

How I Used AI to Build a Real Sales and Marketing Plan: A Case Study



A fractional CIO's detailed account of using AI to create an actionable sales and marketing plan for a B2B manufacturing analytics company—showing the actual process, failures, and breakthroughs that transformed generic AI output into a professional deliverable.

Introduction

For over a year, I have read countless articles about AI. Every single one follows the same pattern. They tell you AI is transformative. They tell you it will change your business. They tell you everyone is using it successfully.

Not once have I seen an actual example showing how someone used AI to solve a real business problem. I have never seen anyone demonstrate the process from start to finish. I have only seen people claiming success without showing the work.

This case study shows you exactly that work.

I am a fractional CIO. I have been a CIO for 16 years. One of my clients needed a comprehensive sales and marketing plan. They are a small B2B company that sells data analytics products in the manufacturing sector. For this case study, I will call them Widgets Inc.

The company was facing serious challenges. Their core product line had been commoditized by competitors. They needed to pivot their go-to-market strategy immediately.

The Challenge



I decided to see if AI could produce an actual executable plan that could be handed to a CEO and implemented. The first attempts produced exactly what they always produce when you try to use AI for serious business problems. Useless output that sounds intelligent but provides no path forward.

But I kept pushing. I tried different approaches. I forced the AI into different roles. I made it critique its own work. I eventually used multiple AI systems to hold each other accountable.

By the end of this process, I had a professional, well-researched, actionable sales and marketing plan. It included quarterly execution plans, specific channel strategies, budget breakdowns, email scripts, technical readiness assessments, and risk mitigation protocols.

This case study shows you the actual prompts I used. It shows you the failures I encountered. It shows you the three critical turning points that transformed useless output into something genuinely valuable.

I am writing this to help you understand the early failures, the processes I went through, and how I managed to get AI systems to essentially hold each other accountable. The result was something I believe is implementable in a real business.

Chapter 1: The First Attempts Produced Nothing Useful

I started with what seemed like a reasonable approach.

I provided an AI system with a strategic roadmap document for Widgets Inc. This document outlined their business challenges, market position, and the direction they needed to move. I asked the AI to produce a tactical sales and marketing plan.

My First Prompt:

"Assume two roles: the senior sales leader at Widgets Inc. and the senior marketing leader at Widgets Inc. In your virtual boardroom, argue, banter, discuss, and ultimately agree on 3 items for sales and 3 items for marketing that will give Widgets Inc. the best chance of acquiring the highest possible revenue collection of customers. Describe each tactical solution, the artifacts, and the metrics."

This seemed like a solid request. I was asking for specific deliverables, clear metrics, and a structured approach. I thought giving the AI two perspectives would force more rigorous thinking.

What the AI Produced

The AI created a fictional boardroom debate between two characters it called Marcus from Sales and Sarah from Marketing. The dialogue looked sophisticated on the surface.

Marcus argued that the company needed to focus on larger enterprise deals instead of small customers. Sarah countered that small customers close faster and provide immediate revenue. They went back and forth for several paragraphs.

Eventually, the AI produced six solutions. Three for sales, three for marketing. Here is an example:

AI Output - Sales Plan Sample:

"The Reputation Rescue Approach

Strategy: Identify widget users with operational inefficiencies and position our analytics as the solution.

Artifacts:

- A 60-second audit video
- A follow-up call script

Metric: Meeting conversion rate greater than 15%"

I looked at this output and felt exactly what I have felt many times before when trying to use AI for real work...

There was no channel strategy. No sequencing. No actual instructions a sales team could follow. No definition of messaging. No timelines. No budget. No owners. Nothing connecting these ideas into an executable plan.

The Useless Output Problem



The output was not wrong. It was simply completely useless.

I have run into this time and time again when trying to use AI for serious business problems. You get output that sounds intelligent but provides no actual path forward.

But I decided to give the AI proponents a fighting chance. They seem to think an AI approach actually works. So instead of dismissing it and moving on, I decided to keep pushing. This was an important task for my client. I had time. This would be a learning exercise regardless of whether it succeeded or failed. I could write about either outcome. So I kept at it to see where it would go.

I went through several more iterations with the sales leader and marketing leader in their virtual boardroom. Each time I pushed for more detail. Each time the output improved slightly. But it never really got good enough.

Chapter 2: Forcing Accountability Changed Everything

After looking at multiple rounds of inadequate output from the sales and marketing leaders, I got frustrated.

The fictional debate between two leaders was not producing an executable plan. In this particular case, there was simply a lack of something...focus, context, something. The debate approach was simply not working.

Out of frustration with the lack of progress, I decided to try something different. Instead of asking the AI to produce content as two leaders debating, I positioned myself as the person who had to approve the work.

My Next Prompt:

"I am the CEO reading these deliverables. I understand targeting certain widget users as a strategy. What I am unclear on are the tactics. What are you going to do with these tools you mention? What do they look like? What will you do with them? I need the specific artifacts, the specific actions, the specific tactical plan, and the specific metrics for what our sales team will do every day."

This created an immediate shift in quality.

The AI, still operating as the sales and marketing leaders, suddenly started producing much more detailed responses. Adding accountability evidently made it realize the previous output would not survive scrutiny.

What Came Back

The revised plan suddenly included:

Quarterly timelines with specific months for each action

Target prospect segments with actual criteria

Channel prioritization with reasoning

Budget considerations broken down by activity

Concrete next steps with assigned owners

This was dramatically better. The content now read with much more tactical depth. The plan had actual sequencing and resource allocation.

I am not a sales expert. I am not a marketing expert. I am a CIO. I understand business and strategy, but I do not run sales and marketing teams daily.

There is a quote from an AI researcher that captures an important truth about AI. **AI seems incredibly smart until you ask it about something you actually know about.** When AI produces content outside your area of expertise, it can sound perfectly reasonable even when it has significant problems.

I was looking at a sales and marketing plan as someone who understands business but lacks deep expertise in those specific functions. It looked reasonable to me. But I had no way to verify whether an actual sales leader or marketing leader would find it adequate.

That concern aside, I had a more immediate problem. The plan was written as internal strategy notes. It was not formatted as something I could hand to my CEO client.

I needed it packaged as a professional deliverable.

Chapter 3: Switching to Consultant Mode for Proper Formatting



By this point, the output from the sales and marketing leaders had improved significantly through the accountability I imposed by positioning myself as the CEO demanding better tactical detail.

The plan now included quarterly timelines, target segments, channel strategies, and budget considerations. It looked reasonable to me given my background as a CIO with general business knowledge.

But the plan was written as strategic thinking notes rather than a polished deliverable. I needed something formatted professionally that said "Mr. CEO, here is the plan I recommend you implement" rather than "here are some thoughts about strategy."

This was the obvious next step. I needed to change the AI's role to match the output format I required.

The Consultant Role Prompt:

"Thanks for the work. Now assume a new role. You are a senior consultant for a high-end management and leadership consulting agency. Take this tactical plan and package it into a comprehensive, timeline-based document for delivery to the CEO who is paying for this work."

What Changed

The plan suddenly had proper consulting document structure:

- Executive summary sections
- Clear section headers and subsections
- Professional formatting conventions
- Timeline matrices showing quarterly breakdowns
- Deliverable specifications
- Budget tables
- Measurement frameworks

It now looked like something a consulting firm would produce.

But even with the improved formatting, I still faced the fundamental expertise problem. As a CIO looking at a sales and marketing plan, it appeared solid to me. However, I had no way to validate whether someone who actually implements these strategies daily would find it adequate.

More importantly, I could not afford to look foolish in front of my CEO client by delivering a plan that fell apart under expert scrutiny. I needed validation from perspectives I did not possess.

I needed to find ways to get AI to validate other AI work, since I lacked the domain expertise to do it myself. This realization led me to try something. Some of the talking heads literally say "speak to it like a person." — Ok. Let's try that then. I decided to try the tactic of "Would YOU accept this work?"

Chapter 4: I Made the AI Refuse to Pay

I had been using one AI system for everything up to this point. It had played the sales leader, the marketing leader, and then the consultant producing the formatted deliverable.

I decided to open a different conversation window using the same AI platform. This gave me what I call a sibling conversation with (hopefully) no shared context about what had been produced so far.

I did this deliberately. I wanted to pit two instances of the AI against each other without contaminating the contexts. I was not certain whether they shared memory or tokens behind the scenes, but I wanted to try keeping them separate. This would let me copy content back and forth between the consultant conversation and what I was about to create as the CEO reviewer conversation.

I gave this second conversation a very specific role, and I fully expected it to reject the work. I had done similar things before with varying success, so I knew there would likely be multiple rounds of back and forth.

The CEO Refusal Setup:

"You are the CEO of Widgets Inc. A very expensive consulting firm just delivered this sales and marketing plan to you. They want you to pay them a significant amount of money for this work. Read the plan carefully. Are you going to pay them? If no, provide specific feedback on what is missing or inadequate. Be direct about what would prevent you from paying."

Then I pasted the entire consultant-formatted plan into this conversation.

What happened next reinforced my decision to separate the conversations.

CEO Response - First Rejection



CEO Response - First Rejection:

"I am not paying for this yet. Here is why.

The plan tells me to send 500 physical field guides to prospects in February. Then it tells me to follow up with personalized video emails 48 hours later. **But where am I getting their email addresses?** The prospect list you specified only has physical mailing addresses. If I cannot email them, the entire video strategy collapses.

The plan also instructs me to place non-responsive prospects into a 3-email nurture sequence. Where are those emails? You wrote the landing page copy. You wrote the video script. But you did not write the nurture emails. I cannot execute a nurture sequence that does not exist.

Fix these gaps and resubmit."

This was dramatically different from anything the consultant conversation had produced. The CEO was not trying to make the plan work or defend choices. It was acting like someone who had to implement the work and was identifying every operational gap that would cause failure.

I took this feedback back to the first conversation window with the consultant and told it to address the problems.

The Revision Cycle

📄 My Instruction to the Consultant:

"The CEO has rejected your plan and is refusing to pay. Address every issue he identified. Resubmit a revised version."

The consultant produced a new version that added:

- A data enrichment protocol explaining exactly how to obtain email addresses from physical mailing lists
- Complete nurture email sequences with subject lines and body copy
- Budget confirmation that the tech stack allocation covered data enrichment costs

I sent this revised version back to the CEO conversation.

CEO Response - Second Rejection:

"Better. But I am still not paying.

You have me printing a 20-page Field Guide that makes specific promises about our software capabilities. The guide states we can detect Widget Production Anomalies and Monitor Real-Time Assembly Efficiency. **Does our engineering team actually have these features built?** If a prospect asks for a demonstration and we cannot show it, we lose all credibility.

Also, the plan promises we have a dataset of 50,000 manufacturing efficiency analyses. Do we actually have that data? Or is this aspirational? If it is aspirational, the language needs to change immediately.

Resubmit with technical validation."

This second rejection caught something critically important. The marketing materials were making promises about product capabilities without verifying those capabilities existed.

This is exactly the kind of problem that destroys credibility with customers. The consultant had written compelling copy without checking whether the company could deliver what the copy promised.

Five Rounds of Rejection

I went back and forth between the consultant and CEO conversations three more times. Each round, the CEO caught different problems:

01

Missing artifacts that were referenced but never created

02

Vague instructions that could not be executed without more detail

03

Budget allocations that did not account for all required costs

04

Timeline dependencies that were not properly sequenced

05

Metrics that sounded good but could not actually be measured with available systems

By the fifth round, the CEO finally approved the plan.

But even with CEO approval, I was still concerned about handing this to my actual client. AI is AI. While I am reasonably intelligent and the plan looked good to me after all these improvements, I am still not a sales expert or marketing expert.

I worried that someone with deep expertise in those areas might look at this and immediately identify fundamental problems I could not see. I was hunting for more ways to validate the plan before delivery.

That is when I thought about board-level review. What could be better than having something survive not just CEO scrutiny, but board scrutiny as well?

Chapter 5: The Board Review That Caught What We Missed



At this point, I had:

- A consultant conversation producing professional documents
- A CEO conversation that had rejected the plan five times until operational gaps were fixed
- A plan that had survived significant scrutiny

The plan was dramatically better than what I started with. But I was still operating with my own limitations as a CIO rather than a sales and marketing expert.

I was still concerned about looking foolish in front of my CEO client. I kept thinking there had to be more ways to validate this work. I needed another perspective that would scrutinize it differently than the CEO had.

That is when I decided to try board-level review. I deliberately opened a different AI platform for this. I used (chatbot 2) instead of the system I had been using for the consultant and CEO conversations.

I did this on purpose. I wanted a completely different AI model with a different training background to review the work. Different viewpoints meant different potential catches.

Before I gave it the board role, I actually had it design an appropriate board first.

Setting Up the Board

📄 My Setup Prompts:

"Create a board of directors appropriate for a B2B manufacturing analytics company."

It produced a list of board member profiles with various expertise areas.



"Here is the company's strategic roadmap for 2026. How does this change who should be on the board?"

It refined the board composition based on the strategy, swapping out some members and bringing in others with more relevant backgrounds.

Only after establishing the board composition did I give it the actual review role.

📄 The Board Directive:

"You are this Board of Directors for Widgets Inc. You have just received the sales and marketing plan drafted by a consulting firm and approved by the CEO. Your job is not to be polite. Your job is to identify every flaw, risk, assumption, unrealistic element, missing component, and execution gap. Be thorough. Be unforgiving. Nothing gets implemented unless you sign off. If you have any questions for the CEO, ask me. I will ensure they are answered."

I also told the board that the organization itself was facing serious challenges. The company had limited resources and was potentially facing layoffs if the pivot did not succeed. The success of this plan was paramount. I have done this many times before with AI. I try to create urgency and high stakes, hoping it will be more critical in its analysis.

The board asked me a dozen questions. It wanted to know about team size, current workload, technical capabilities, budget constraints, and timeline pressures. I answered these questions, providing context about the company's actual situation.

Board Response - Strategic Risks Identified

Board Response - Excerpt:

"The Board cannot approve this plan for execution. The following issues must be resolved:

Product Capability Misalignment:

The Field Guide explicitly states the software 'detects Widget Production Anomalies and Monitors Real-Time Assembly Efficiency.' Has engineering validated these capabilities exist in the codebase? If these features are not built, every prospect demo will fail.

Dataset Claims:

Multiple sections reference a dataset of '50,000 manufacturing efficiency analyses.' Does this dataset exist? If not, all references must be removed immediately. Claiming data assets you do not possess is a credibility disaster.

Market Assumptions:

The plan assumes a 60-70% email match rate when enriching physical mailing lists. What is this assumption based on? If the match rate is actually 40%, the entire Q1 volume model collapses.

Timeline Realism:

The plan targets 4 new contracts in Q1 from a single-person sales team while also rebuilding the website, writing all marketing materials, and managing existing customers. This is not realistic for a team of this size.

Technical Readiness:

Before any marketing materials are printed, engineering must validate every capability claimed in those materials. Add a technical validation checklist to the critical path."

The Board response continued for several more paragraphs. It identified problems at a completely different level than the CEO had caught.

Different Perspectives, Different Problems



The CEO had focused on operational gaps. Can we execute this plan with the team and resources we have?

The Board focused on strategic risks. Are we making promises we cannot keep? Are our assumptions validated? What happens if key assumptions are wrong?

This was the moment I truly understood how well this multi-conversation approach was working. I had used similar techniques before, but never at this depth or with these stakes.

Each conversation caught different categories of problems:

Consultant

Tries to produce impressive work but has natural blind spots

CEO

Catches operational and execution gaps

Board

Catches strategic risks and credibility threats

The Board had identified something that could have been catastrophic. Without this layer of review, the plan might have made technical promises the company could not demonstrate, destroying credibility in the first prospect conversation.

Chapter 6: The Iteration Process and Final Deliverable

After the Board rejected the CEO-approved plan, I sent the Board feedback directly back to the consultant conversation.

I made a deliberate choice here. In a real organization, the feedback would flow Board → CEO → Consultant. But these were my AI conversations. I could skip that intermediary step.

This decision might not have been optimal. Going through the full chain might have produced additional refinement. But I chose directness for efficiency.

My Instruction to the Consultant:

"The Board has rejected the plan. Revise to address every issue they identified. Do not defend your previous work. Fix the problems and resubmit."

The consultant produced a revised version that included:

- A technical readiness table distinguishing current capabilities from roadmap items
- Messaging guardrails specifying what could be claimed versus what had to be positioned as future development
- Revised marketing copy that did not promise features still in development
- Dataset language changed from specific numbers to "building the data foundation"
- Realistic Q1 targets reduced from 10 contracts to 4 pilot partners to account for team size

I sent this back to the Board. The Board approved it with some minor additional clarifications.

The Final Deliverable

The final plan included:

Strategic Components:

- Executive summary explaining market context and pivot rationale
- Quarterly execution roadmaps with specific tactics, owners, and deadlines
- Budget breakdown showing allocation across channels with expected ROI
- Risk assessment identifying key dependencies and mitigation approaches

Tactical Depth:

- Complete prospect lists with data enrichment protocols
- Email nurture sequences with subject lines and body copy
- Video audit scripts for sales outreach
- Direct mail specifications including printing and postage requirements
- Landing page copy aligned with all campaign messaging
- Technical readiness tables showing current state versus roadmap

Governance Controls:

- Messaging guardrails preventing overpromising
- Validation checklists ensuring marketing claims matched product reality
- Phased rollout language for capabilities still in development

The final deliverable included multiple appendices with ready-to-use templates and scripts.

I still had to read every single word carefully. The AI conversations left artifacts I had to clean up. They would include phrases like "revised and updated" in odd places. They would add notes like "Appendix B unchanged" that I had to remove. Some language from the role-playing scenarios remained embedded in the text.

What the process did was enable me as a non-sales and non-marketing person to go through extensive research, editing, and improvement cycles. The result was something that actually seemed like a reasonable sales and marketing plan built from the strategic roadmap we already had.

I built the governance system that forced the quality. I did not write the content, but I created the accountability structure that made the content better.

About the Author

David Logan

Founder & CTO, Your Software CTO

Contact:

david@yoursoftwarecto.com

720-820-2286

[linkedin.com/in/davidjlogan](https://www.linkedin.com/in/davidjlogan)

yoursoftwarecto.com