

# Building Walls.

What ZoomInfo and SciLeads teach us about defending B2B intelligence businesses in the age of agentic AI – and how *GTMplus*, *GTMbench*, and the OmniTech portfolio should respond.

*Anything that looks like a database loses. Anything that looks like a network, a point of view, or earned judgment wins.*

## § 00 CONTENTS

# What's *inside* this memo.

§ 01	Executive summary	03
§ 02	The three-stage pattern – aggregation, workflow, agentic compression	04
§ 03	The 2026 moat taxonomy – what's dying, what's strengthening	06
§ 04	GTMplus – the four pillars, scored	07
§ 05	GTMbench – moat versus table stakes	10
§ 06	The portfolio flywheel – how OmniTech's brands compound	11
§ 07	Agents as distribution, not just competition	13
§ 08	Where to spend the next 90 days	15
§ 09	Risks and counter-moves	17
§ 10	Closing principle	19
Appx	Revised copy – Signals, Playbooks, Advisors	20

## § 01 EXECUTIVE SUMMARY

## The *one-page* version.

Two companies, two ends of the market, the same structural story. ZoomInfo, the \$1.2B public-market leader in horizontal B2B intelligence, spent a decade building contact and company data moats that once commanded a category premium. SciLeads, a Belfast-based vertical-intelligence business serving Agilent, PerkinElmer, Qiagen, and Bio-Rad, spent a decade building the equivalent moat in life sciences – NIH grants, publications, clinical trials, lab purchasing data. Both succeeded. Both are now facing the same compression.

In 2026, agent-native entrants can assemble equivalent per-account context in thirty seconds from the same public sources these incumbents spent years stitching together. Clay, Common Room, and Default are doing it to ZoomInfo. Salesmotion and BioIntelli are doing it to SciLeads. The aggregation layer that was the moat is becoming a commodity input to an agent, not a product in its own right. ZoomInfo trades at a fraction of its 2021 multiple because the market has figured this out.

This memo uses both as case studies – horizontal and vertical, public and private, \$1B and \$20M – to show the pattern is structural. Three conclusions for OmniTech Capital.

### **One – the dying moats are not the ones OmniTech is betting on.**

Data aggregation and workflow UI – the moats that powered ZoomInfo and SciLeads alike – are the moats agents compress most aggressively. OmniTech has wisely not bet the portfolio on either.

### **Two – the strengthening moats are already present, but some need to be made more defensive.**

First-party signals, network effects, trust, and taste are all getting *more* valuable as AI commoditises everything else. These live in the GTMplus community, the Inner Circle, the Summit, and the GTMBench operator network. Several surfaces – the Signals and Playbooks pillars especially – are currently framed in ways that invite agent competition unnecessarily.

### **Three – the portfolio itself is the biggest moat.**

Five operating units plus one ecosystem, compounding into a closed loop, is something no single competitor and no single agent can replicate. The opportunity is not to build a GTM ZoomInfo. It is to build the GTM operator network – and let the agents plug into it as distribution.

## § 02 THE THREE-STAGE PATTERN

# Every B2B intelligence category moves through *three stages*.

The category does not disappear. Its moats change.

## STAGE 01

**Aggregation moat**

2010 → 2020

## STAGE 02

**Workflow moat**

2018 → 2025

## STAGE 03

**Agentic compression**

2024 → now

## 01 Aggregation moat

First-generation players won by ingesting, cleaning, and entity-resolving fragmented public data. ZoomInfo's moat was never "contact data" in the abstract – that's available everywhere. It was the pipeline that scraped, cleaned, and verified direct-dial phone numbers, titles, and org charts across millions of companies. SciLeads did the same in life sciences, pulling grants, PubMed, ClinicalTrials.gov, and tradeshow delegate lists into one queryable layer *with cross-entity resolution* – this PI equals this publication equals this grant equals this conference talk. Expensive. Slow. Hard to copy. A good moat for a decade – at either scale.

## § 02 · cont. STAGES TWO AND THREE

## 02 Workflow moat

Once aggregation commoditises, incumbents layer on workflow: saved searches, alerts, CRM sync, email sequencing, dashboards. The *platform* becomes the workflow, not just the data. Switching costs come from rep habit and Salesforce integration rather than data uniqueness. This is the ZoomInfo model in its mature phase, and the SciLeads model today. It is also where Gong, Apollo, and most vertical SaaS incumbents sit in 2026.

## 03 Agentic compression

Agents can now ingest the same public sources on demand, reason over them, and produce the output a rep actually wants – a personalised account brief, a drafted email, a prioritised call list – without the user ever touching a search builder. The aggregation layer becomes a commodity input to the agent. The workflow layer gets collapsed into natural-language instruction. Clay, Common Room, and Default are attacking ZoomInfo this way. Salesmotion and BioIntelli are doing the same to SciLeads. The pattern scales: same compression, different end of the market.

The ten-year head start matters less than it did. Not zero – data quality, entity resolution, and trusted provenance still matter. But the compression is real and accelerating, and it is coming to every B2B intelligence category – horizontal and vertical, public-data-driven and proprietary – not just life sciences.

## § 03 THE 2026 MOAT TAXONOMY

# Not all moats are compressing. Some are strengthening.

The table below is the scorecard OmniTech should navigate by. Read it as an investment allocation rule.

MOAT TYPE	EXAMPLE	2026 TRAJECTORY	WHY
Data aggregation from public sources	ZoomInfo contact data · SciLeads NIH pipeline	DYING	Agents compress years into hours
Workflow / UI layer	Saved searches, alerts, dashboards	COMMODITISING	LLM-native interfaces win on speed
First-party signals	Private community data, Summit behaviour	STRENGTHENING	Agents cannot scrape what they cannot access
Network effects	Trust graph between senior operators	STRENGTHENING	Humans still hire humans they know
Trust & reputation	"The person who's been in that seat"	STRENGTHENING	Boards don't hire AI; they hire judgment
Brand & taste	Distinctive editorial point of view	STRENGTHENING	Agents hedge; brands take positions

## THE ALLOCATION RULE

Every calorie spent on the top two rows compounds *slower* than it used to. Every calorie spent on the bottom four compounds *faster* than it used to – because the AI tide is lifting human-centred moats relative to data-centred ones. Invest accordingly.

## What this looks like in practice

At a portfolio level, this means OmniTech should resist every instinct to "aggregate data" and every instinct to "build better workflow tooling." Both are rational responses to stage-one and stage-two competition – and both are wrong answers for stage three. The right answers: deepen the community, deepen the operator network, strengthen named authorship, and let the agents plug in.

## § 04 GTMPLUS – THE FOUR PILLARS, SCORED

# Signals. Systems. Playbooks. Big Ideas. *Not equally exposed.*

The scoring below translates the 2026 moat taxonomy into pillar-level vulnerability and prescribes a defensive move for each.

## Signals

◆ HIGHEST VULNERABILITY

If Signals reads as *"we aggregate the GTM signals you'd otherwise miss"* – funding, hiring, tool-stack changes, pipeline events – it is in a fair fight with Clay, Common Room, Default, and the next wave of agentic prospecting tools. They monitor the same public sources on demand.

**DEFENSIVE MOVE** First-party. Reposition Signals around what GTMplus can see that nobody else can – the Slack community, the VP+ Inner Circle, Summit attendee conversations, the questions senior operators are actually asking each week. An agent cannot scrape a private community. The pillar description should say so explicitly: *"Signals only the GTMplus network can see."* This changes the pillar from a database bet to a network bet.

## Systems

◆ MEDIUM VULNERABILITY

Systems content – how to build the RevOps stack, how to instrument PLG, how to wire Clay → Salesforce → Gong – is prescriptive, and agents handle prescription reasonably well. The moat is not the information; it is authorship.

**DEFENSIVE MOVE** Every Systems piece carries a named operator byline from the GTMBench network. *"How the VP of RevOps at a \$30M ARR fintech rebuilt lead routing"* is defensible. *"How to rebuild lead routing"* is not. This also creates an audience-to-placement flywheel: every Systems piece becomes a trailer for the operator behind it.

## § 04 · cont. PILLARS THREE AND FOUR

## Playbooks

◆ HIGH APPARENT · MEDIUM ACTUAL

LLMs generate generic playbooks trivially. The LLM coverage strip already on the pillar – ChatGPT, Claude, Grok, Gemini – silently acknowledges this. The defensive question is whether Playbooks look like generic AI output or like something an AI *cannot* produce.

**DEFENSIVE MOVE** Brand every Playbook as first-person, operator-authored, stage-specific, battle-tested. The *"My First 30 Days as an AI-Native CMO"* deck is the template: first-person field notes from a fractional CMO at a \$50M ARR fintech ISV, week by week, named. An agent will not stake its reputation on a specific 30-day sequence at a specific company stage. A named operator will. That is the moat.

## Big Ideas

◆ LOWEST VULNERABILITY · UNDER-WEIGHTED

Agents synthesise existing ideas; they do not originate category-defining frames. *"Go To Market is being rebuilt."* *"The AI GTM Engineer is the new critical hire."* These live on taste, timing, and willingness to stake a strong position. Irreducibly human.

**DEFENSIVE MOVE** Double down on named, opinionated, contrarian authorship. Bylined pieces. Strong calls. Positions that could be wrong. This is also the pillar that does the most for GTMplus as a *media* brand – it is what gets quoted, screenshotted, and argued about on LinkedIn.

## § 04 • Synthesis WHAT THE SCORING TELLS US

## The pattern is *consistent*.

Across all four pillars, the same principle applies. The pillars that *describe a database* – "we aggregate signals", "we publish playbooks" – are the vulnerable ones. The pillars that *describe a network or a point of view* – "signals from our community", "playbooks from named operators", "opinions from our editorial voice" – are the defensible ones.

The good news: the pillars are already built on the right substrate. GTMplus has a community. It has an Inner Circle. It has Summit London. It has operators in the GTMBench network. Every ingredient for a network-based moat is present. The work is not to build new infrastructure; it is to surface the existing infrastructure more clearly in the copy, the product surface, and the editorial cadence.

### THE FRAMING SHIFT

Today, GTMplus reads as "*a media brand with a community attached.*" It should read as "*a community whose insights we publish.*" The nouns reverse. The community becomes the product; the media is the distribution.

## What to change on the landing page

Three surfaces carry most of the framing risk: the four-pillar strip, the Who It's For grid, and the hero subheading. Revised copy for the Signals and Playbooks pillars is attached in Appendix A. The hero subheading can follow the same logic – lead with the network, not the newsletter. "*The signals, systems and playbooks from the senior GTM operators building the category*" is closer to the defensible framing than "*Signals, systems and playbooks for modern GTM,*" which still reads like a content product.

## What to leave alone

The visual identity is working – the dark theme, the electric green, the terminal-like kineticism – and it reinforces the *intelligence-layer* positioning rather than undermining it. The risk is copy-level, not design-level.

## § 05 GTMBENCH – MOAT VERSUS TABLE STAKES

# More *structurally* defensible. Three surfaces, three profiles.

GTMBench is structurally more defensible than GTMplus because the core product is human placement, not intelligence. But the revenue-engine-hub surface has three components with very different exposure profiles.

## Fractional Operators – not compressible

Placement is a trust-and-vetting business. Agents do not fix this. The moat is the network, the vetting rigour, and the reputation for matching well. Keep investing here. *This is working.*

## Fractional Industry Advisors – OmniTech's single strongest wedge

The premium advisor card – 20+ years' experience, ex-Fortune 500, Tier 1 strategy firm background, covering Retail, CG, Healthcare, Finance, Manufacturing – sells deep pattern recognition from actual deals in actual industries. That is the single thing an agent cannot do. An LLM has read the case studies; it has not sat in the boardroom when the CPG deal nearly collapsed.

### RECOMMENDATION – MAKE THIS SECTION LOUDER

It is currently framed as a premium tier *below* operators. The sharper frame is a clean dichotomy – **operators for execution, advisors for judgment** – with both at equal visual weight. The revised two-card layout is in Appendix A. "*Judgment you can't prompt*" is the kind of line that gets screenshotted and shared.

## AI-Native CMO Playbooks – same framing risk

If it reads "*here are some playbooks*," it's commoditised. If it reads "*first-person field notes from the fractional CMO we placed at a \$50M ARR fintech*," it's defensible. One-line fix.

## Table stakes – not moats

The platform UI, searchable directories, standard marketplace features – useful, but not defensible. Build them to parity. Scarce calories belong upstream in the network and advisor bench.

## § 06 THE PORTFOLIO FLYWHEEL

# A single brand's defensibility is bounded. *A portfolio's is not.*

This is the part most founders building B2B intelligence brands miss. The single-brand analysis of the previous sections is useful. The portfolio-level analysis is where the real moat lives.

OmniTech Capital is structured as **five operating units plus one ecosystem**:

**GTMplus.ai** – Audience, community, events (Slack, Inner Circle, GTMInstitute, GTM Summit London)

**Decision-Maker Network** – Buyer intelligence (BoardroomAI.io, VerticalAI.media, GTMBench Review)

**IndustryGeniuses** – Demand generation

**ENAI** – Execution engine

**GTMBench** – Operator placement platform

Individually, each is attackable. *Together*, they form a closed loop that no single competitor – and no agent – can replicate, because each brand generates an asset the others monetise.

## THE LOOP

- **GTMplus** brings in the audience (media, community, events)
- The community produces **first-party signals** – Inner Circle dialogue, Summit behaviour, question patterns, deal stories
- Those signals feed the **Decision-Maker Network** (Boardroom briefings, VerticalAI intelligence)
- The same community surfaces high-quality operators, who get placed via **GTMBench**
- Placed operators author **Playbooks** that flow back to GTMplus – content without content-marketing cost
- The **Summit** reinforces the trust graph live – signal capture plus relationship capital
- **GTMInstitute** certifies the operators GTMBench places, locking in quality
- **ENAI** executes for companies the Decision-Maker Network surfaces as in-market

Each brand creates a moat that strengthens the others. GTMplus without GTMBench is a newsletter. GTMBench without GTMplus is a marketplace with no audience. Together, they are something a single competitor cannot build.

## § 06 • Synthesis WHAT THIS IMPLIES FOR STRATEGY

# The portfolio is the *strategy*. The brands are instruments.

When founders run portfolios, two failure modes dominate. The first is *drift* – each brand wanders into its own market logic, the compounding stops, and the portfolio becomes five stand-alone businesses sharing a corporate parent. The second is *over-integration* – the brands get merged too tightly, the audiences resist the cross-sell, and the portfolio collapses into one product.

The OmniTech design threads the needle if – and only if – the flywheel stays in motion. That requires three operating disciplines:

## Discipline one – asset handoffs must be codified, not assumed

The flywheel only works if GTMBench placements actually produce Playbooks, if Summit attendees actually feed the Decision-Maker Network, if community questions actually surface operator opportunities. These handoffs do not happen automatically. They require process. *Every placement produces one Playbook within 90 days* should be a contract term, not an aspiration.

## Discipline two – audience reuse, not audience collision

The audiences overlap but are not identical. A CRO reading GTMplus is not necessarily the same person buying BoardroomAI briefings, even if they sit in the same office. Cross-pollinate by insight, not by email list. Treat each brand's audience as sovereign.

## Discipline three – resist scope creep at the brand level

The most expensive mistake for a portfolio like this is GTMplus drifting into operator placement, or GTMBench drifting into media, or VerticalAI drifting into execution. Each brand has a single job. When a brand tries to do two jobs, it does neither well, and the flywheel slows. Protect the job description.

### THE ONE-LINE REFORMULATION

Other GTM media brands publish content. Other placement platforms place operators. Other intelligence products aggregate signals. *OmniTech Capital is the only company where those three activities compound into a single network effect.* That is the positioning. Everything else is execution.

## § 07 AGENTS AS DISTRIBUTION, NOT JUST COMPETITION

# The reframe that *flips* the threat into the opportunity.

Agents are not only the competitive threat. They are becoming the new distribution layer. Getting this right is worth more than any single defensive move in this memo.

When a VP of Marketing asks her AI assistant "*what's a good 30-day plan for a new fractional CMO at a \$50M ARR fintech,*" the agent is going to cite something. Today, it cites a mixture of Medium posts, McKinsey summaries, and generic SaaS blogs. In twelve months, the citations will consolidate around sources the agents treat as authoritative. The question is which sources.

GTMplus Playbooks, and every OmniTech editorial surface, should be built to be the thing agents cite. That means five specific moves:

## One — named authors.

Agents cite humans more readily than anonymous content. Every Playbook carries the name, role, company, and credentials of the operator who wrote it. No ghostwriting, no masthead-only bylines.

## Two — specific context.

Stage, industry, and company shape. Agents need to match user context to content. A Playbook labelled "*Series-B fintech ISV, US-based, 40 reps*" retrieves better than a Playbook labelled "for SaaS companies."

## Three — structured formats with clear claims.

LLMs retrieve structured content better than prose. Use headings, enumerated steps, labelled sections, explicit claims. Make it easy for an agent to extract a quotable excerpt.

## § 07 · cont. MOVES FOUR AND FIVE

**Four – public, crawlable, archive-quality URLs.**

The agent needs to reach the content. Don't gate Playbooks behind email capture. Make them open web, indexable, permanent. The email capture moves to adjacent assets – newsletter signup, Community invite, event RSVP – not the Playbook itself.

**Five – distinctive framing other sources don't replicate.**

Agents surface unique voice. A GTMplus Playbook that reads like a McKinsey deck will be passed over in favour of the actual McKinsey deck. A Playbook that reads like no other source will be cited because it is the only source for that framing.

## THE INVERSION

The same agents threatening SciLeads and ZoomInfo are the agents that will increasingly send traffic – and authority – to named, specific, operator-authored content. GTMplus is well-positioned for this shift if the content is built for it. *The agent becomes distribution, not competition – but only for the brands that design for it.*

## § 08 WHERE TO SPEND THE NEXT 90 DAYS

# Scarce calories. *Specific bets.*

Two columns. What to invest in, what to stop. Both matter; the stop-list is where most portfolios bleed out.

---

## ◆ INVEST HERE

### **Community infrastructure.**

The Inner Circle and open Slack are the first-party signal generator. Budget for moderation, curation, retention, and structured capture of what members discuss.

---

### **Operator-to-Playbook pipeline.**

Make Playbook authorship a condition of GTMBench placement. One named, first-person, stage-specific Playbook per operator within 90 days.

---

### **Summit as signal layer.**

Capture who's there, what they're asking, who's meeting whom. Feed back into the Decision-Maker Network as proprietary buyer signal.

---

### **Industry Advisor bench – made louder.**

Restructure the revenue-engine-hub section as "Execution or judgment" with operators and advisors at equal visual weight.

---

### **Named bylines on every surface.**

No anonymous content. Every piece ships with a named author and their operator credentials.

---

### **Agent-readable Playbooks.**

Open web, structured formats, specific context tags. Built to be cited by the research agents that arrive in 2026–27.

## § 08 · cont. THE STOP-LIST AND THE BUDGET TEST

## ◆ STOP HERE

**Don't build a GTM database.**

The instinct to aggregate tool stacks, org charts, hiring, funding will be strong. Resist it. A GTM ZoomInfo is a stage-1 moat that will be compressed before you finish building it.

**No generic playbooks.**

No unnamed Playbooks. No "10 tips for..." content. Every piece = named operator writing specific, stage-matched field notes.

**Signals-as-feed framing.**

Stop presenting Signals as a data product. Frame them as network-sourced insight. The pillar copy must say so.

**In-house workflow tooling.**

Let agents plug into OmniTech. Don't build the workflow layer yourself – that's a stage-2 bet losing to stage-3 players.

**Scope creep between brands.**

GTMplus doesn't place operators. GTMBench doesn't publish editorial. VerticalAI doesn't execute. Each brand = one job.

**Gated Playbooks.**

Drop email-gates on the Playbooks themselves. Move the capture moment to adjacent assets – newsletter, Community, Summit.

## BUDGET TEST

Before every spending decision – feature, hire, content piece, platform – ask: *does this deepen the network, the operator bench, the trust graph, or the editorial voice?* If yes, fund it. If it deepens the database or the workflow UI instead, reconsider. The test is blunt on purpose.

## § 09 RISKS AND COUNTER-MOVES

# Three risks worth *naming*.

## Risk 01 – Community fragmentation

If GTMplus's Slack ceases to be the best place for senior GTM operators to talk to each other, the first-party signal source dries up. Private communities are brittle – they thrive until they don't. A quieter competitor with better moderation, a better cohort, or a better in-person cadence can hollow out a community in six months.

**COUNTER-MOVE** Treat the Inner Circle like the product it is. Premium programming, tight membership criteria, high-quality moderation, regular in-person moments. *Under-admit rather than over-admit.* The scarcity is the asset. Most communities die by growth, not by neglect.

## Risk 02 – An agent-native entrant builds "the GTM research agent"

A Chrome extension or Slack bot that answers any GTM research question in natural language, ingesting the same sources GTMplus and GTMBench use. The Salesmotion-for-GTM scenario. It happens in 2026 or 2027, funded by a Series A that prices in the aggregation layer's collapse.

**COUNTER-MOVE** Do not compete on research. Compete on the thing the agent needs – the network, the named operators, the trust graph. Publish Playbooks in structured, agent-readable formats (per §07). When the research agent gets built, OmniTech content is what it cites. *The agent becomes distribution, not competition.*

## Risk 03 – Playbook commoditisation

LLMs become good enough at generating operator-quality playbooks that named authorship stops being a differentiator. The "AI can't write this" defence erodes faster than expected.

**COUNTER-MOVE** Ship faster than agents can learn. By the time LLMs reliably produce a 30-day CMO plan at the quality of the GTMBench Playbook, ship the 90-day one. By the time they catch that, ship the stage-specific post-Series-B version. *The content moat is not static knowledge; it is live operator access.* Agents train on yesterday; operators live in today.

## Three more risks worth *monitoring*.

Not yet requiring counter-moves, but each deserves quarterly re-assessment.

### **01 – A horizontal intelligence incumbent enters GTM content**

ZoomInfo, Gong, or Apollo buys into the GTM content category and runs a scale play – using their data pipes and distribution to flood the same audience GTMplus is building. They do not need to win the category outright; they need to slow OmniTech's compound rate enough for the flywheel to stall.

### **02 – A major consultancy productises fractional GTM advisory**

Bain, McKinsey, Accenture, or a boutique packages the "fractional CMO / CRO" offering at price points that make GTMBench look expensive relative to the brand and the Rolodex. The advisor bench is the defence here – OmniTech's advantage is the operator network and the mid-market price point, not the institutional brand.

### **03 – A platform shift in how B2B buyers discover operators**

LinkedIn changes its algorithm, a new professional network emerges, or agent-first discovery surfaces (ChatGPT-style recommenders) replace search. The operator-discovery layer is currently LinkedIn-shaped. If that shape changes, the acquisition funnel for GTMBench needs to change with it – fast.

## § 10 CLOSING PRINCIPLE

*In 2026,  
anything that looks like a  
database loses.  
Anything that looks like a  
network,  
a point of view,  
or earned judgment wins.*

ONE SENTENCE WORTH TATTOOING ON THE WALL

---

Every product decision, every positioning decision, every investment decision at OmniTech Capital should be tested against that principle. When there is doubt, pick the thing that looks less like a database and more like a network. Pick the thing that cannot be assembled by an agent in thirty seconds. Pick the thing a board would pay a human to do.

The ZoomInfo of 2012 and the SciLeads of 2016 both built the right businesses for their years. *The GTMplus of 2026 has the chance to build the right business for this year* – if it refuses to fight yesterday's war.

## Appx · AREVISED COPY BLOCKS

# Ready-to-ship copy for the *most exposed* surfaces.

These blocks slot directly into the existing GTMplus landing page and the GTMBench revenue-engine-hub without layout changes. Only the Industry Advisors section needs a structural restructure, shown below as a two-card layout.

## GTMPUS · PILLAR 01

## Signals

HEADLINE First-party. Network-only.

BODY The shifts operators are seeing inside their own pipelines – surfaced in the GTMplus community and VP+ Inner Circle before they hit LinkedIn, a press release, or anyone's agent. What you can't scrape, you can't automate around. Start here.

ALT HEADLINE *What agents can't see.*

## GTMPUS · PILLAR 03

## Playbooks

HEADLINE Written by operators. Not by AI.

BODY First-person field notes from the fractional CMOs, CROs, and RevOps leaders who actually shipped the motion – week by week, stage by stage, named. ChatGPT will write you a generic playbook in thirty seconds. It won't tell you what the VP of RevOps did when the board meeting went sideways.

ALT BODY *First-person field notes from the fractional operators we've placed – week by week, stage by stage, named. The playbook and the person who ran it. Hire them if you want the next chapter.*

Appx · AGTMBENCH · REVENUE-ENGINE-HUB

# Execution or judgment. *Often both.*

Restructure the current advisor section from a stacked hierarchy (operators primary, advisors premium add-on) into a clean dichotomy – both at equal weight, each with a distinct job. Section headline and subhead below; two-card layout follows.

SECTION HEADLINE

## Execution or judgment. Often both.

SUBHEAD

Two ways to work with *GTMBench* – pick the one that matches the problem.

### ◆ FOR EXECUTION

#### When you need someone to run the motion.

Senior GTM operators placed in days, shipping in weeks. CMOs, CROs, VPs of RevOps, GTM Engineers – embedded into your team, running the playbook, moving the number.

### ◆ FOR JUDGMENT

#### Judgment you can't prompt.

20+ years in Retail, CG, Healthcare, Finance, Manufacturing. Ex-Fortune 500. Tier 1 strategy firms. For the decisions an LLM can't make – when the board's wobbling, when the category's shifting, when the CEO needs someone who's sat in that room before.

## AI-Native CMO Playbooks – card copy fix

Small change, big lift. Replace the current card copy with language that signals first-person operator authorship rather than generic content.

REVENUE-ENGINE-HUB · PLAYBOOKS CARD

EYEBROW Field notes from the seat

TITLE My First 30 Days as an AI-Native CMO

BODY Week-by-week field notes from the fractional CMO we placed at a \$50M ARR fintech ISV. Not a framework. Not a "best practice" deck. What actually happened, what actually worked, and what the board pushed back on.

– End of memo.