AccountMade

Sales has become the default motion

earlier than the playbook admits

A corpus-only study of 201 companies, each scored on ten go-to-market motions across six funding stages and twelve categories.

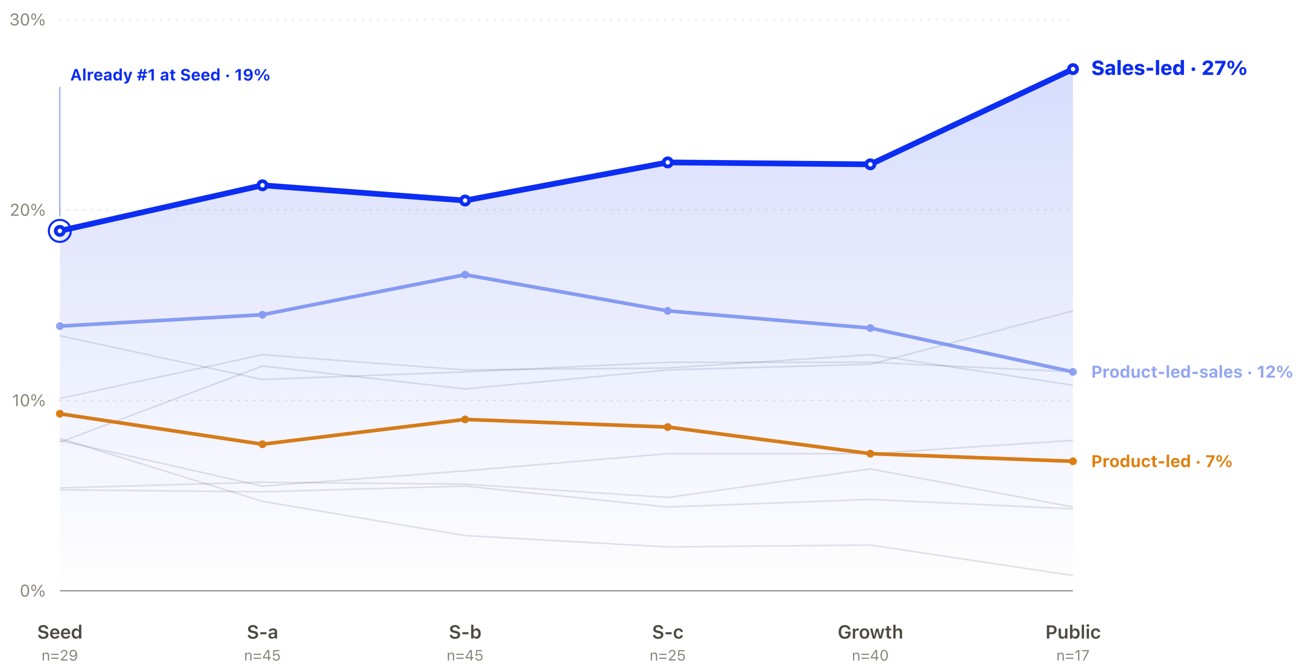

Jake Jinyong Kim

FIGURE 1 · MEAN MOTION WEIGHT BY FUNDING STAGE · N = 201

Contents

01	The playbook is wrong about motion	the story vs the record
02	How we mapped it	method · selection · verification
03	What the corpus shows	Figure 1 · the tightening
04	The exception map	Figure 2 · category × motion
05	How they actually reach buyers	Figure 4 · the channel stack
06	How they grow, and how they charge	Figure 5 · expansion + billing
07	What works — and what doesn't	Figure 3 · outcomes + efficacy
08	What's changing	Figure 6 · vintage cohorts + AI
09	The twelve machines, in one view	Part II · Figure 8
10	Deep dive · Security & compliance	why trust leads early
11	Deep dive · Developer tools	why product leads
12	Inside the labels	Part III · Figure 7 · motion anatomy
13	Three that break the rule	Part IV · deep specimens
14	Six laws	Part V · synthesis
15	Limits, and the road to conclusive	honesty · what's next

1

What holds across all 201 companies

Eight chapters on the corpus-wide record — the findings that survive every caveat, before we go sector by sector. How motion moves by round, how companies actually reach buyers, how they grow and charge, what works and what doesn't, and what changed in the last year.

01	The playbook is wrong about motion	the story vs the record
02	How we mapped it	method · selection · verification
03	What the corpus shows	Figure 1 · the tightening
04	The exception map	Figure 2 · category × motion
05	How they actually reach buyers	Figure 4 · the channel stack
06	How they grow, and how they charge	Figure 5 · expansion + billing
07	What works — and what doesn't	Figure 3 · outcomes + efficacy
08	What's changing	Figure 6 · vintage cohorts + AI

1 The playbook is wrong about motion

The tidy story goes like this: a young software company starts product-led, lets users self-serve, bolts on sales-assist around Series B or C, and becomes a full hybrid once enterprise procurement shows up. A clean sequence, one stage at a time.

The record does not support it as a general rule. Across 201 companies scored on ten motions, the centre of gravity is already **sales-led in five of six stage bands — including Seed**. By the time a company leaves enough public evidence to study, it is already carrying a sales motion. The pure product-led phase, if it existed, happened off-camera.

> **The transition isn't product-led growth giving way to sales. It's a hybrid stack that gets tighter, and more legible, each round.**

This is not the death of product

Product-led-*sales* is the second-strongest stage-weighted motion at every stage from Seed through Growth. Product still matters enormously — as a wedge, a credibility device, an expansion surface. It is simply not, in most companies, the centre of gravity. **Pure product-led leads for only 8%** of the corpus (16 of 201). The other 92% put a human in the loop.

8% lead with pure product-led (16 of 201)	**152** of 201 are pure hybrids (top motion < 0.30)
0.26 mean weight of a company's strongest motion	**7** companies run essentially one motion

2 How we mapped it

Most GTM commentary ranks a leaderboard of the biggest companies and reasons backward. That corpus is ~80% late-stage and public, and it has converged on the same hybrid — so the picture comes out muddy and non-transferable. We did the opposite: we designed a sample to fill a matrix.

One reframe: fill a matrix, don't rank a list

The analytical unit is a **(stage × category) cell**. Companies were selected to populate those cells, weighted toward the ICP — Seed-to-Series-B founder-sellers and their first few reps — not chosen for size. **Stage is the primary axis**, because motion is genuinely stage-dependent: validation at Seed, building a repeatable channel at Series A, proving it across quarters at Series B, optimizing at Growth. The insight the corpus was built to expose is the *transition* — at which round, in which category, the centre of gravity shifts, and what triggers it.

Four strata on every company

Each company is stamped on four axes: **stage** (seed → public), **category** (twelve, from devtools to regulated health), **ACV band** × **buyer** (self-serve developer through enterprise InfoSec), and an **outcome signal** — winner, steady, or stalled/down-round. Motion itself is scored as a ten-way vector summing to ~1, over company-neutral archetypes (product-led, product-led-sales, sales-led, marketing-led, community/developer-led, partner-ecosystem-led, procurement/trust-led, founder-network-led, usage-led, embedded-API-led).

Anti-survivorship, by design

The single least-fakeable decision in the study: **15–20% of every stage band is a stalled or down-round company** — 27 of 201 overall (13.4%). Without a failure cohort you can only learn what winners did; you can never say what *doesn't* work, which is half the point.

A study that only looks at winners can describe success; it cannot diagnose failure. So the failures were sourced on purpose.

The research pipeline

Every company was first sourced from public, stage-tagged lists as a *hypothesis* — a provisional stage and outcome. Then each ran through a fixed pipeline. **Research** rebuilt the record against an evidence-required schema: every load-bearing claim carries a dated public source or is moved to a named list of gaps — stage is *verified* from evidence, not assumed, and unknowns are declared rather than guessed. **Adversarial verification** then took a second pass whose only job was to refute the thesis; where it found the record over-claimed, a repair step corrected it — that pass caught, among others, a company filed as "Series B" that had in fact already been acquired. The output is a structured record per company, aggregated into cells.

What we did not do

No number in this report comes from free web search or an unstated benchmark — every figure traces to the aggregated corpus, which is what makes it a primary source. And efficacy here is an **observed-pattern signal, not a controlled experiment**: it records which motions concentrate among stronger outcomes, not that a motion *caused* an outcome.

METHOD & BIAS, NAMED

Corpus. 201 structured research units across 67 (stage × category) cells. Motion is a 10-way vector per company; cell figures are weighted centroids.

Verification is uneven. ~129 units passed adversarial verification; ~72 are research-grade (single-pass, evidence-required but not adversarially refuted). Treat single-cell claims accordingly.

Known biases. Survivorship (public evidence is richer for bigger, louder companies — stalled firms are included but quieter by nature); public-evidence-only selection; and the uneven verification above. All three are named; none are corrected away.

3 What the corpus shows

N=201

Sales-led is the weighted top motion in five of six stage bands, and it climbs steadily — from **18.9% of the Seed motion centroid** (N=29) to **27.3% at Public** (N=17). Product-led-sales holds the second tier throughout, peaking at 16.6% in Series B.

FIGURE 1 N=201 · WEIGHTED MOTION CENTROID

The motion stack tightens toward sales, round by round

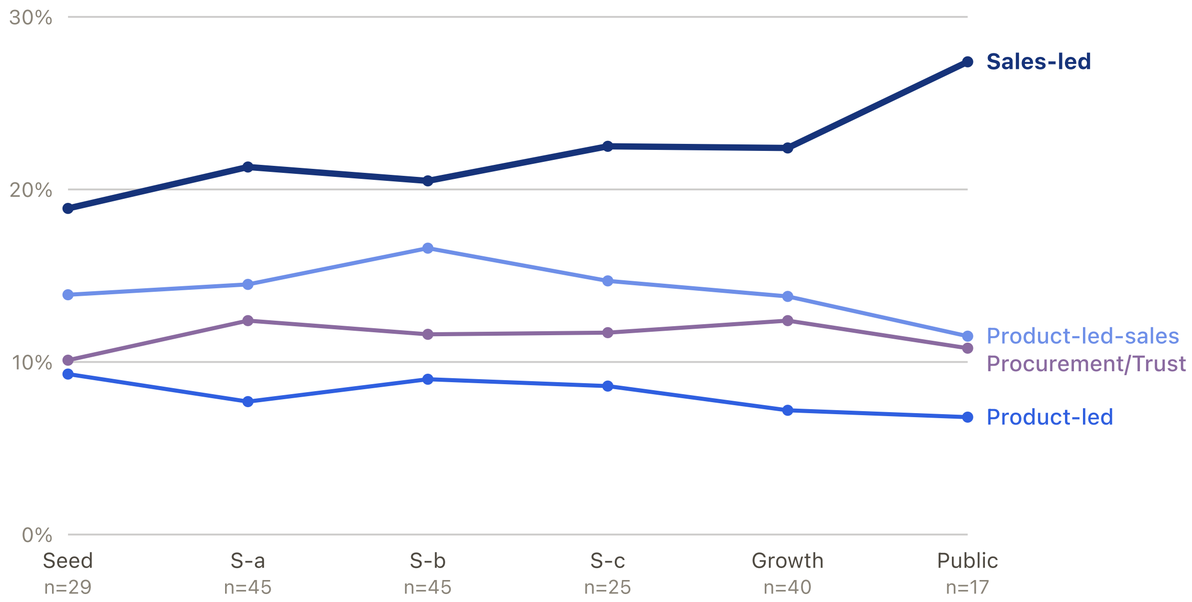

Sales-led leads at every stage and rises into the public band; product-led-sales holds second; pure product-led stays a minority throughout. The stack broadens with scale but never flattens.

weighted mean motion weight per stage

A tightening, not a pivot

Seed companies are not mostly self-serve experiments. Across 29 Seed units, sales-led is already the top motion (18.9%), ahead of product-led-sales (13.9%) and marketing-led (13.4%). Series A and B don't *introduce* sales so much as make it legible: the conspicuous change is product-led-sales climbing 14.5% → 16.6%, a product surface increasingly used to qualify, route, and support selling. By Growth the stack is broad but still tilted to sales (22.4%); at Public, partner-ecosystem-led rises to become the second motion (14.7%).

Nobody runs a single motion. The average company's strongest play carries just a quarter of the weight — the other three-quarters is the strategy.

4 The exception map

N=201 ACROSS 12 CATEGORIES

The average says sales-led. The category map says where that average hides useful timing. Most lanes present as sales-led early with adjacent support; a few show a genuine round-by-round handoff worth naming.

FIGURE 2 12 CATEGORIES × 10 MOTIONS

Where each motion concentrates, by category

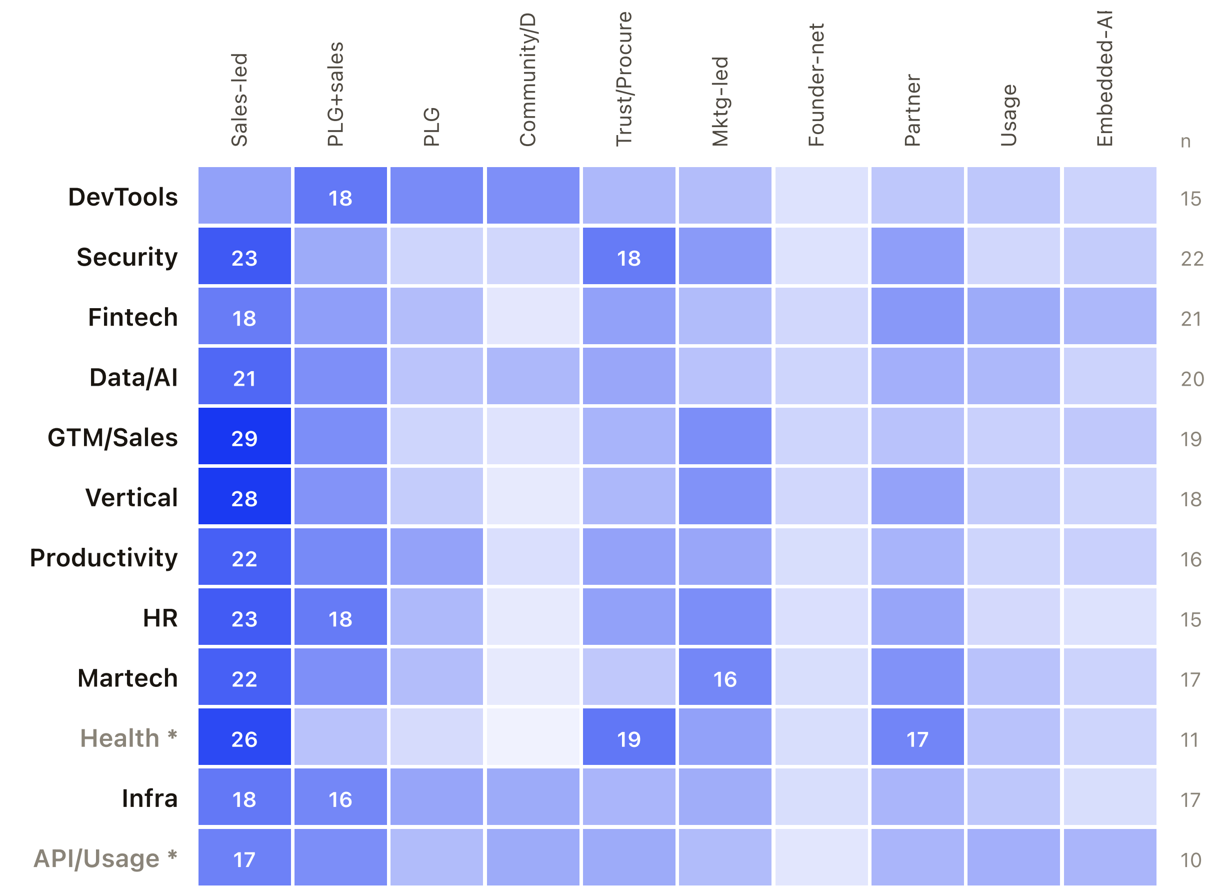

*Cell shade = share of the category's weighted motion centroid. Devtools is the one clearly product-leaning row; security/compliance and regulated health light up procurement/trust; GTM-sales-tech and vertical SaaS are the most sales-heavy. Thin categories (n<12, marked *) are directional only.*

motion centroid ×100 per category

Devtools is the clean transition

If a product-to-sales handoff exists anywhere in the corpus, it lives in devtools (N=15): product-led at Seed, moving to product-led-sales by Series A and B. Its category centroid leads with product-led-sales (18.3%), then product-led (15.5%) and community/developer-led (14.9%) — the only row where the developer motions out-weigh sales.

Trust shows up early where the buyer is InfoSec

Security and compliance companies carry a sales-plus-trust posture from Series A onward — procurement/trust-led runs well above its corpus baseline. When the buyer is InfoSec, the trust envelope is not a late-stage add-on; it is the entry motion.

Fintech flips fast

Fintech reads product-led at Seed but moves to sales-led by Series A and stays there. The corpus can't name the ACV threshold that triggers it — it doesn't expose per-cell ACV — but the flip is early and durable.

The one buyer that bends the motion

Across every stratum, one variable moves the motion more than stage or category: **who signs**. When a developer sits in the buying group (78 companies), sales-led falls to **0.17** and product-led rises — against **0.25** sales-led for everyone else. No other buyer type discriminates, because nearly every company lists practitioner, department-lead, exec, and InfoSec together. The developer is the exception that proves the rule: product-led is a *buyer* property, not a stage.

5 How they actually reach buyers

1,026 CHANNEL OBSERVATIONS

Motion vectors are an abstraction. Underneath them sit concrete channels — and when you count only the *load-bearing* ones, rated dominant or strong, the picture sharpens into something a founder can act on.

FIGURE 4 DOMINANT + STRONG CHANNEL MENTIONS

The channels that actually carry the motion

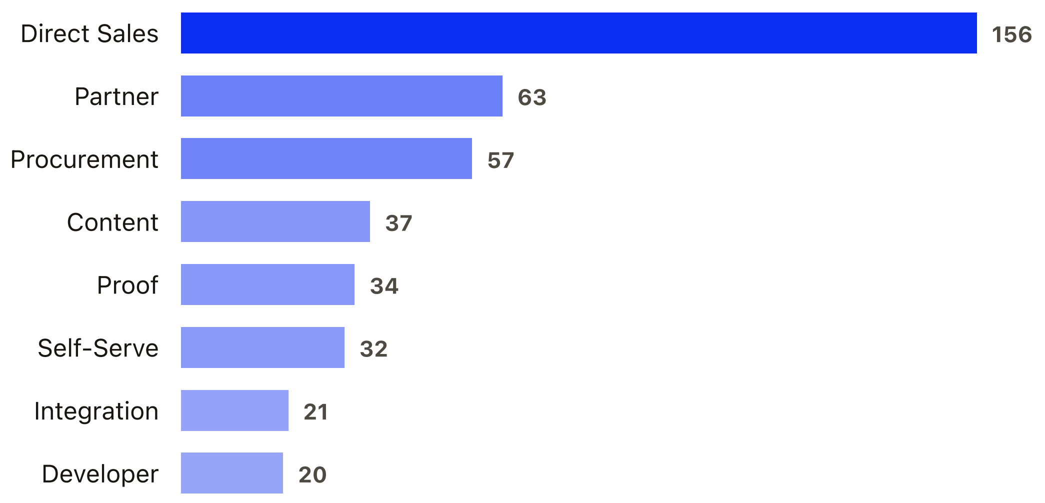

Channels rated dominant or strong across 201 companies. Direct sales leads by a wide margin; partner ecosystems and a procurement/trust channel form the enterprise-access layer; content and customer proof are the second tier; self-serve and developer channels are real but a distinct minority.

channel observations · intensity: dominant or strong

Direct sales carries the deck; self-serve is a minority

The corpus does not, on the whole, reach buyers by letting them sign up. Direct enterprise sales is the single most load-bearing channel, and the two that follow — partners and a distinct procurement/trust channel — are the mechanics of getting through a security review, not of demand generation. Self-serve exists, but as a wedge feeding a sales motion far more often than as the motion itself.

The dominant channel in modern B2B isn't a signup page — it's a security questionnaire, answered well.

6 How they grow, and how they charge

N=201

Two mechanisms decide whether a landed account becomes a franchise: how it expands, and how it is billed. On both, the corpus has moved off the classic seat-based subscription.

FIGURE 5 EXPANSION + MONETIZATION MECHANICS

Growth and billing have moved off the seat

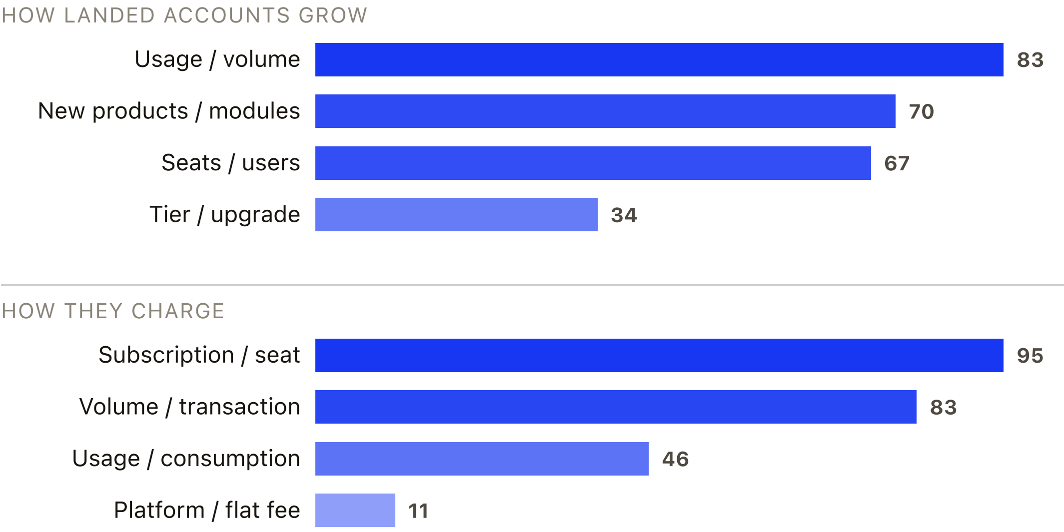

Top: how landed accounts expand. Bottom: how companies charge. Usage/volume marginally leads expansion, ahead of new modules and additional seats; and in billing, volume/transaction and usage pricing together rival the per-seat subscription.

expansion and billing mechanisms

Growth is usage as much as seats

Expansion splits three ways — usage and volume, new products and modules, and additional seats — with usage/volume just ahead. Billing tells the same story: the pure per-seat subscription barely leads volume- and transaction-based pricing, and usage/consumption is a substantial third. For a large share of the corpus, revenue now grows with the customer's *usage*, not their headcount — which quietly changes what "expansion" even means.

7 What works — and what doesn't

N=201

Outcomes are not evenly distributed across stages. The corpus has a clear proving ground and a clear graveyard.

FIGURE 3 OUTCOME MIX BY STAGE

Series B is where the motion proves out; public is where it derates

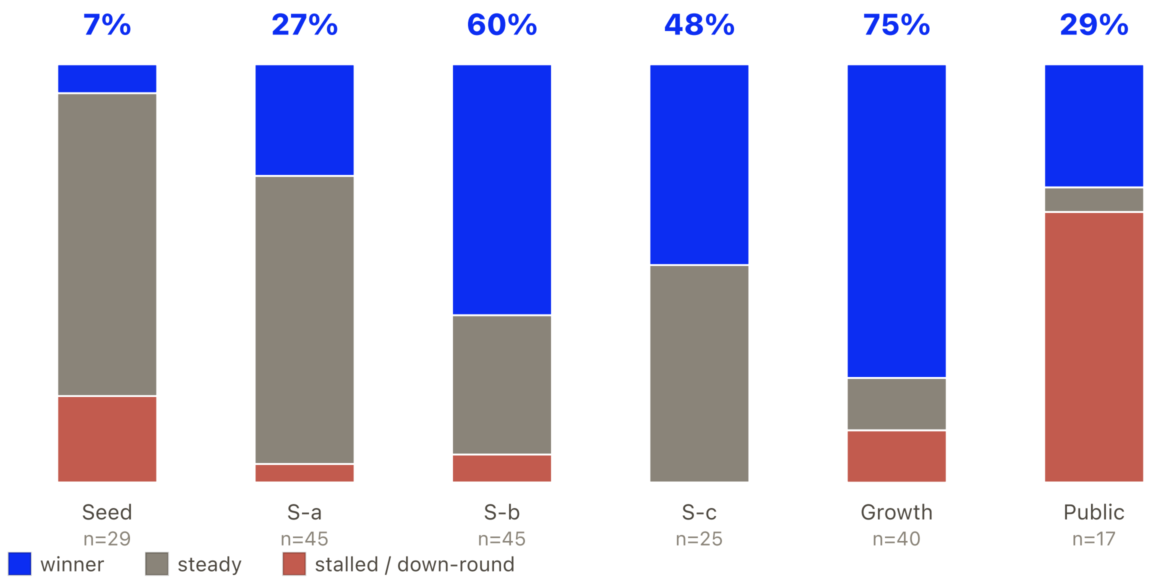

Win-rate (number above each bar) jumps from 27% at Series A to 60% at Series B and peaks at 75% in Growth, then falls to 29% at Public — where 65% of the observed cohort is stalled or down from peak. Seed is almost all "steady" because a Seed outcome is not yet knowable; the study does not fabricate one.

outcome by stage · ~18% stalled, seeded by design

Motion does not predict the outcome

Here is the corpus's most useful negative result. Compare the average motion of the 88 winners with the 27 stalled companies and they are, to the decimal, the same posture — sales-led around 0.22, product-led-sales around 0.15. You cannot read a company's fate from its motion mix. "**Which motion should we run" is the wrong question**: the winners and the failures ran the same ones.

The faint signal that does exist points one way, and it is worth naming. Winners carry slightly more procurement/trust weight (0.12 vs 0.10); the stalled cohort carries slightly more marketing-led (0.13 vs 0.12). On the margin, trust and proof beat volume and noise. It is a small effect on thin numbers — but it is the only directional tilt the outcome data offers, and it agrees with everything else in the corpus.

Winners and failures ran the same motions. What separated them was fit and timing — not the play they picked.

What the stalled cohort shares

The 27 stalled or down-round companies are not concentrated in a single motion — you cannot blame "too much PLG" or "too much sales." What they share more often is a **stage-mismatch**: a motion running ahead of, or behind, the trust its buyer required — self-serve machinery aimed at a procurement buyer, or an expensive sales motion with no proof to hand the reviewer. Motion without enough trust for the buyer is the pattern that recurs on the failure side.

> **Failure in the corpus rarely looks like the wrong motion. It looks like the right motion, mistimed for the buyer in the room.**

8 What's changing

FOUNDED-YEAR ON 189 OF 201 · 179 DOCUMENTED SHIFTS

A snapshot can't watch a company change over time. But it can sort companies by when they were founded, and read the change events each has put on record in roughly the last year. Both point the same way.

The pure-sales playbook is aging out

Sort the corpus by founding era and the motion drifts, cleanly. The oldest cohort — founded before 2015 — leads sales-led at 0.25, with partner-ecosystem second. The newest — 2022 and later — is down to 0.20 sales-led, with product-led-sales risen from 0.12 to 0.16 and founder-led motion up with it. Newer companies still put a human in the loop, but the human is increasingly a *founder plus a product surface*, not a partner channel and a full sales org.

FIGURE 6 MOTION CENTROID BY FOUNDING ERA

Newer companies lean off sales and partners, toward product and founder

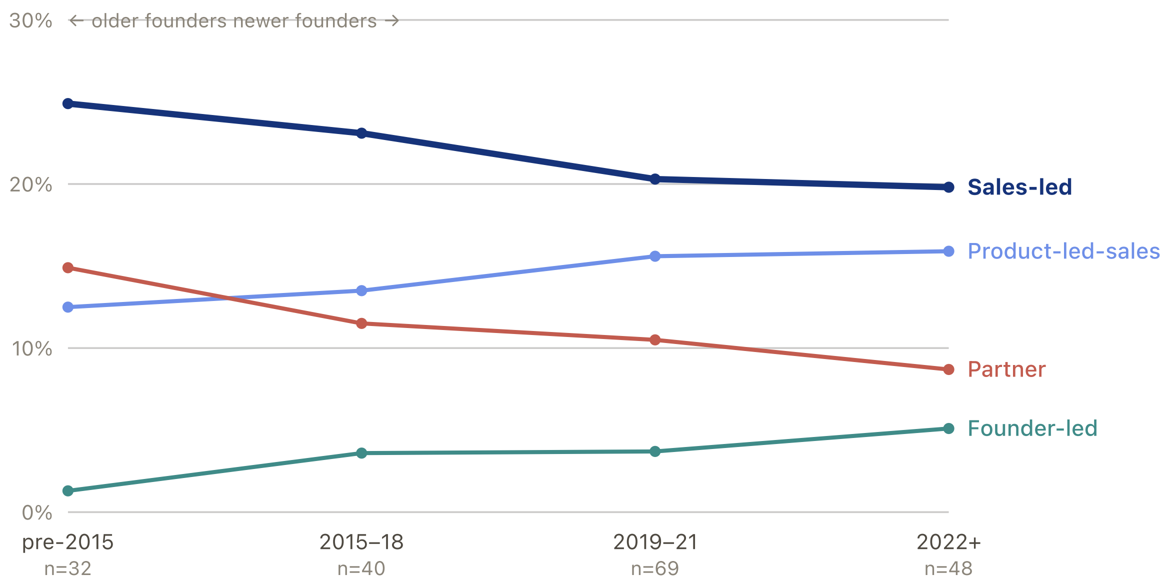

Weighted motion by founding cohort, oldest to newest. Sales-led and partner-ecosystem fall with each cohort while product-led-sales and founder-led rise — the clearest quasi-temporal signal a snapshot can offer.

founded-year enrichment · 189 of 201 with a verified year

The last year's move was AI, not a new motion

Read the 179 documented go-to-market changes from the last twelve months and one thing dominates: companies expanded their *product surface*, and did it with AI. 61% shipped a new product or platform; 42% launched or repositioned specifically around AI. Far fewer touched the sales machine — 16% moved upmarket or added sales, 15% added partnerships, and only a handful changed pricing or opened self-serve. The headline change of the year is not a new motion; it is AI-shaped product expansion bolted onto the motion a company already had.

61% shipped a new product or platform	**42%** launched or repositioned around AI
16% moved upmarket or added sales	**5%** opened self-serve / product-led

The motion of the last year didn't change. The product did — and it grew an AI surface.

A trap worth naming

It is tempting to read outcomes by vintage and conclude newer companies win more: 2021-and-later founders show 1% stalled against 22% for pre-2021. That is survivorship, not success — a company founded in 2023 has not had time to fail or take a down-round, while the older cohort has. Read the vintage trend for what companies *do*, never for how they *end*. On the change data, the outcome clock simply hasn't run yet.

2

Twelve categories, twelve different machines

The corpus-wide average is a useful lie. Underneath it, each sector runs a recognisably different motion — and the two most instructive, security and devtools, run near-opposite ones for the same reason: their buyer. This part maps all twelve at a glance, then goes deep on those two.

09	The twelve machines, in one view	Figure 8 · every sector at a glance
10	Deep dive · Security & compliance	why trust leads early
11	Deep dive · Developer tools	why product leads

9 The twelve machines, in one view

12 CATEGORIES · N = 10–22 EACH

Every sector's motion mix on one axis, ordered from the most sales-led to the most product-led. The spread is the point: "software company" hides a factory floor of different machines, and two of them barely resemble each other.

FIGURE 8 WEIGHTED MOTION MIX PER CATEGORY

From most sales-led to most product-led

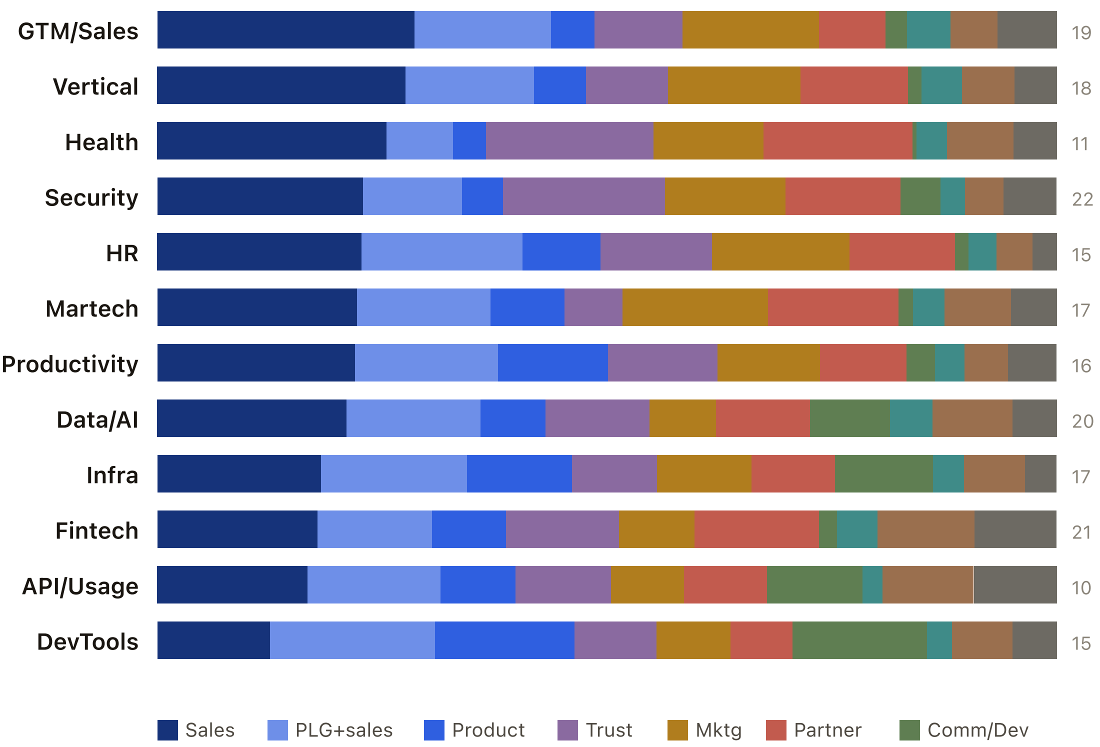

Each bar is one category's weighted motion centroid, stacked left-to-right by motion. GTM/sales-tech and vertical SaaS sit at the sales-led extreme; devtools is the lone product-led outlier; security and regulated health carry a visible trust band (purple) that no other sector shows.

motion centroid per category · sorted by sales-led weight

Two sectors are worth stopping on — not because they are the biggest, but because they are the clearest natural experiment the corpus contains. Security and devtools sell to opposite buyers, and run opposite motions as a direct result.

10 Deep dive · Security & compliance

N=22 · INFOSEC BUYER ON 100% OF COMPANIES

Security is the cleanest proof in the corpus that motion follows the buyer. Every one of the 22 security companies sells, among others, to InfoSec and procurement — the only sector where that buyer is universal. When the person who signs the contract is the same person paid to run the security review, the trust envelope stops being a late checkbox and becomes the entry motion.

100% list InfoSec / procurement as a buyer — the only sector at 100%	**#2** trust/procurement is the second motion from Seed onward (0.18)
0.22 trust weight already at Seed — vs ~0.10 corpus-wide	**0.33** sales-led at Public, as the trust-gated machine industrialises

Trust is the ante, not the upsell

Sales-led still leads the sector's centroid (0.23), but procurement/trust rides second (0.18) — and, unlike every other sector, it is there from the *start*: 0.22 at Seed, 0.19 at Series A, where elsewhere trust barely registers before Growth. The procurement channel is the second-most load-bearing in the sector (12 of 22 companies), behind only direct sales. This is a market reached through a security questionnaire, not a signup page — and the questionnaire is answered before the deal, not after.

The stage arc: trust holds, sales industrialises around it

Across rounds, trust stays high and early while the sales machine assembles around it. By the public stage, sales-led jumps to 0.33 and partners rise to 0.17 as the motion scales — but trust does not *grow* into the picture the way it does elsewhere, because it was never absent. It is the ante paid from the first deal.

What the winners are actually doing

The exemplars decode the motion precisely. **ConductorOne** does not sell access reviews; it reprices identity as "governed work," and the security review is what qualifies the buyer in the first place. **Veza** sells an authorization graph as the missing security *data layer* — trust rendered as infrastructure, which is why ServiceNow wanted its identity context inside AI controls. **Sublime Security** runs a practitioner PLG wedge — free analyzer, open docs, community — whose real job is to *prove transparency* against black-box incumbents, then converts into an enterprise replacement sale. In all three, the free or product surface exists to establish trust; the revenue lives in the governed enterprise motion.

In security, "prove it" comes before "buy it." The motion is trust-first and sales-closed — and it is structural, dictated by a buyer who is paid to be skeptical.

11 Deep dive · Developer tools

N=15 · DEVELOPER BUYER ON 100% OF COMPANIES

If security proves motion follows the buyer, devtools proves it from the opposite pole. It is the single sector in the corpus where sales-led is *not* the top motion — it sits fourth (0.12), behind product-led-sales (0.18), product-led (0.15), and community/developer-led (0.15). The reason is security's reason, inverted: every devtools company sells to a developer, and the developer is both the user and the first buyer. Adoption is bottom-up, hands-on, and usually free; the org purchase is a consequence, not the entry.

4th where sales-led ranks here — the only sector it isn't #1	**8/15** lean on a developer channel — uniquely load-bearing
0.21→0.20 product-led at Seed hands off to sales-led by Public	**80%** carry a "mixed" ACV — a self-serve bottom under an enterprise top

The one place the classic playbook actually holds

Devtools is where you can watch the textbook transition happen. Product-led leads at Seed (0.21); product-led-sales takes over at Series A and B (0.21, 0.20); sales-led finally leads at the public stage (0.20). Round by round, a product motion hands off to a sales motion — the exact arc the rest of the corpus refuses to show. The developer channel is uniquely central here (8 of 15, second only to direct sales), and ACV is overwhelmingly "mixed" — the signature of a free self-serve base feeding an enterprise ceiling.

Why it is product-led-*sales*, never pure PLG

Even at its most bottom-up, devtools converges on a sales motion — because the enterprise upgrade forces it. **GitLab** is "easy to misread as open-source PLG, but the economic center is enterprise standardization: the free surface creates gravity, the money concentrates in governance-heavy Ultimate." **CodeRabbit** installs free at the pull request, but sells *neutrality* — the independent, agent-agnostic quality gate an enterprise can mandate. **Momentic** makes engineers own quality inside the repo, then expands on CI volume and organisational control. The free developer surface is the wedge; the compliance-and-governance layer is the business.

Product-led isn't a strategy you choose. It is what a developer buyer allows — and even then, it converges on a sales motion the moment the org, not the individual, has to sign.

3

"Sales-led" is not one thing

A ten-way motion vector is a useful abstraction and a lossy one. A single label — "sales-led" — hides five or six genuinely different machines. This part opens the biggest labels and names what is inside them, from the companies' own channel evidence.

12 Inside the labels Figure 7 · motion anatomy

12 Inside the labels

5,646 DOCUMENTED CHANNEL MECHANISMS

Take the 111 companies whose dominant motion is sales-led. Their *channels* tell you what that sales motion actually is — and it is overwhelmingly relationship- and proof-driven, almost never cold.

FIGURE 7

N=111 SALES-LED COMPANIES

What "sales-led" is made of

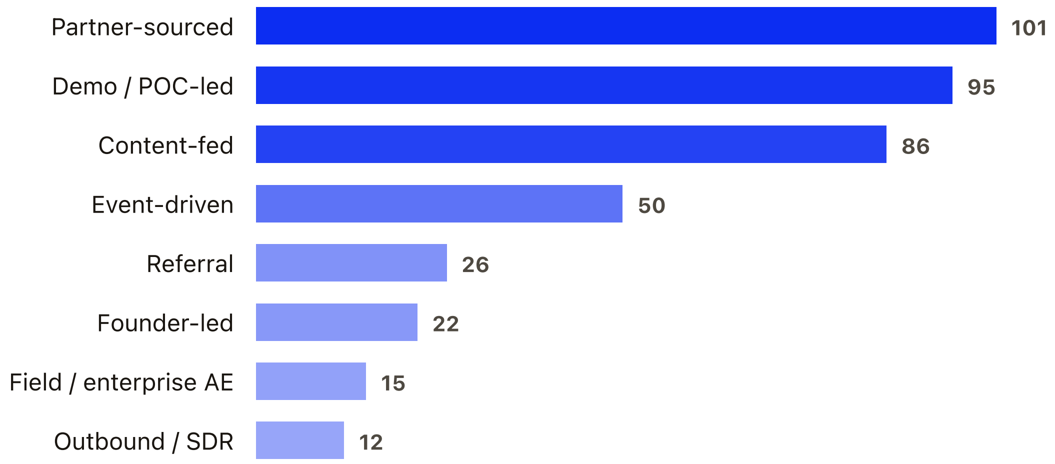

Companies whose top motion is sales-led (n=111), counted by the sub-channels they actually operate. Partner-sourced, demo/POC-led, and content-fed dominate; cold outbound is the rarest. "Sales-led" here means warm, partnered, and proven — not a bank of SDRs dialing.

channel operating motion and specific mechanisms

Sales-led, decomposed

Sales-led

the umbrella over 111 companies · six machines underneath

Partner-sourced. The most common by far. Resellers, SIs, marketplaces, and ecosystem referrals open the door; the AE closes what someone else warmed. most prevalent

Demo- and POC-led. A guided demo or a time-boxed proof-of-concept is the conversion event. The product does the persuading; a human orchestrates the trial.

Content-and-proof-fed. Benchmarks, case studies, and category content generate inbound intent that routes to a rep. Marketing warms, sales closes.

Event-driven. Conferences, roadshows, and field events manufacture the pipeline — heaviest in security and vertical SaaS.

Founder-led. The founder is still the closer. The early-stage default, and it never fully disappears.

Cold outbound / SDR. The rarest sub-type in the corpus (12 of 111). Almost no company here leads with a bank of SDRs dialing lists — a quiet rebuke to the outbound-first orthodoxy.

The other labels, briefly

Product-led-sales

n=39 · the "bolt-on," itself four things

Enterprise-gate. Self-serve until a security or compliance need (SSO, SCIM, audit logs) forces a sales conversation. The most common trigger.

Usage-threshold. Free or cheap until a volume limit; the limit is the sales event.

Team-viral, then routed. Spreads inside an org bottom-up, then a rep engages the budget owner.

Integration-anchored. The product embeds into a workflow, and the depth of integration is what the AE expands.

Product-led & Community/Developer-led

n=16 & n=6 · the developer motions (thin, directional)

Developer adoption. Individual engineers adopt in their own workflow; the org follows.

Free-tool wedge. A genuinely useful free product is the top of the funnel, not a trial.

Open-source / ecosystem. The project is the distribution; commercial sits on top of community trust.

Partner-led & Procurement/Trust-led

n=10 & n=4 · the enterprise-access motions (thin, directional)

Channel / reseller. A partner owns the customer relationship; the vendor arms them.

Marketplace-native. Cloud or app marketplaces are the primary transaction and discovery surface.

Compliance co-sell. A certification or regulatory approval (a FedRAMP-class credential) is simultaneously the moat and the motion.

The most useful thing a motion label can do is dissolve. "Run a sales motion" means nothing until you know which of six it is.

4

The companies that break their cell

Averages describe the middle. The edges are where the learning is. These are the companies whose motion vector sits farthest from what their (stage × category) neighbours do — each a documented, source-verifiable case of running against the grain, and each with a reason.

13 Three that break the rule — deep specimens

13 Three that break the rule

THE THREE LARGEST CELL-OUTLIERS, READ IN FULL · INDIVIDUALLY VERIFIABLE

Averages describe the middle; the edges teach. Rather than catalogue every outlier, here are three companies that each overturn a different assumption — each a single, source-verifiable record, and the highest-confidence material in the report.

Rain AI — a sales-led cell, run on access

Series A · Data / AI infrastructure · expected sales-led, runs founder-networked. Rain builds compute-in-memory AI accelerators, and its cell — Series-A AI infrastructure — runs sales-led. Rain runs on the founder's network instead. Its strongest asset is not a product funnel but *access*: elite AI investors, advisors, and strategic buyers, cultivated personally. At the hardware frontier, where there is no repeatable motion yet and every design win is bespoke, relationship access substitutes for go-to-market. That is also the fragility the record names outright — the reported failure point is precisely the conversion of that access into a repeatable hardware sale. **The lesson:** capital and relationship access can *be* the motion at the frontier — right up until the moment it has to become one.

Solace — a sales-led cell, run like consumer marketing

Growth · Regulated health · expected sales-led, runs marketing-led. Regulated health reads as a slow, institutional, sales-led sale. Solace is the opposite: a Medicare-covered patient-to-advocate conversion funnel that runs, mechanically, like a consumer performance-marketing machine — Webflow/Statsig CRO, Meta and Google, influencer hiring, conditioned SEO. The first conversion is not a demo; it is an insurance-eligibility check. The regulated reimbursement engine is the moat; the *motion* is high-velocity consumer acquisition. **The lesson:** a regulated category does not dictate a regulated motion. The buyer and the payment mechanism do — and here a public payer quietly underwrites a consumer funnel.

Gamma — a sales-led cell, run product-led

Series B · Horizontal productivity · expected sales-led, runs product-led. By Series B, horizontal productivity has mostly converged on sales. Gamma stays product-led, and the reason is instructive: its distribution is baked into its output. It is easy to misread as an AI PowerPoint; the sharper truth is that Gamma makes the *artifact itself* distributable — a web-native page generated from messy input, shared as a link, embedded in workflows, copied by creators. Every deliverable is a piece of marketing. When the product's output *is* the channel, product-led growth doesn't run out of road at scale — it compounds. **The lesson:** the most durable product-led motions are the ones where using the product distributes it.

5

The laws of modern go-to-market

Six patterns recur across every cut of the corpus — sector, stage, channel, outcome, vintage. None is a controlled result; each is the compression of everything above into something a founder can carry into a room.

14	Six laws	synthesis
15	Limits, and the road to conclusive	honesty · what's next

14 Six laws

Law 1

Motion is a buyer property, not a stage.

Stage barely moves the motion; the buyer moves it. When a developer is in the buying group, sales-led falls to 0.17 against 0.25 for everyone else — the single largest swing in the corpus. Design the motion around who signs, not around your round.

Law 2

The security questionnaire is the channel.

Direct sales, partners, and a distinct procurement/trust channel are the three most load-bearing ways companies reach buyers; self-serve and developer channels are a real but distinct minority. The dominant channel in modern B2B is not a signup page — it is a security review, answered well.

Law 3

Trust beats hype — but only at the margin.

Motion does not predict outcome: winners and stalled companies run the same mix. The only tilt that survives is small and directional — winners carry slightly more procurement/trust weight (0.12 vs 0.10), the stalled slightly more marketing. On the margin, proof beats noise; nothing bigger separates success from failure.

Law 4

Growth left the seat.

Expansion splits across usage/volume, new modules, and additional seats — with usage just ahead — and in billing, volume/transaction and usage pricing together rival the per-seat subscription. For a large share of the corpus, revenue now grows with the customer's usage, not their headcount.

Law 5

Nobody runs one motion; the blend is the strategy.

The average company's strongest motion carries just a quarter of its weight (0.26); 152 of 201 are pure hybrids and only 7 are monoline. "Pick your motion" is the wrong instruction. The winning object is a stack, tuned — and it tightens toward sales, round by round.

Law 6

When companies change, they change the product — and lately, into AI.

Of 179 documented shifts in the last year, 61% shipped a new product or platform and 42% launched or repositioned around AI; only 16% changed how they sell. The motion of the year did not change. The product did — and it grew an AI surface. Meanwhile the pure-sales playbook is aging out: newer cohorts lean off sales and partners, toward product and founder.

15 Limits, and the road to conclusive

This report is an agenda, honestly scoped — not a closed argument. It is worth being precise about what it can and cannot yet carry.

The **corpus-wide** findings are solid: sales-as-default, the hybrid-stack reality, the trust-and-proof vocabulary, the Series-B proving ground. These rest on the full 201 and survive the caveats. The **cell-level matrix** is thinner: the design target was N≥15 per cell, but the current aggregate has a maximum cell N of 6, a median of 3, and zero cells at threshold. Every per-cell claim here is therefore directional — strong enough to set the agenda, not broad enough to close it.

The arithmetic is simple: 67 analytically interesting cells at N≥15 is roughly a thousand companies. That is the distance between this study — a legible, honest first map — and a conclusive one. The map is drawn correctly; it is not yet drawn densely.

> **The map is drawn correctly. It is not yet drawn densely. Both are true, and saying so is the point.**

How this was made

Every figure in this report traces to 201 structured company records — each scored on ten go-to-market motions across six funding stages and twelve categories, with a dated public source or a declared inference. No external benchmark, survey, or vendor list contributed a single number. Efficacy here is an observed-pattern signal, not a controlled experiment; thin-N cells are hedged, and the failure cohort was sourced on purpose.

Cite as — The State of Go-To-Market 2026. Accountmade, July 2026. Corpus of 201 companies; figures traceable to source.

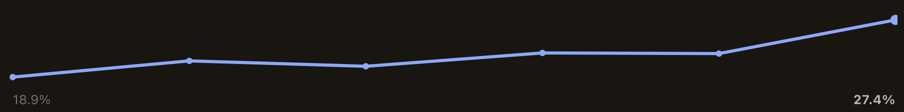

THE FINDING, IN ONE LINE · SALES-LED WEIGHT, SEED → PUBLIC