From *legacy GTM* to the *agentic GTM* universe

Every major function of go-to-market, what it was, what it becomes, and what has to be true underneath for the new version to actually work.

AUGUST 2026 25 PILLARS · 6 CLUSTERS

01 The frame: what actually changed

Most "AI is changing GTM" content lists tools. That's the wrong altitude and it dates in ninety days. Three structural things changed, and every function-level shift in this document is downstream of them.

Shift 01

The cost of a GTM action went to ~zero

The scarcity the whole operating model was built to ration simply isn't there any more.

scarce attention → **abundant action**

Shift 02

The bottleneck moved to judgment and trust

Anyone can execute now. Grounding, permission, and compounding are the scarce parts.

execution → **judgment**

Shift 03

The unit of work became the loop

Batch became continuous — but only where the outcome actually feeds back in.

campaign → **loop**

Shift one

The marginal cost of a go-to-market action went to ~zero

A researched account, a personalised sequence, a competitive brief, a call summary, a territory plan. Each of these used to cost a person a chunk of a day. Legacy GTM was, at its core, **an allocation system for scarce human attention.** Segmentation, lead scoring, MQL thresholds, territory carving, SLAs — every one of those exists to answer the question *"we can't do this for everyone, so who gets it?"*

LLM inference costs dropped more than 90% in roughly two years (Stanford HAI, 2025). When the constraint disappears, the machinery built to manage the constraint becomes overhead. **Half of the legacy GTM operating model is scarcity infrastructure that no longer has a scarcity to manage.**

Shift two

The bottleneck moved from execution to judgment and trust

If anyone can generate a thousand personalised emails, personalisation is no longer differentiation — it's table stakes, and shortly after that, it's noise. The scarce resources are now:

- **Grounding** — is the agent working from what's actually true about this account, this quarter, this product?
- **Judgment** — is this the right action, on the right account, at the right moment?
- **Permission** — is this action allowed, by whom, with what audit trail?
- **Compounding** — did we learn anything from the outcome, or did we just do it again?

None of those are solved by adding another agent. This is the part most teams have not internalised yet.

Shift three

The unit of work changed from the campaign to the loop

Legacy GTM was **batch**: plan the quarter, build the campaign, launch, wait, measure, retro, repeat. A ninety-day cycle with a review at the end.

Agentic GTM is **continuous**: signal arrives, agent acts, outcome returns, system updates, next signal is judged better than the last. The cycle time collapses from a quarter to minutes — but only if the outcome actually feeds back. Without a feedback path you don't have a loop, you have a faster batch. Most "AI GTM" deployments in market today are faster batches.

the shape of the shift ↓

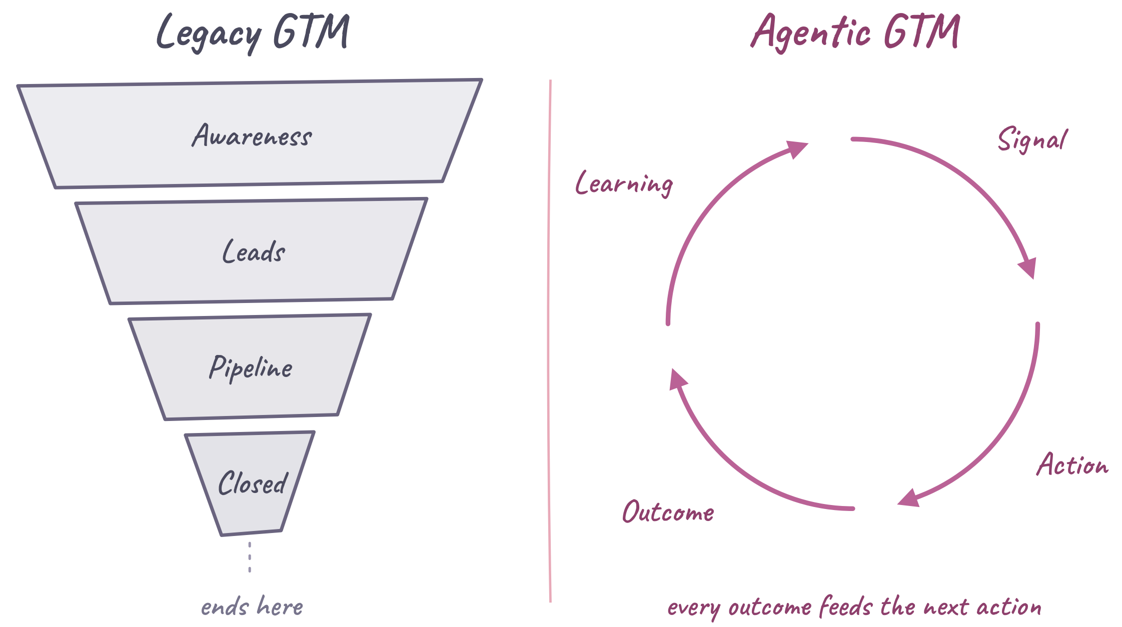

The one-line version: Legacy GTM optimised for *how much a team could execute.* Agentic GTM optimises for *how fast a system can learn.*

02 The twenty-five pillars, legacy vs. agentic

Grouped into six clusters that follow the actual lifecycle. For each pillar: what it was, what it becomes, and the shift in one phrase.

A KNOW

Understanding the market and the account

01 Market & category definition

artifact → living hypothesis

LEGACY	AGENTIC
Annual positioning exercise. Agency-led. A deck that ages the day it ships.	Continuously tested against live win/loss, call language, and buyer vocabulary. Positioning is a dataset, not a document.

02 ICP & segmentation

static filter → learned model

LEGACY	AGENTIC
Firmographic filter — industry, size, geo. Set once a year, encoded in list-pull criteria.	Behavioural and contextual. Defined by what closed accounts *did*, refined every deal. ICP is an output of the system, not an input to it.

03 Data & enrichment

record append → living entity

LEGACY	AGENTIC
Buy a vendor record. Append fields. It's stale on arrival and decays year over year. Multiple vendors, conflicting truth, no provenance.	Continuous synthesis across CRM, calls, email, product usage, and the public web into a persistent entity that has history. You don't enrich a record — you maintain an understanding.

04 Competitive intelligence

wiki page → ambient signal

LEGACY	AGENTIC
Battlecard in a wiki. Updated when someone loses a big deal and complains loudly.	Assembled from actual call transcripts, pricing seen in the field, and public signals — refreshed continuously, surfaced in the moment of need.

What this cluster requires underneath

- One shared representation of an account that every agent and every rep reads from — otherwise each agent enriches separately and you pay for the same retrieval a dozen times.
- **Provenance.** Where did this claim come from, when, how confident? Ungrounded enrichment at machine scale is just fabrication at machine scale.
- Memory that persists beyond a session, a thread, or a vendor contract.

B ATTRACT

Creating and capturing demand

05 Messaging & content

calendar → constrained generation

LEGACY	AGENTIC
Quarterly content calendar. Human-written, human-bottlenecked. One message per segment, at best.	Continuous generation constrained by an approved message model. Volume is free; *consistency* becomes the hard problem and the moat.

06 Discoverability (SEO → AEO/GEO)

ranking → being cited

LEGACY

Rank on page one of Google. Keywords, backlinks, domain authority. Optimise for a human who will click.

AGENTIC

Be the answer an AI gives when it's asked. Optimise for retrieval, citation, and structured extractability. A growing share of buyers never reach a results page.

07 Demand generation

channel spend → signal response

LEGACY

Channel-and-budget planning. Spend, wait, attribute, argue about attribution.

AGENTIC

Signal-triggered. Demand is detected in behaviour rather than purchased in aggregate, and budget follows observed intent in near-real-time.

08 Brand, trust & proof

marketing asset → product requirement

LEGACY

Logos, case studies, analyst placement. Trust as a marketing asset.

AGENTIC

Trust becomes a *buying requirement about your AI*, not just your company: how does your system handle data, what can it do autonomously, what's the audit trail? Procurement asks this now.

What this cluster requires underneath

- A canonical message model agents inherit — otherwise every agent invents its own version of what you do and the brand fragments in a hundred directions at once.
- Structured, machine-readable truth about the company. AEO/GEO is not a content tactic; it's a data-structure problem.
- Content that is *attributable to a source*, because the compliance surface of generated claims is real.

C ENGAGE

Reaching and qualifying buyers

09 Inbound

queue → conversation

LEGACY

Form → MQL → SLA → SDR call. A queue with a service-level promise, because humans were the scarce resource.

AGENTIC

Instant, conversational qualification at the moment of interest, with full account context already loaded. The form was always a rationing device.

10 Outbound

list + sequence → signal + response

LEGACY

Build list. Load sequence. Send volume. Measure reply rate. Personalisation was a variable field.

AGENTIC

Triggered by an observed event, not a list membership. Research, relevance, and timing are computed per account. Cadence is a consequence, not a plan.

11 Lead scoring & routing

point system → explained judgment

LEGACY

Points for a whitepaper download. A model built once, tuned rarely, trusted by nobody.

AGENTIC

Continuous scoring on behavioural and contextual evidence, with the *reasoning* exposed so a human can overrule it — and the overrule teaches the model.

12 Sales development (SDR/BDR)

volume execution → signal judgment

LEGACY

Volume role. Dial and email targets. Ramp a human for six months to send messages they didn't write.

AGENTIC

The role compresses hard. What survives is judgment: which signals matter, which agent output ships, which account deserves a human. This is the most disrupted job in GTM — say so plainly.

What this cluster requires underneath

- Governance on outbound action. When sending costs nothing, the only thing standing between you and burning your domain, your brand, and your TAM is an explicit policy layer.
- Human-in-the-loop on anything that writes to a system of record or reaches a prospect — with approvals captured as training signal, not just as a gate.
- Shared context, so the prospect doesn't get three different agents asking the same discovery question.

D CONVERT

Running the deal

13 Discovery & qualification

self-reported → observed

LEGACY

MEDDIC/BANT in a rep's head, half-entered into CRM the Friday before forecast call.

AGENTIC

Extracted automatically from the actual conversation. Gaps flagged in the moment, not discovered at close. Qualification becomes an observed state rather than a claimed one.

14 Deal execution & multithreading

rep discipline → system monitoring

LEGACY

Rep-dependent. Whoever's diligent wins. Champion leaves, deal dies, nobody sees it coming.

AGENTIC

Continuously monitored — engagement breadth, stakeholder coverage, sentiment drift, silence. Risk surfaces with the specific next action attached.

15 Proposals, pricing & deal desk

approval chain → encoded policy

LEGACY

Approval chains and spreadsheets. Days of latency to answer "can I do 18%?"

AGENTIC

Policy encoded once and applied instantly, with exceptions escalated by rule. The bottleneck was never the decision — it was the routing.

16 Enablement & ramp

training event → in-the-moment grounding

LEGACY

Two-week bootcamp, a content library nobody opens, six-month ramp to quota.

AGENTIC

Assistance at the point of need, grounded in what actually wins here. Ramp becomes a competence curve, not a calendar.

What this cluster requires underneath

- Grounding in *your* deal reality — a generic model gives generic advice, which is worse than none because it's confidently wrong.
- Write-approval controls. An agent updating opportunity stages without governance corrupts the forecast, and a corrupted forecast is worse than no forecast.
- Outcome linkage: which suggestion actually preceded a stage advance? Without that, enablement stays a faith-based function forever.

E RETAIN & EXPAND

After the close

17 Onboarding & handoff

handoff → continuity

LEGACY

Slide deck thrown over the wall. Customer re-explains everything they already told sales.

AGENTIC

Continuous context. The account's history travels with it — no re-discovery, no cold restart.

18 Customer success & health

lagging score → leading signal

LEGACY	AGENTIC
Health score built from logins and NPS. Lagging, coarse, quietly ignored.	Composite signal from usage, support, sentiment, and stakeholder change. Leading rather than lagging, with a specific intervention attached.

19 Renewal & expansion

date-driven → signal-driven

LEGACY	AGENTIC
Calendar-driven. Start ninety days out and hope.	Continuous. Expansion signals surface when they occur, renewal risk surfaces when it emerges — typically two quarters before the date on the contract.

20 Partner, channel & advocacy

portal → shared context

LEGACY	AGENTIC
Portal, deal registration, quarterly enablement webinar.	Partners plug into the same context and the same guardrails. Advocacy is identified from behaviour rather than solicited by email blast.

What this cluster requires underneath

- Memory that survives the handoff between teams and systems — the single most common failure point in the entire customer lifecycle, and the one AI is best placed to fix.
- Attribution to outcomes, not activity, so you can tell interventions that worked from interventions that merely happened.

F RUN THE BUSINESS

The operating discipline — where legacy GTM breaks hardest, and where almost nobody is looking

21 Forecasting & pipeline

self-reported → evidence-based

LEGACY	AGENTIC
Rep commit + manager gut + a spreadsheet. Accuracy defended, rarely measured.	Evidence-based, from observed deal behaviour rather than self-reported confidence. The forecast becomes falsifiable.

22 Attribution & measurement

channel credit → outcome attribution

LEGACY

Multi-touch attribution models everyone distrusts and argues about at QBR.

AGENTIC

The measurement question inverts: not *which channel got credit*, but *which agents and actions actually moved deals* — and generic AI observability (tokens, latency, cost) does not answer it.

23 Governance, risk & audit

didn't exist → first-class discipline

LEGACY

Not a GTM function. Legal's problem. A CRM permission matrix nobody reviewed.

AGENTIC

A first-class GTM discipline. Who can act, on what, with whose approval, with what audit trail. This is the single biggest net-new function in agentic GTM.

24 The stack itself

point-to-point → shared layer

LEGACY

Buy best-of-breed per function. Integrate point-to-point. Accept the sprawl.

AGENTIC

The integration surface becomes the problem: the same connections wired separately into every agent means paying for the same retrieval a dozen times, with a dozen versions of the truth. Connection consolidates.

25 Team design & comp

headcount-scaled → compute-scaled

LEGACY

Headcount scales with pipeline. Comp rewards activity and quota attainment.

AGENTIC

Headcount decouples from volume. The org rebalances toward judgment, orchestration, and trust. Comp design for a team whose leverage is a system, not their calendar, is a genuinely unsolved problem.

What this cluster requires underneath

- Per-agent permissions, human-in-the-loop, write-approval queues, and a complete audit trail — treated as infrastructure, not as a policy PDF.
- Outcome attribution tied to deal stages, not token counts.
- One connection layer rather than N×M integrations.

The number worth quoting here: Gartner projects **40% of agentic AI projects will be cancelled by the end of 2027** (Gartner, 2025) — largely on cost and governance, not on model capability. The models work. The stack around them doesn't.

03 The eight shifts that cut across all twenty-five

If the audience remembers nothing else, this is the slide. Every pillar above is an instance of one of these.

01 Campaigns → **Continuous loops**
Cycle time collapses from a quarter to minutes — *if* outcomes feed back

02 Lists → **Signals**
Membership in a segment is a weak reason to act; an observed event is a strong one

03 Scarcity rationing → **Abundance governance**
The job changes from *deciding who gets attention* to *deciding what's allowed*

04 Activity records → **Outcome memory**
CRM logged what happened; the new system has to know what worked

05 Funnel → **Loop**
Stages become states you can enter, exit, and re-enter — measured continuously

06 Tools per function → **A layer beneath all functions**
Twenty-five agents with twenty-five memories compound nothing

07 Managing people → **Governing agents**
A wholly new management discipline, and nobody has a playbook yet

08 Headcount leverage → **System leverage**
Growth decouples from hiring — which breaks planning, comp, and career paths

The through-line: every one of those requires the same four things underneath — **shared grounding, governed action, reusable capability, and an outcome feedback path.** Twenty-five functions, one substrate. That's the whole argument.

04 The diagnostic: three questions

Useful in the room. Every GTM leader can answer these about their own org in thirty seconds, and most don't like their answers.

If you added a new agent tomorrow, would it know everything your existing agents know?

If no — your memory is per-agent, and nothing you deploy compounds. You are buying tools, not building capability.

Can you name which of your AI investments moved a deal stage last quarter?

If no — you have observability of cost, not of outcome. You cannot manage what you cannot attribute, and you will not defend the budget in the next cycle.

Who approved the last thing an agent wrote to your CRM?

If nobody knows — you don't have governance, you have exposure. And this is the question your CFO asks at renewal.

Three no's is the normal answer today. That's the point — this isn't a maturity failure, it's a missing layer.

05 Where this leads — and the honest answer on mapping

On whether the GTM pillars should map one-to-one to wysdym's five pillars: deliberately don't — and make that the payoff of the piece.

A one-to-one map would be a mistake for three reasons:

- **It's not true.** Cortex isn't "the enrichment pillar." It serves all twenty-five. Forcing the map breaks the actual architecture story, which is that one layer sits beneath every function.
- **It converts education into a pitch,** and the room will feel the turn. The credibility of this piece comes from the first four sections being pure market analysis with no product in them.

- **The grid is the stronger visual anyway.** GTM functions across the top, the operating layer underneath all of them. The map is a *dependency*, not a correspondence — and dependency is a far better thing to own than a feature match.

So the relationship to state is:

> Twenty-five GTM functions are all changing. Every one of them needs the same four things underneath: **grounding** in what's true, **capability** that's reusable, **governance** on every action, and **feedback** from every outcome. Those four aren't features of any agent. They're the layer the agents run on.

That's where **wysdym** sits: the operating layer for agentic go-to-market. Cortex, Skills, Governance, Observe, and Operator — connected through the Gateway, one MCP door to your stack. It grounds every agent in your GTM truth, runs the motion under your governance, and gets sharper with every deal you close.

Stated once, at the end, after the audience has already drawn the shape. That's the whole play.

The one-page version of this argument — six clusters, the legacy/agentic split, and the operating layer running beneath all of it — is here: The agentic GTM universe, one page ↗

---

wysdym · THE OPERATING LAYER FOR AGENTIC GO-TO-MARKET AUGUST 2026