

The Promise and What's *Underneath* It

What fifteen years of go-to-market technology taught us, and why the lesson matters more with agentic AI.

A MARKETING AFFECTS POINT OF VIEW

For revenue leaders across marketing, sales, and customer success.

By Phylis Miquel · Marketing Affects · 2026



WHAT YOU'LL TAKE AWAY

- 1 Why fifteen years of go-to-market technology underdelivered.
- 2 What AI speed amplifies, for better and for worse.
- 3 The one move to make before your next AI investment.

Every few years, go-to-market makes a serious technology bet. Over the past decade it was marketing automation, then data platforms, then revenue intelligence suites. Each arrived on a clear and reasonable promise: more pipeline, a tighter forecast, stronger retention, higher close rates, less wasted effort. Leaders invested accordingly, and they were right to. This year the bet is artificial intelligence, and for the next few years it will dominate every roadmap and budget line. It is the biggest one yet, and it arrives with the highest expectations and the most stretched resources of any wave so far.

This paper is not about whether that spending was worth it, or whether every tool was fully used. That debate is both unfair and beside the point; the investments were rational, and much of the capability is real. It is about something quieter and more useful. Across every one of those waves, what did we learn about the distance between a technology's promise and the value it actually delivers? And are we about to apply that lesson to AI, or spend straight past it?

The lesson is remarkably consistent, and it is not about the tools. A technology keeps its promise only when the underlying system is ready to carry it: when the words that drive revenue mean the same thing across teams, when the data is clean and connected, when the people have the time and context to run it. That readiness is what turns a capable tool into a realized result, and it is the one thing a tool cannot define for you.

AI changes what that lesson costs. Every technology before it reported and then waited for a person to act. AI does not wait; it acts. The gap between what a tool

promises and what the underlying system can actually support used to surface slowly, in a review, where someone could catch it. Now it plays out on its own, at machine speed, across the whole revenue motion. That is what this paper is about, and what to do about it before the next wave of spending.

01 THE PATTERN

The promise, wave after wave

Start with the scale of it. In 2011, when Scott Brinker first mapped the go-to-market technology landscape, there were roughly 150 marketing technology products. By 2025 there were 15,384, about a hundredfold increase in fifteen years. Sales and customer success built out their own stacks alongside marketing's. Each new category made a promise, and organizations bought against it, integrating tools to fit their go-to-market strategy.

150 → 15,384

marketing technology products mapped between 2011 and 2025, about a hundredfold rise.

SOURCE: BRINKER

Now the honest part. By Gartner's measure, organizations use less than half of the capabilities they have paid for, and in some years as little as a third. It is tempting to read that as underuse, but that framing misses what actually happened. The technology usually worked. The promise was kept where the underlying system was ready to carry it, and it stalled where that system was not.

under half

of the martech capabilities organizations pay for are in active use, and in some years as little as a third.


SOURCE: GARTNER

The pattern in what stalled was consistent, and it was rarely the software. It was the readiness beneath it. The data had not been connected, the problem leaders name as their single biggest stack challenge. The definitions had not been reconciled across teams. And the people were already stretched by the day-to-day, without the hours to run one more system well. A capable tool fed by disconnected data and run without time to learn it does not fail loudly; it produces confident-looking output on a shaky base, and quietly underdelivers on the promise that sold it.

None of that is a knock on the tools or the leaders who bought them. It is the nature of technology inside a revenue system. What actually turns a capable tool into a realized result is not something you can put on a purchase order. It is three things:

- a single definition of the customer that every team agrees to act on;
- clean, connected data and the operating discipline to keep it that way;
- people with the business context and market knowledge to direct the work and judge whether it worked.


These take collaboration and agreement, and they call for judgment a tool does not have. No software has ever included them, and none ever will.



The readiness the promise assumed

Of everything the underlying system requires, one piece matters most and is the least visible: agreement. Not agreement on tactics, but agreement on what good looks like, what a good customer experience should be, and what to do next, shared across every team that touches the customer.

Take a single word as an example. *Engaged*. To marketing it means activity on its channels. To sales it means a real person returning calls. To customer success it means the product is being used deeply enough to renew. One word, three readings, one account. And engaged is only the start. Qualified, healthy, at risk, ideal customer, good revenue, ready to expand: each fractures the same way. A real go-to-market system depends on dozens of these being defined, agreed, and connected across the lifecycle.



The answer is not to force every function to measure the same thing; they shouldn't. Customer success should read an account differently than marketing does. The answer is a shared definition: a single agreed meaning that each function translates

into its own view, so three readings still ladder up to one truth. This is the readiness every tool has quietly assumed and few organizations have built. Without it, each new system layered on top simply renders the same disagreement in higher resolution.

This is where the gap between promise and readiness stops being abstract and starts costing real money, and you can see it in pipeline, in revenue, and in wasted hours.

WHAT THE READINESS GAP COSTS, IN HARD NUMBERS

15–25%

of annual revenue lost to poor CRM data quality

30–40%

of the average pipeline is phantom, kept alive by stale records

~\$32K

per rep each year, lost working around bad data

\$12.9M

average yearly cost of poor data quality per organization

SOURCES: VALIDITY, STATE OF CRM DATA MANAGEMENT 2025; GARTNER.

It also quietly raises the cost of acquiring a customer: when the same buyer shows up as three records, the spend still counts and the outcome does not, so effective CAC drifts up while no one ever decides to let it. We have written more about that gap between data and trust in [More Dashboards Did Not Make You More Confident](#).

03 A QUICK TEST

The ten-minute one-account test

You do not need a methodology to test for that readiness in your own business. You need one account and three people. Pick an active account everyone knows. Separately, before anyone confers, ask the marketing owner, the sales owner, and the

customer success owner the same two questions: in one word, what is this account's status, and what is the right next move? Then put the three answers side by side.

Function	What it sees this week	Status	The next best action
Marketing	Opens, clicks, and content downloads are up	Engaged	Enroll the account in a new nurture campaign
Sales	Two missed meetings, silent for three weeks	Stalled	Drop priority and stop outreach
Customer Success	Product usage down, support tickets climbing	At risk	Trigger the churn-save playbook

One account, one week. Three functions, three one-word answers, three contradictory next best actions for the same account. In most companies one team formally owns the account at any given stage, yet each function still holds its own read, and each function's tools still act on that read. Today a person reconciles it, usually in a review that comes too late to matter. The question this paper turns on is what happens when those next moves stop waiting for the review and start firing on their own. Run the test with one account and three colleagues, and you will find your version of the gap within the hour.

04 THE SHIFT

Then AI arrives, and the promise starts acting

Now bring AI into that system. It arrives from two directions at once. The first is the AI already inside the tools you own: the copilot in the martech suite, the scoring model in the CRM, the sentiment engine in the customer success platform. It is there for a

reason that has little to do with your readiness to use it. Vendors are under real pressure to add AI, because the market now discounts those without an AI story: software valuations have fallen to decade lows, down roughly 40 percent year over year, and boards and investors expect an AI narrative on every roadmap. As a result, every product you already pay for now ships an AI feature, whether or not the system underneath is ready to make use of it. The second direction is the agents you are being encouraged to add on top, to draft, research, and reach out at scale. Our companion guide, *The Agentic Workforce*, lays out the difference between a tool you operate, an assistant you prompt, and an agent you direct.

The difference that matters is what each does with a gap. Every technology before AI, however smart, mostly reported. It scored a lead, drew a chart, flagged a risk, and waited for a person to act. That pause was where a shaky definition got caught before it did any harm. AI does not pause. Given a goal, it decides and acts. The same unreconciled definition that once produced a misleading dashboard now produces an action, taken before anyone reviews it.

This is why the lesson from the last fifteen years matters now more than it ever did. AI does not build the underlying system. It runs on it, faster, and it acts on whatever it finds. Where the readiness is there, AI compounds the advantage, which is the subject of *The Mesh*. Where it is not, AI executes the gap, the case we made in *AI Compounds What's Underneath It*.

The promise and the readiness underneath it used to be separated by a human and business context. Now they are separated by nothing.

Adding AI is not the same as adding intelligence

Two cautions come with the AI wave, and both are about that same distance between promise and delivery. The first is that not all of the promise is ready to deliver yet. Gartner separates genuinely agentic products from what it calls agent washing, existing automation or a chatbot described in agent language, and finds the genuinely agentic still rare today. That is less a reason to distrust vendors than a reminder to look past the label at what a tool actually does. And even where the capability is real, buyers are finding it underdelivers at first, not because the product is broken, but because it is switched on over a system that is not yet ready to feed it. The mood is turning, from racing to deploy toward asking what it takes to make it pay off, and that turn is not a verdict on the technology or the vendors. It is what happens when a promise outruns the readiness underneath it.

~130

of the thousands of vendors marketing agentic AI are genuinely agentic today, by Gartner's count.

SOURCE: GARTNER, 2025


45%

of martech leaders say the AI agents already built into their tools are not yet delivering the performance they expected, a readiness gap, not a broken product.

SOURCE: GARTNER, 2025

The second caution is that even real AI rarely coordinates. The AI inside your martech tool and the agents you add on top mostly do not talk to each other. Standards are emerging, so they increasingly can, but connection is not coordination. Wiring two models together does not make them agree on what qualified, engaged, or healthy means. Each one still acts on its own reading.

Here is what that looks like from the only seat that matters, the customer's. Picture a customer six months into using your product. A sales rep's outreach pitches them the very thing they already bought, because the record never caught up. A marketing track enrolls them in a nurture campaign and a webinar for that same product, ignoring the support ticket they opened last week. A customer success play opens an expansion and renewal nurture, even as the account's own usage has been sliding for weeks. Three different systems reached out that month. Each was fluent. None knew about the others.



Nobody designed that journey. Three systems each acted from a different read of who the customer is. That is not a technology failure. It is the same missing agreement, now reaching customers automatically.

The same pattern, function by function

That was the customer's view, from the outside. From the inside, the same gap shows up function by function. The table is the short version, and the narrative underneath adds the why.

Function	What now runs on its own	What used to pause it
Marketing	Scoring, routing, and nurturing leads	A rep's read on whether the lead was real
Sales	Updating stages, shaping the forecast, enriching records	A manager's read on the number
Customer Success	Triggering renewal and expansion plays off the health score	A CSM's call on how the account is really doing

Read down the table and the through-line is simple. In every function, work that now runs on its own used to pause on someone's desk. Lead routing waited for a rep's judgment. The forecast waited for a manager's read. The renewal play waited for a CSM's call. Those pauses were friction, and they were also the last point at which a human applied a shared sense of the account before the company spent money or made a commitment.

Human judgment runs optimistic in every function, not just one. Marketing counts engagement that flatters the funnel. Sales forecasts the deal it wants to win. Customer success rates an account green that churns a month later. The health score is the vivid case, because it compresses usage, support, sentiment, and tenure into one number that, by Gainsight's 2025 benchmark, three quarters of teams still assemble by hand, but the pattern is the same everywhere. When AI acts on any of these optimistic reads, the optimism scales with it. The opportunity is the same in every row: get the shared

definition right first, and what acts on it inherits good judgment instead of a hopeful guess.

07 COMPOUNDING

Where a small gap becomes a large one

There is one more effect, and it only appears when you step back from the individual functions and look at the whole motion.

A go-to-market system is a chain of handoffs. Marketing hands to sales, sales hands to customer success, and each handoff is a translation between those three readings of the same account. When a person makes the handoff, the mismatch usually surfaces as friction, a reworked deal, a surprised renewal, a debate in the quarterly review. Slow, but visible, and catchable.

When the handoffs run automatically, each step acts on its own reading and passes the result to the next, which treats it as settled and acts again. The small mismatch at the first handoff is not corrected. It is built upon. By the time the quarterly review meets, the system has already routed, spent, and renewed against it.



The gap grows at each handoff



Each step treats the last step's read as settled, and acts again.

The review that was meant to catch the drift runs in months. The automation runs in minutes.

The opportunity, again, is upstream. Reconcile the definition once, and the handoffs carry good information at machine speed instead of compounding a small error into a quarter. This is the difference between a system that gets smarter as it runs and one that just runs harder, which we explored in [*Your Team Is Executing. Your System Is Not Learning.*](#)

And the cost does not stay inside the company. When the business is sold or raised against, buyers look past the revenue to the operating system that produced it. Revenue that compounded cleanly holds up under that scrutiny. Revenue assembled automatically from three conflicting reads raises questions, and questions lower the price. The gap you left unresolved upstream is the one that shows up in the valuation. For a public company it is the same story on a faster clock: analysts re-rate the model behind the forecast quarter after quarter, long before any sale.

08 THE PEOPLE

The capability nobody is hiring for

The readiness has a human dimension too, and it is the part most companies are least prepared for. When you put AI into a function, the human job changes. The marketer stops drafting and starts directing the system that drafts. The sales leader stops running the play and starts coaching the systems that run it. McKinsey calls the people who can do this M-shaped talent, generalists who can orchestrate a blended team of people and machines and steer it toward an outcome. The titles are starting to appear, and quickly: by hiring-tracker data, postings for the GTM engineer alone more than doubled in 2025, from roughly 1,400 to over 3,000. But most of these roles, GTM engineer, AI operations lead, agent orchestrator, are scoped for technical implementation, deploying prompts and wiring systems together. The scarcer capability is not technical, and IDC puts the cost of the broader skills shortage at 5.5 trillion dollars by 2026, with only about a third of organizations feeling ready to adopt AI-driven ways of working.

1,400 → 3,000+

GTM-engineer job postings more than doubled in 2025, even as the scarcer skill, business judgment, goes unhired.

SOURCES: APOLLO; GTMNOW

It helps to be precise about what that scarcer skill is, because the market got it wrong once already. The early advice was to learn prompting, the syntax of asking a model. The models have largely solved that. They read intent well enough now that clever phrasing is table stakes. What separates the people who get results is business context: knowing what the company actually sells, what the customer actually needs, and how to connect the two. They can say what good looks like, give the work the right picture of the business to run from, and coach it when the output misses. That is judgment, and it does not show up on a resume as a tool.


It shows up at every stage of the lifecycle. In marketing, it is knowing which message actually matters to which segment, not just generating more of them. In sales, it is knowing which deal needs a human in the room and which can be carried automatically. In customer success, it is telling a real risk from a noisy signal before the renewal play fires. Same skill, business context applied as judgment, wearing three different uniforms.

Which is the same lesson, one more time. Automation inherits whatever definition of the business its human director hands it. Give it a shared, reconciled one and the directing job becomes teachable, because good is defined once for the whole system instead of reinvented by every manager. Compete for AI engineers instead, and you have bought capability without the capacity to aim it.

The fork

Put it together and the choice in front of every operator is not whether to deploy AI. That spend is committed. The choice is what the AI lands on.


Start with deployment, because switching AI on alone is not the win. Each function can turn it on at its own pace, and a function working alone only sharpens its own view. The value shows up where the functions share a read of the customer and coordinate the handoff between them. That overlap is the whole game.



Turning AI on inside a silo is limited. The overlap, one shared definition of the customer, is where the value lives.

Now follow two companies forward. For private-equity operators we made this case portfolio company by portfolio company in [AI Compounds What's Underneath It](#); the divergence is the same for any business. They commit the same budget and buy similar tools. One has the readiness in place, end to end. The other does not. Both rise

at first, because AI delivers early efficiency that is easy to celebrate and easy to mistake for value. Then they part.



On a system with the readiness in place, AI compounds, because everything acts on the same picture and good information travels cleanly across the handoffs. On a system without it, AI does not simply plateau. It accelerates the gap, and the early efficiency is what hides it.

Same investment, opposite outcomes, decided by what sits underneath rather than by the AI.

10 THE VENDOR PROMISE

What no platform can do for you

There is a version of all this you will hear from the vendors, and it is the newest and most confident promise of them all. The agentic go-to-market platforms say they will

unify your customer profile, standardize your definitions, and orchestrate your handoffs for you. Buy the platform, the pitch goes, and the readiness comes built in.

It is a good pitch, and it is half true. A platform can enforce a definition. It cannot author one. It can make every function read from the same field, but it cannot decide what that field should mean for your business: which accounts are good revenue, what healthy looks like in your model, when a lead is genuinely ready. Those are judgments about your economics, and they are exactly the readiness the platform assumes you already have. Hand it an unreconciled definition and it will enforce the disagreement faithfully, at machine speed.

This is the same lesson, one more time, wearing the newest logo. The tool has never been the readiness. Buying a platform to create your shared definition is this year's version of buying martech to fix the forecast. It can carry the promise once the thinking is done. It cannot do the thinking.

11 WHAT TO DO

What this points to

Everything in this paper is a requirement the last fifteen years already taught us to expect. A shared definition across functions. Clean, connected data. The policies and feedback loops that keep a system honest. The people who can direct the work. None of it is new. What is new is that AI raises the cost of not having it, because the system now acts, automatically, on whatever is underneath it.

The first move, then, is not another purchase. It is to capitalize the lesson we have already paid to learn. Before the AI is switched on, the work is to make the words that drive revenue mean one thing across marketing, sales, and customer success, to decide what good looks like across the lifecycle, and to give every system, human and machine, the same definition to act on. It is unglamorous, it is mostly not technology, and it is the readiness every promise has assumed and few have built.

Revenue Integrity Architecture™ is the practical, named approach to putting that readiness in place: the shared definition across the revenue motion, the operating discipline that keeps it honest, and the learning layer that lets a go-to-market system compound rather than strain. It is the difference between AI that produces Compounding Enterprise Value™ and AI that simply runs the gap faster.

We have already paid for this lesson, wave after wave. The only question left is whether we capitalize it before the AI starts acting, or after.

Where does your system stand?

If you want a fast read on where your own go-to-market system stands, the two-minute readiness check is a place to start, and the [Revenue Integrity Assessment™](#) is the full diagnostic behind it.

Take the readiness check

Companion reference: The Agentic Workforce, a plain-language guide and a go-to-market map of human versus agentic work.

Related reading: [AI Compounds What's Underneath It](#); [More Dashboards Did Not Make You More Confident](#); [The Mesh](#); [Your Team Is Executing. Your System Is Not Learning.](#)

SOURCES AND REFERENCES

1. Marketing technology landscape, 150 to 15,384 products, roughly 100 times growth since 2011: Scott Brinker and Frans Riemersma, 2025 Marketing Technology Landscape, [chiefmartec.com](#) and [martechmap.com](#).
2. Martech capability utilization below half, 33 percent in 2023 and about 49 percent in 2025, with data integration the top constraint at 65.7 percent: Gartner Marketing Technology Survey, 2023 to 2025.
3. Cost of poor CRM data quality, 15 to 25 percent of annual revenue; 30 to 40 percent of pipeline phantom from stale data; roughly 32,000 dollars per rep per year in lost time; 37 percent of users lost revenue directly; about 16 opportunities lost per quarter: Validity, State of CRM Data Management 2025, and related industry data-quality research. Average cost of poor data quality at 12.9 million dollars a year per organization: Gartner.
4. Customer success health-scoring still manual at most firms, and the optimism bias in manual scoring: Gainsight, 2025 Customer Success Benchmark.
5. Vendor pressure to embed AI, and the decline of software valuations to decade lows (down roughly 40 percent year over year) for companies without a credible AI story: Sapphire Ventures, 2026 Software x AI; Octus, US Software Sector, 2026.

6. Agent washing, roughly 130 genuinely agentic vendors, and the forecast that more than 40 percent of agentic AI projects will be canceled by 2027: Gartner, 2025. Vendor-supplied AI agents underdelivering, 45 percent of martech leaders: Gartner, October 2025.
7. M-shaped talent and the agentic organization: McKinsey and Company, People and Organizational Performance, 2025 to 2026. AI skills-shortage cost, 5.5 trillion dollars by 2026, and AI-readiness levels: IDC.
8. Growth in GTM-engineer hiring, postings more than doubled in 2025 from roughly 1,400 to over 3,000: industry hiring trackers (Apollo; GTMnow), 2025 to 2026.
9. Shift from prompt engineering to context engineering and outcome clarity: industry coverage (Open Data Science and others), 2026.

The architectural concepts and methodology behind this argument are the proprietary work of Marketing Affects.

marketingaffects.com

