scope language loosening. Run the scope stress test before sending.
Churn signal digest	Weekly	Scans usage and support history, flags accounts whose behaviour changed	Alert fatigue. Cap it at five accounts a week or it gets ignored.
Invoice and expense sort	New document received	Extracts fields, categorises, flags anything outside normal ranges	Silent failures on unusual formats. Log what it could not read.
Candidate screen	New application	Summarises against the actual requirements, flags for human review	Bias amplification. Never let it reject. Only rank for a human to read.

Tier Three: Compound Builds

AUTOMATION	WHAT IT DOES	WHY IT IS WORTH THE EFFORT
Competitive monitor	Watches competitor sites, pricing pages and announcements, reports only material changes weekly	Replaces a task nobody does consistently and everybody says they should.
AI visibility tracker	Asks a fixed set of buyer questions across assistants on a schedule, logs who gets named	The only way to see the discovery layer described in Part Three moving over time.
Client reporting pipeline	Pulls each client's data, applies their brief, drafts the narrative report for review	In an agency this is often a full day a month per account manager.
Knowledge base from tickets	Turns resolved tickets into documentation, identifies undocumented recurring issues	Documentation that maintains itself is the only kind that stays current.
Onboarding sequencer	Generates the client or employee onboarding plan from the specifics of their situation	Removes the single biggest source of first-month churn, which is a generic start.



THE HONEST WARNING

Roughly a third of automations you build will be abandoned within two months, usually because the underlying process changed and nobody updated the automation. Build them so they fail loudly rather than quietly, and review the whole set once a quarter. An automation producing wrong output silently is worse than no automation at all.

Connectors Worth The Setup Time

Which integrations earn their keep, and which are a demo that impresses nobody twice.

Connecting your AI to your systems is the difference between an assistant that talks about your business and one that works in it. But most connector projects stall because the wrong ones get connected first.



The Selection Test

A connector is worth it when it satisfies all three: the data is something you reference often, it changes frequently enough that a static export goes stale, and pulling it manually is annoying enough that you skip it. Fail any one and you are building a demo.

CONNECT	BECAUSE	FIRST JOB TO RUN
Calendar	It changes hourly and drives everything else	Auto-brief before every external meeting
Email	It is where commitments hide	Weekly extraction of what you promised and to whom
Document storage	It is your institutional memory, unindexed	Answer questions from your own policies and past work
CRM	Deal context is the highest value context you own	Pipeline hygiene and pre-call briefs at scale
Team chat	Decisions get made there and recorded nowhere	Capture decisions from channels into a searchable log
Ticketing	The richest source of product truth in the company	Root cause clustering, weekly
Database	Removes the analyst bottleneck for simple questions	Read-only queries for the questions people ask constantly

Connect In This Order

1. **Read only, low stakes first.** Documents and calendar. Nothing can break.
2. **Then read only, high value.** CRM and ticketing. This is where the analysis payoff is.
3. **Then write access, narrow scope.** Let it create drafts and tasks before you let it send anything.
4. **Never grant write access to customer-facing channels without an approval step.** Not because the model is unreliable, but because the one time it is wrong is the time that costs you a client.

Governance That Does Not Kill Adoption

The two failure modes are equally bad: no policy at all, so people paste customer data into consumer tools, or a policy so restrictive that everyone quietly ignores it. What works is a short, specific document

that tells people what they can do rather than only what they cannot.

- Name the approved tools explicitly. Ambiguity pushes people to whatever is on their phone.
- Define three data tiers: freely usable, usable in approved tools only, and never. Give examples of each, not categories.
- State what must be human reviewed before it leaves the building. Keep the list short enough to remember.
- Require disclosure where it matters to the relationship, and say plainly where it does not.
- Put one named person in charge of answering questions. Most policy violations are people guessing.

The Ninety Day Rollout

How to deploy this across a team without a change management disaster.

Most AI rollouts fail the same way: a big launch, a training session, a prompt library nobody opens, and silence. This plan is built the opposite way, small and evidenced, because proof spreads and mandates do not.



Days 1 to 30: Prove It Somewhere

Do not roll out. Pick one team, ideally one with a visible, measurable, painful recurring task. Three to five people. Make it work there completely.

- **Pick the task, not the team.** The right first task is frequent, universally disliked, and has an output someone can see.
- **Measure before.** Time it. Count it. Whatever it is, get the number now, because after it works nobody will believe the before figure.
- **Build the context brief for that team.** Not the company. That team's actual work.
- **Meet weekly, thirty minutes.** What worked, what failed, what we are keeping. Write down the keepers.
- **Do not talk about it outside the group yet.** Premature evangelism kills more rollouts than skepticism does.

THE ONLY DAY 30 METRIC

Would this team be annoyed if you took it away? If yes, you have something. If no, more training will not fix it. You picked the wrong task, and picking a new one is cheap.

Days 31 to 60: Systematize

Now turn what worked into something transferable. The pilot team's tacit knowledge has to become explicit or it will not survive contact with the second team.

- **Write the playbook.** The five to ten prompts that actually get used, with the context they need and an example of good output.
- **Name the owner.** Someone whose job includes maintaining this. Not a committee.
- **Add the second team.** Choose an adjacent function. Have the pilot team teach them, not you. Peer proof travels further than executive endorsement.
- **Publish the number.** The before and after from the pilot, stated plainly, including anything that did not work.
- **Ship the governance doc.** Short. Now, while the stakes are still low.

Days 61 to 90: Make It The Default

- **Open access with a real onboarding.** Not a webinar. A thirty minute session where each person applies it to their own live work and leaves with one working prompt.

- **Build the first two automations.** From Tier One in Part Six. Something visible.
- **Run the AI visibility audit from Part Three.** It reframes AI from an internal efficiency story to a revenue one, which is what keeps executive attention.
- **Set the review cadence.** Quarterly: what are we still doing manually that we should not be, and what automations have quietly broken?
- **Decide what you stop doing.** The rollout is not finished until something was removed. If nobody's workload changed, nothing changed.

A rollout succeeds when people would fight to keep it, and fails when people had to be reminded to use it.



The Prompt Library

*The structure underneath every prompt in this manual,
plus the shortlist worth memorising.*

Once you understand why these work, you stop needing a library. You write the right prompt on the first try because you know what a prompt is made of.



The Five Part Structure

Every prompt in this manual is built from the same five components. Missing components are exactly where weak output comes from, and you can usually diagnose a bad result by spotting which one you left out.

COMPONENT	WHAT IT SETS	WHAT BREAKS WITHOUT IT
Role and stance	Whose judgement is being applied and how skeptical they are	Generic, agreeable output that reads like a summary of the internet
Context	The situation, constraints and what you already know	Advice for a company that is not yours
Task	The precise operation, one verb	It answers a nearby question instead of yours
Artifact	The shape, length and format of the output	An essay when you needed a decision
Standard	What good looks like and what to avoid	Technically correct output you have to rewrite anyway

THE SAME REQUEST, BOTH WAYS

WEAK

Help me write a follow up email to a client who has gone quiet.

BUILT

You are a senior account director who is direct and does not chase. Context: this client signed a six month retainer, engaged heavily for the first three weeks, and has not replied to two emails in eighteen days. Their stated goal was pipeline. We have delivered the first milestone. Task: write one email that re-opens the conversation. Artifact: under 120 words, no subject line pleasantries, one clear ask. Standard: it should give them a reason to reply that serves them. Do not say checking in, do not apologise for following up, and do not ask if now is a bad time.

Make It Show Its Work

For anything involving reasoning rather than recall, force the steps into the open. This is not a stylistic

preference. Reasoning that has to be written down is reasoning you can check, and reasoning you can check is reasoning that gets caught when it is wrong.

REASONING FIRST

Work through this step by step before you give me an answer. Lay out what you know, what you are inferring, and where the inference is weakest. Then give me the conclusion, and separately tell me what would have to be true for the conclusion to be wrong.

Twelve Prompts Worth Memorising

- **Make it worse first.** Write the terrible version of this so I can see what to avoid.
- **The missing question.** What should I be asking that I have not asked?
- **Steel man.** Give me the strongest possible version of the position I disagree with.
- **The specific objection.** What would the most informed skeptic in the room say to this?
- **Cut by half.** Remove 50 percent without losing a point, then tell me what you regret cutting.
- **In their words.** Rewrite this using only language my customer would actually use.
- **The tell.** What does this document accidentally reveal about the person who wrote it?
- **Second order.** If this works, what breaks downstream?
- **The cheapest test.** What is the fastest way to find out if I am wrong about this?
- **Name the trade.** What am I giving up by choosing this, stated plainly?
- **Before I read it.** What are the two things I should know before reading this document?
- **Confidence check.** Rate your confidence in each claim above and tell me which one you are least sure of.

One Morning, Start To Finish

If you do nothing else with this manual, do this tomorrow. It takes three hours and it will tell you more than another month of reading about AI.

TIME	DO THIS	YOU WILL END UP WITH
9:00	Build your context brief using the generator in Part One	A reusable asset that improves every conversation after it
9:30	Run the same real task twice, with and without the brief	Evidence, not opinion, about whether this matters

10:00	Run the AI visibility audit from Part Three on your own company	An uncomfortable and extremely actionable list
10:45	Take your least favourite recurring task and build it into a five part prompt	One workflow you will use every week
11:30	Pick one Tier One automation from Part Six and scope it	A build that pays for itself inside a month
12:00	Write down the one thing that surprised you	The reason you will actually keep going

You do not need to master AI. You need to stop starting from scratch, and everything after that compounds on its own.



ABOUT THE AUTHOR

Roger Wong Won

Roger Wong Won is the founder and CEO of Capture That Media, San Antonio's AI Visibility specialist and an award winning digital marketing agency serving brands across Texas and around the world.

A U.S. Army veteran, he spent two decades producing brand campaigns, live events and music across Texas, South Korea and Indonesia before bringing that work home. He shot and produced Red Bull's BC One All Stars campaign in South Korea, campaigns for City Bank's IPO and Reebok's CrossFit Sentinel line, and the rebrand of Wonderful Indonesia's national tourism website, filmed on the ground in Bali.

In 2018 he won the AAF Gold Addy for creative work produced in Texas. In 2020 he founded Capture That Media on Radium Street in San Antonio and built AI into client delivery from the founding week. The agency was named Best of San Antonio in 2023. He runs creative direction personally, and every engagement starts with a founder led discovery workshop, which is where most of the material in this manual was pressure tested before it was written down.

Find out what AI says about you

Part Three asks you to audit your own AI visibility by hand. We also run it automatically. The free AI Visibility Audit scores three layers: whether AI can find you, whether it recommends you, and whether a buyer can act on the answer. It takes about a minute and you get the actual transcript of what the assistants said.

capturethatmedia.com/free-audit

Capture That Media, LLC · 12015 Radium Street, San Antonio, TX 78216
roger@capturethatmedia.com · www.capturethatmedia.com
Branding · Social · SEO & AI Visibility · Web · E-Commerce · Marketing

VISIBILITY, EVERYWHERE.
