

# Kill the Battle Card

The AI method for dramatically improving enterprise win-loss rates — by replacing static battle cards with a knowledge base and generating account-specific positioning for every deal.

Daniel Glickman · Product Marketing & AI Leader

## EXECUTIVE SUMMARY

Competitive enterprise deals turn on one thing: whether the seller can answer the buyer's real criteria and show, precisely, where the competitor falls short. Static battle cards do not do that — reps do not use them — so most teams leave the work undone.

This paper describes the method I built to do it at scale: replace the generic battle card with an AI-consumable competitive knowledge base, extract each buyer's stated criteria from call transcripts, and generate account-specific positioning for every deal. In the enterprise segment where I ran it — on deals above \$30,000 where the competitor was known, with no-decision outcomes excluded — the method held a roughly 74% win rate, and the average discount on competitive deals fell from 18% to 4%. It gets there with less manual work, not more: the AI does the slow analysis, and the strategist solves what is genuinely new.

## 01 · THE PROBLEM

### **The battle card that never got opened**

Every product marketer has built a battle card. Most have built dozens. Sales asks for them by name — a one-page competitive cheat sheet, one card per competitor, a few bullets on how to win — so you build them, you keep them current, and you post them where the team can find them.

Then you talk to the reps, and you hear the same thing I did: they don't use them.

Not because the cards are wrong. Because a card is the wrong shape for the moment a seller is actually in. A battle card is generic by construction. It answers "how do we talk about Competitor X" in the abstract, for every deal at once, which means it answers no specific deal well. The rep in front of a real buyer does not have a Competitor X problem. They have a *this account, these three requirements, this stakeholder, this stage* problem. The card was written before any of that was known, so it stays in the drawer.

There is a reason companies fall back on the card anyway. The thing that actually works — bespoke competitive strategy built for one specific deal — is exactly what the best account strategists and most proficient enterprise AEs already do. But by hand it is slow. Pulling a buyer's real requirements off the call recordings, matching each one against what you know about the competitor, and turning that into positioning takes hours of expert time per deal, so it happens on a few marquee opportunities and nowhere else. The battle card is the compromise you reach when the real work does not scale. It is not the thing that works; it is the thing that is cheap to produce.

That is the gap I set out to close. Not "make better battle cards." Bridge the distance between the competitive content a company produces and what an enterprise seller can actually use in a live deal — which turns out to be a different artifact entirely, produced a different way. And the reason it can scale now, when it could not before, is that AI does the slow part. This method is not more work layered onto the deal. It removes the manual cost that kept real competitive strategy off most deals in the first place.

The rest of this paper is the method I built to do that, and the AI system I built to run it. The claim is simple, and I will support it with the one number I can stand behind: when a team knows exactly who it is competing against and responds to that competitor on the buyer's own criteria, it wins far more often. In the enterprise segment where I ran this — on deals above \$30,000

where the competitor was known, with no-decision outcomes excluded — the method held roughly a 74% win rate. The lever is not louder messaging. It is competitive clarity, applied to one deal at a time — on every deal, not just the ones important enough to justify the hours by hand.

02 · THE REFRAME

## **Enterprise selling is account strategy, not messaging**

The reason battle cards fail is that they treat competitive selling as a messaging problem. It is not. In an enterprise deal, it is an account-strategy problem.

A messaging problem has one answer that you broadcast to everyone: here is how we describe ourselves against this competitor. An account-strategy problem has a different answer for every deal, because every deal has a different buyer, a different set of stated requirements, and a different reason the competitor is in the room. The work is not to say the right thing in general. It is to do the right analysis for *this* account and position against *this* competitor on the criteria *this* buyer already told you matter.

So the process starts with the account, not the collateral. You analyze the account in detail. You map the buyer's stated criteria. You establish what actually matters in this deal — not what would be convenient for us to pitch. Then, and only then, you position against the named competitor, criterion by criterion. That is more steps than skimming a card, and if a person had to do all of it by hand on every deal, it would never happen — which is exactly why, historically, it hasn't. The point of the system in the sections that follow is that a person no longer has to.

*Rely on what the customer says matters, not on what your company says matters.*

Those are two different lists, and the gap between them is where deals are lost. Your internal team has a list of features it is proud of and wants to lead with. The buyer has a list of requirements they will actually judge you on. When you pitch the first list, you sound like every vendor. When you answer the second, you sound like the only one who was listening.

Getting the buyer's real list is not guesswork. It is in the record. The requirements are stated on the calls — in the discovery, in the demos, in the offhand line where a stakeholder says what would make this easy to sign off. Pull the requirements from the transcripts. Treat what the buyer said as the specification. Everything downstream — the positioning, the proof, the materials — is built to answer that specification and nothing else.

03 · THE CORE SHIFT

## **Replace the battle card with a knowledge base**

If the account is the unit of work, the battle card is not just too generic — it is the wrong *artifact*. And this is the shift that everything else depends on.

A battle card is optimized for a human to skim. It is short by design: one competitor, a handful of talking points, a page you can glance at before a call. That format has a hard ceiling. A talking point is a conclusion with the analysis stripped out, so a rep can either repeat it or not, but cannot reason with it. The moment a real deal needs something the card did not anticipate, the card has nothing left to give.

Replace it with a knowledge base: a structured repository of competitive intelligence, optimized for an AI agent to consume rather than for a human to memorize. Where a card has a bullet, the knowledge base has the full picture behind it — specific feature-level differentiation, technical detail, the competitor's public claims, and the documented gaps between what they

claim and what they deliver, each one anchored in fact and technical analysis rather than in a confident sentence. It does not have to be short, because nothing is reading it end to end. An agent retrieves exactly the parts that match the deal in front of it and leaves the rest.

That difference in depth is the whole point. A human-readable card can hold one thing about a competitor. An AI-readable knowledge base can hold everything you know about that competitor, at full resolution, and hand back only the slice that answers this buyer's stated criteria. The card forces you to decide in advance what will matter. The knowledge base lets the deal decide.

The non-negotiable element of this method is that the repository exists — deep, current, and honest. *How* you build and maintain it — the sourcing, the pipelines, the update cadence, the human review that keeps it accurate — is real work, and it is deliberately outside the scope of this paper. What matters here is the principle: the thin-talking-point paradigm gives way to a fact-anchored knowledge-base paradigm, and a person stays in the loop over both what goes into the repository and what comes out of it. That is the substrate. The next section is the loop that runs on top of it.

#### 04 · THE METHOD

### The loop

On each active deal, the method runs as a short loop — five steps, then one output.

- 1 Extract the buyer's criteria.** Take the call transcripts for the deal and pull out what the buyer actually said they need, in their words. This is the specification the rest of the loop answers to.

- 2 **Identify the competitor.** Determine who else is in the deal. If the competitor is named on a call, that is the answer. If they are not named yet, that itself is a finding — an unknown competitor is the most expensive gap in a competitive deal, and the loop flags it for the team to resolve rather than guessing.
- 3 **Map each criterion to the competitor.** For every requirement the buyer stated, pull from the knowledge base what the competitor claims about it and where the competitor is actually weak. The output is a criterion-by-criterion picture: here is what they need, here is what the competitor says, here is the gap, here is where we are strong against that gap.
- 4 **Generate deal-specific positioning.** Turn that map into the positioning for this deal, and only this deal — not the generic competitive narrative. For each stated requirement, the seller gets what the competitor claims, where they are weak, and how to position us directly against that weakness. Everything that does not bear on this buyer's criteria is left out.
- 5 **Escalate the unknowns.** When the knowledge base does not have an answer — a requirement no one has solved for, a competitor claim we cannot yet counter — the loop does not paper over it. It routes the gap to a human — the strategist and their team. They work out how to actually meet the requirement, which can mean brainstorming, pulling in services, or engineering an answer that is real and defensible. Then that answer goes back into the knowledge base, and the next deal inherits it.

The output of the loop is not a document. It is a customized buying experience. Once you know what the buyer needs and how you win against the competitor on those terms, you make that show up where the buyer feels it: in the demo, in the collateral, and in the proof you put in front of them. For a six-figure enterprise deal, that tailoring is worth the effort — and, as the next sections show, AI makes it cheap enough to do every time.

**An illustrative example.** The following is hypothetical – invented to show the mechanics, not drawn from any real deal or competitor. A buyer is evaluating us against an established single-product incumbent. Three criteria surface from the transcripts, and the system maps each one against the knowledge base to produce the account-specific positioning. In practice, that output is a table:

ACCOUNT-SPECIFIC POSITIONING – HYPOTHETICAL, FOR ILLUSTRATION

THE BUYER'S CRITERION	THE COMPETITOR'S KNOWN GAP	HOW WE POSITION AGAINST IT
The rollout cannot create friction with employees	Deployment is heavy-handed and visible to end users	A low-visibility rollout designed to avoid employee pushback
The data has to stay in the buyer's region	No regional data-residency option	A residency guarantee the incumbent cannot match
It has to fit the HR system already in place	Only a shallow, unsupported integration	A native, supported integration

That table is the account-specific battle card – and it resolves the tension the paper opened with. The generic battle card fails because it is made once, for everyone, in the abstract. This one is its opposite: built for a single deal, drawn from fact, scoped to what this buyer actually said, and produced by the system in seconds rather than by a person over hours. This is the version reps use, because it is about the deal in front of them.

And it is the seed, not the finish line. From that same positioning, the system generates the rest of the deal's materials automatically, each off a known template: a slide for the demo deck, a summary slide for the buying committee, proof-of-concept design instructions where a POC applies, and whatever other collateral the deal calls for. The strategist decides what is true and what matters; the system produces the artifacts that carry it into

the buying experience. That is what customizing everything for an enterprise deal looks like once the customizing is no longer manual.

05 · THE BUILD

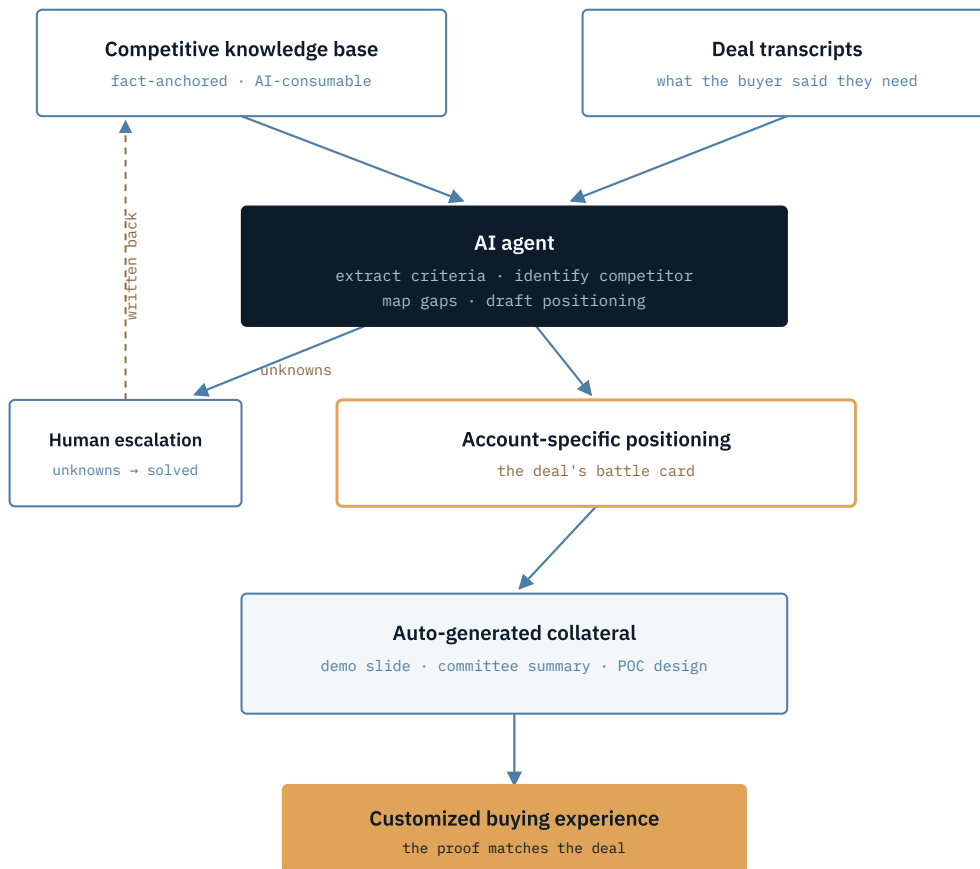
## **The system that runs it**

The loop is a method; what makes it run on every deal instead of a few is that it is a system, and I built it.

Start with the knowledge base. An automated monitoring layer tracked the competitive field — 24-plus competitors — and kept their claims, their positioning, and their gaps current, with a person reviewing what went in so the repository stayed accurate rather than just large. Keeping it fresh was itself automated, not a quarterly scramble.

On an active deal, agents did the per-deal work. An account-intelligence agent ingested the transcripts and assembled the picture of the account. A MEDDIC agent ran the qualification and pulled the buyer's stated criteria out of the conversation. Those criteria were matched against the knowledge base, and the deal-specific positioning was drafted — the criterion-by-criterion map from the loop, produced in the time it takes to read this sentence rather than the hours it would take by hand. Unknowns were routed to the team. All of it lived inside the sales workflow the reps already used, not in a separate tool they had to remember to open.

This was not a prototype. It ran as part of a set of roughly ten production AI agents doing real go-to-market work in a real organization, with a human in the loop at every point that mattered — over what entered the knowledge base, and over anything the system was not sure about.



The AI-Powered Win-Rate Method: the competitive knowledge base and the deal's transcripts feed an agent that extracts the buyer's criteria, identifies the competitor, maps weaknesses, and drafts the account-specific positioning. Unknowns route to a human and are written back. From the positioning, the system auto-generates the deal's collateral – the customized buying experience.

## Less work, not more

It is fair to read the loop and think it sounds like a lot. It is a lot — that is the point. It is the work that great account strategists and enterprise AEs have always known they should do and mostly could not, because doing it by hand costs hours per deal, and hours per deal do not scale across a full pipeline.

AI changes the economics in two directions at once. It removes the manual cost of the analysis — the transcript reading, the criteria mapping, the competitor matching, and the first draft of the positioning — so the expensive part is no longer expensive. And it removes the manual cost of the output — tailoring a demo, branding a one-pager, assembling proof for a specific buyer — so per-deal customization, which used to be reserved for the biggest opportunities, becomes the default. The net is not more work added to the deal. It is less: the strategist's judgment goes to the hard, human parts, and the slow parts are gone.

I can put a number on the direction, though not on this method specifically. Across the AI program I ran, that shift showed up as roughly 50% faster delivery and about twice the output, with no added headcount. That is the shape of the change — the same team doing markedly more, faster. Applied to competitive deals, it is what let a practice that used to fit a handful of marquee accounts run on every deal in the segment.

There is a second economic effect, and it shows up in price. When a seller can answer a buyer's real criteria and show exactly where the competitor falls short, the deal stops being a discount negotiation — you win on fit, not by giving margin away. In the segment where I ran this, the average discount on competitive deals fell from 18% to 4%. That is the clearest sign that precise positioning, and not price, is doing the work.

win-rate	competitor known · \$30K+ segment · no-decision excluded	~74%
discount	average discount on competitive deals – re-engineered	18% → 4%
throughput	AI program – faster delivery / more output, no added headcount	~50% · ~2x

## **The human layer is the point**

None of this replaces the strategist. It aims their time.

The system is good at exactly the parts that are mechanical: retrieving what a competitor claims, matching it against stated criteria, and drafting a first pass of the positioning. It is not the thing that decides what to do when the answer is not known. That is the human layer, and it is not a fallback — it is the point. When a buyer states a requirement no one has solved for, or a competitor makes a claim the knowledge base cannot yet counter, the system escalates rather than inventing. A person works out how to genuinely meet the need — brainstorming it, bringing in services, or engineering an answer that is real and defensible — and the resolution is written back so the system is smarter on the next deal.

This is also the guardrail against the failure mode everyone worries about with AI in a sales motion: generic, confident, wrong output. I do not want a seller armed with plausible competitive language that falls apart under a technical question. The knowledge base is fact-anchored, the positioning is scoped to what the buyer actually asked, and the unknowns are handed to people instead of guessed at. The AI does the volume. The humans keep it honest and solve what is genuinely new.

08 · THE PLAYBOOK

### **How to start**

You do not need the full system on day one. The method has a maturity curve, and every stage is useful on its own. Find where you are, and take the next step.

#### STAGE 0

### **Battle cards**

Static, generic, one card per competitor. If reps are not using them, that is not a content problem to fix with better cards. It is a signal to change the artifact.

**next move:** ask five reps whether they open the cards in a live deal, and listen to the answer.

#### STAGE 1

### **A real knowledge base**

Replace the cards with a structured, fact-anchored repository built to be read by a machine: feature-level differentiation, technical detail, competitor claims, and documented gaps.

**next move:** pick your top competitor and write down everything you actually know about them at full depth, not as talking points.

#### STAGE 2

### **Criteria from transcripts**

Start pulling buyer requirements from call recordings instead of assuming them.

**next move:** on your next competitive deal, extract the buyer's stated criteria from the transcript before you decide how to position.

#### STAGE 3

### **Agent-generated positioning**

Let an AI agent match those criteria against the knowledge base and draft the deal-specific positioning.

**next move:** automate the match for one deal and compare the output to what a rep would have improvised.

#### STAGE 4

### **Customized buying experience**

Push the positioning into the demo and the collateral for each significant deal, and route what you cannot answer to a human to solve.

**next move:** build one tailored demo for one six-figure deal and measure whether it moves.

CLOSE

### **The magnet is the method**

I am Daniel Glickman, a product marketing and AI leader. I build the AI systems that run go-to-market rather than only talk about them — work the Product Marketing Alliance recognized with its 2024 AI Marketing Innovation Award. The method in this paper is one of those systems.

The claim it rests on is narrow on purpose: in the enterprise segment where I ran this — on deals above \$30,000 where the competitor was known, with no-decision outcomes excluded — answering the buyer's own criteria held roughly a 74% win rate. The rest of the paper is how to get there without adding work.

This paper is free. Download it, use it, and send it to whoever runs competitive deals in your organization. If you want to talk about applying it — to your team, or to a role — leave your email below and I will follow up. No gate, no sequence.