HelixScale

RESEARCH REPORT · 2025–2030

What Becomes Scarce

Redesigning go-to-market from first principles, now that execution costs are falling rapidly.

ABUNDANT EXECUTION SCARCE ADVANTAGE

PUBLISHED BY HELIXSCALE V2.0 · 2026

AI lowered the cost of producing go-to-market. It did not lower the cost of being worth listening to.

Nearly every activity that can be written down as a procedure is becoming abundant: research, drafting, personalization, sequencing, reporting, routine analysis. Abundance removes those activities as sources of advantage, because a capability every competitor buys in the same quarter is not a capability.

What remains scarce is anything that has to be earned over time. This report applies a three-part test to each candidate: is it scarce, is it ownable, and does it compound. Buyer attention is scarce but nobody owns it. Judgment is scarce and ownable but walks out with the person who holds it. Four assets pass all three: proprietary context, relationships, ecosystem standing, and the rate at which an organization learns.

The organizational conclusion follows from that. An AI-native go-to-market function is not the current org with tools bolted on, and it is not the current org with fewer people. It is an org rebuilt around a shared context layer, with continuous sensing in place of campaigns, and humans concentrated on the work that cannot be generated.

ABUNDANT	SCARCE, NOT OWNABLE	SCARCE AND OWNABLE
Research, content, personalization, sequencing, reporting, workflow execution.	Buyer attention, timing, market trust. Competed for, never held.	Context, relationships, ecosystem standing, organizational learning.

Strip away the org chart and four jobs remain.

Go-to-market has accumulated a large amount of structure: functions, tools, campaign calendars, handoff stages. Almost none of it is fundamental. Underneath, a revenue organization does four things.

01 **Find out where a problem is live** Sensing. Which organizations have the problem right now, and how would you know.	02 **Reach the people it affects** Access. Getting a hearing from a named person who can act.
03 **Give them reason to believe you** Proof. Evidence that survives contact with a skeptical evaluator.	04 **Get better at all three each cycle** Learning. Whether the next attempt starts from a stronger position than the last.

Three assumptions the current structure rests on, and their status

Execution capacity is the scarce input, so add headcount to add pipeline.

No longer holds. Capacity is now purchasable at rapidly falling marginal cost by everyone at once.

Work should be split by stage, with specialists at each handoff.

Weakening. Specialization existed to make manual work efficient. The efficiency argument is gone; the context loss it causes remains.

Demand arrives in campaign cycles, so plan quarterly and launch in waves.

Obsolete. Buying windows open on the buyer's calendar, not the seller's.

Activity is a fair proxy for progress, because activity was expensive.

Broken. Cheap activity stops correlating with anything and becomes a cost the buyer absorbs.

Everything that can be written down as a procedure.

Gartner puts AI software spending on a path from US$124 billion in 2022 to US$297 billion by 2027, with more than 70% of independent software vendors expected to embed generative capability by 2026. The practical meaning for a revenue team is that these capabilities arrive inside tools it already pays for, on the same schedule as its competitors.

ACTIVITY	WHAT IS GENUINELY ABUNDANT	WHAT IS NOT
Research	Summarising anything public about an account	Knowing what is not public
Content	Producing competent volume in any format	Having something to say
Personalization	Matching surface details to a recipient	Relevance the recipient agrees with
Prospecting	Building and working lists at any scale	Being let in
Analysis	Describing what happened in the data you hold	Deciding what to do about it

The right-hand column is the report

Each row splits at the same seam. The generative half needs only public inputs and a procedure. The other half needs something proprietary, or standing, or a decision someone is accountable for.

Abundance is real, adoption is not free

Gartner expects more than 40% of agentic AI projects to be cancelled by 2027 on cost and unclear value. Capability being available is not the same as it working inside a specific organization.

Cheap execution degrades the channel it is aimed at.

The second-order effect rarely appears in vendor material. Every seller's new capacity lands on the same finite set of buyers at the same time. Volume per buyer rises, attention per message falls, and the response rate that justified the volume declines. The tooling performs exactly as sold while the channel gets worse for everyone using it.

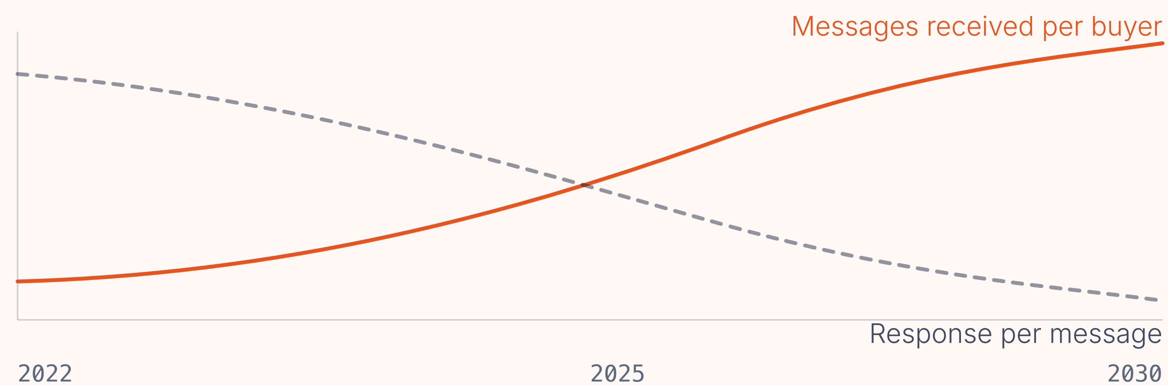

Illustrative. Shape of the relationship, not measured values.

Why it does not self-correct

Each seller sees a falling response rate and rationally sends more, since sending is nearly free. The cost of the added volume lands on the buyer, not on the sender, so the individual incentive keeps pointing the wrong way.

What the buyer does in response

Filters harden, cold channels are delegated or ignored, and buyers fall back on sources with existing standing: peers, communities, partners and people they already know. Reporting on customer preference for human contact points the same way.

> The exit from a noise trade is never more volume. It is a route the noise does not travel.

Scarce is not the same as advantageous.

Most commentary stops at listing what stays scarce. That list is not decision-useful, because several items on it cannot be held by anyone. Three questions separate a real asset from a market condition.

Scarce

Genuinely limited, or only inconvenient to produce.

Ownable

Can your organization hold it, or is it competed for continuously.

Compounding

Does using it leave you with more of it, or does it deplete.

CANDIDATE	SCARCE	OWN	COMPOUND	READING
Buyer attention	Yes	No	No	A market condition. Plan around it.
Timing	Yes	No	No	Not a resource. An output of sensing.
Judgment	Yes	Yes	Partly	Held by people. It leaves with them unless written down.
Proprietary context	Yes	Yes	Yes	An asset. Grows every cycle it is used.
Relationships	Yes	Yes	Yes	An asset, if held by the company and not only individuals.
Ecosystem standing	Yes	Yes	Yes	An asset. Slow to build, hard to copy.
Organizational learning	Yes	Yes	Yes	The rate asset. It governs how fast the other three grow.

Analytical framework, not measured data.

Four assets, and the trust graph that connects them.

Proprietary context

What your organization knows that is not on the public internet: which objections killed which deals, which changes preceded a purchase, what the last hundred conversations actually revealed. Models cannot generate it and competitors cannot buy it.

Relationships

The set of people who will read a message because of who sent it. Roughly four in five buyers still prefer human contact in the buying process, which is precisely the input that cannot be manufactured at volume.

Ecosystem standing

Credibility inside a marketplace, an integrator network, an open-source project or a practitioner community. It transfers to whoever is introduced through it and takes quarters to build, which is exactly why it holds.

Organizational learning

Whether what one person learns on Tuesday changes what the team does on Wednesday. This is the rate variable. Two companies with identical assets diverge on how fast each cycle updates the next.

PATHS INTO A NAMED ACCOUNT, MOST OF THEM UNMAPPED

Cold outreach treats an account as a stranger. It rarely is. The graph exists in most companies already; nobody has written it down, so the strongest path goes unused and the account gets worked cold.

Borrowed standing travels further than borrowed lists.

A contact list gives you an address. An ecosystem gives you credibility that transfers to whoever is introduced through it. Treated as a distribution channel with its own targets rather than a partnerships side project, it produces the same measurable output as outbound: named accounts entered, meetings held, pipeline created.

CHANNEL	WHAT IT UNLOCKS	RAMP
Existing customers	Peer introductions inside the same buying committee and vertical	Days
Investors and advisors	Portfolio access at executive level, usually with a strong default reply	Days
Technology partners	Joint accounts where the integration is already deployed	Weeks
Integrators, resellers	Enterprise procurement paths and budget that is already approved	Months
Practitioner communities	Standing with the technical evaluator before a deal exists	Quarters

> The channels with the longest ramp compound the hardest. They have to be started before the quarter that needs them.

Not the same org with tools attached.

The traditional structure was designed around a scarce input that is no longer scarce. Rebuilt around what is, the differences are structural.

DIMENSION	TRADITIONAL	AI-NATIVE
Organizing unit	Specialist function per funnel stage	Account, owned end to end
Data	Per-tool silos, reconciled in reporting	One shared context layer, written to by everyone
Cadence	Campaigns on the seller's calendar	Continuous sensing on the buyer's
Prioritization	Static ICP list, refreshed annually	Observed change, ranked by how open the window is
Execution	Manual, and the constraint on output	Automated where it is procedural, and no longer the constraint
Human focus	Spread across production and admin	Concentrated on judgment, trust and high-value conversations
Entry route	Cold by default, warm by luck	Mapped relationship and ecosystem paths first
Measurement	Activity volume as a proxy for effort	Accounts entered, paths that worked, pipeline created

This comparison is our model of the change, not a survey of adopted practice. Taking the right-hand column selectively does not work. A team that senses continuously but still routes through stage handoffs loses the timing advantage at the first queue.

The piece most teams skip, and the reason their AI projects fail.

Gartner's forecast that more than 40% of agentic projects will be cancelled by 2027 points to a gap between what the technology can do and the value organizations actually realize from it. Gartner attributes this to cost and unclear business value. Our own reading, offered as a thesis rather than a finding, is that fragmented context is one important contributor: automation applied to an organization with no shared, current picture of its own market produces confident output nobody trusts.

WHAT IT HOLDS	WHAT MAKES IT WORK
Accounts and the changes observed against them. The relationship graph, including who knows whom and how well. Every conversation and what it revealed. Objections, and which ones ended deals. Which entry path produced which meeting.	One place, not per-tool copies. Written to as a by-product of doing the work rather than as a reporting duty. Readable by both people and systems. Current, since a context layer that lags by a quarter describes a market that has moved.

Three failure modes worth naming

The CRM substitute

A pipeline ledger records stage and amount. It does not record what was learned, which is the part that compounds.

The purchased substitute

Bought data is by definition available to competitors. It can seed the layer. It cannot be the layer.

The individual substitute

Context held in one experienced person's head is real and leaves when they do. Ownable by them, not by the company.

Design for higher leverage, not for headcount reduction.

The common claim is that AI-native means smaller teams. That reads the change as a cost programme, and it misstates the design goal. AI-native GTM should be designed for higher leverage per person. Smaller teams may result, but they are not the objective, and a team sized down before its leverage actually improved has simply cut capacity.

HUMANS MOVE TOWARD	HUMANS MOVE AWAY FROM
Deciding which accounts deserve the quarter, and why.	Assembling lists and enriching records by hand.
Asking for and honouring introductions, which is a reputational act.	Drafting the first version of anything routine.
The conversations where a deal is actually won or lost.	Manual sequencing, scheduling and follow-up mechanics.
Judging whether the machine's output is true before a buyer sees it.	Producing internal reporting that describes what happened.
Writing down what was learned so the organization keeps it.	Re-researching what the organization already knew.

Leverage has a ceiling set by trust

Sellers spend roughly 28% of the week actually selling. Recovering that time raises leverage only if the recovered hours go into the four scarce assets. Spent on more sending, they fund the noise trade.

The measurable version

Track accounts meaningfully entered per person per quarter, not messages per person per day. If the first number rises while headcount holds, leverage improved. If only the second rises, nothing did.

From campaigns to continuous sensing.

A campaign assumes the seller chooses when the market is ready. Continuous sensing assumes the buyer does, and organizes the week around changes that already happened. Gartner expects 65% of B2B sales organizations to shift from intuition to data-driven decisions by 2026, which makes the differentiator not whether you watch, but what you watch and how fast you act.

STRUCTURAL Hiring and org change. Budget follows the role.	FINANCIAL Capital raised, a new fiscal plan, a cost programme. Each reopens a settled decision.	TECHNICAL Migrations, dependencies, public roadmaps. Technical buyers signal in public.
COMMERCIAL New region, new pricing, a partnership. Expansion creates unowned problems.	RELATIONAL A former user or champion lands somewhere new. The warmest signal there is.	BEHAVIOURAL Issues, talks, forum posts and job specs that describe the problem you solve.

Two tests before a change counts

Does it plausibly create budget or urgency for the problem you solve, and can you name the person it affects. A change that fails either test is trivia, and acting on it manufactures noise.

Shelf life is the whole point

Most changes are worth acting on for days. By the time one reaches a vendor shortlist the useful window has closed and the conversation is a bake-off.

The operational unit that replaces the campaign is a short weekly list: accounts that changed, the person affected, and the strongest available path in. It stays short on purpose, because a list nobody can work in a week is a backlog.

Cost per message collapsed. Earning a conversation is still expensive.

The economics are frequently misread, because the metric that improved is not the metric that matters. Production cost fell sharply and is measurable. What it costs to earn a hearing does not necessarily fall with it, because attention, trust and standing remain constrained.

WHAT GENUINELY FELL	WHAT DID NOT
Cost to research an account. Cost to draft anything. Cost to run a sequence. Cost to produce a report. These are real savings and they are available to every competitor on the same terms.	Cost to earn a hearing, build ecosystem standing, or be believed by a skeptical evaluator. These remain constrained by competition for attention, trust and standing.

The measurement shift that follows

STOP REPORTING	START REPORTING
Messages sent, sequences launched	Named accounts entered
Contacts touched	Accounts with three or more people engaged
Reply rate in aggregate	Meetings by entry path, warm against cold
Content published	Assets that shortened a live deal
Pipeline created this quarter	Pipeline created, and what the quarter taught

Expect the volume numbers to look worse under this model. That is the intended result, and it is why the change has to be agreed before the first reporting cycle lands.

Vendor-neutral. Buildable with tools you already have.

01

Name 50 accounts, not 500

Where the deal is large enough and the problem is already live. Everything downstream inherits this list's quality.

02

Define five observable changes

Each written as a sentence with a source and a window. If nobody can point at where it comes from, it is not observable.

03

Map the relationship graph once, properly

Every founder, employee, customer, investor and partner connection against the 50. Most teams surface relationship paths that were never documented anywhere.

04

Run one weekly list

Accounts that changed this week, path attached, worked within days. One conversation beats one hundred sends.

05

Write the proof the conversations asked for

Content earns its place by shortening a live deal, not by publishing on schedule.

06

Review by change and by path

Monthly, cut what produced nothing and double what produced meetings. This is where compounding actually happens.

DAYS 1–30

Account list, change definitions, relationship graph. No outreach yet.

DAYS 31–60

Weekly list live. First warm entries. Log every path used and its result.

DAYS 61–90

Cut what failed, build the proof assets, extend into one ecosystem channel.

Five questions worth answering before the next planning cycle.

01	Which 50 accounts would change the year?	If the list is not written down, targeting is happening implicitly at the rep level.
02	What would tell you an account is in motion?	Name the observable changes and where you would see them. Then check how fast you actually do.
03	Who in your network can open those accounts?	The answer is usually larger than expected and almost never documented anywhere.
04	What share of pipeline came in warm?	If nobody tracks entry path, the organization cannot tell which of its assets are working.
05	What compounds if you run this for a year?	Anything that resets each quarter is a campaign. Anything that carries forward is a system.

The teams that win the next five years will not be the ones producing the most.

They will be the ones that see change first, hold the relationships that open the door, and learn faster each cycle while everyone else buys the same execution layer.

How we assemble the model at HelixScale.

The model in this report can be built several ways, and a team can build it without a vendor. This is the arrangement we run, included so the argument has a concrete referent rather than as the conclusion the report was written toward.

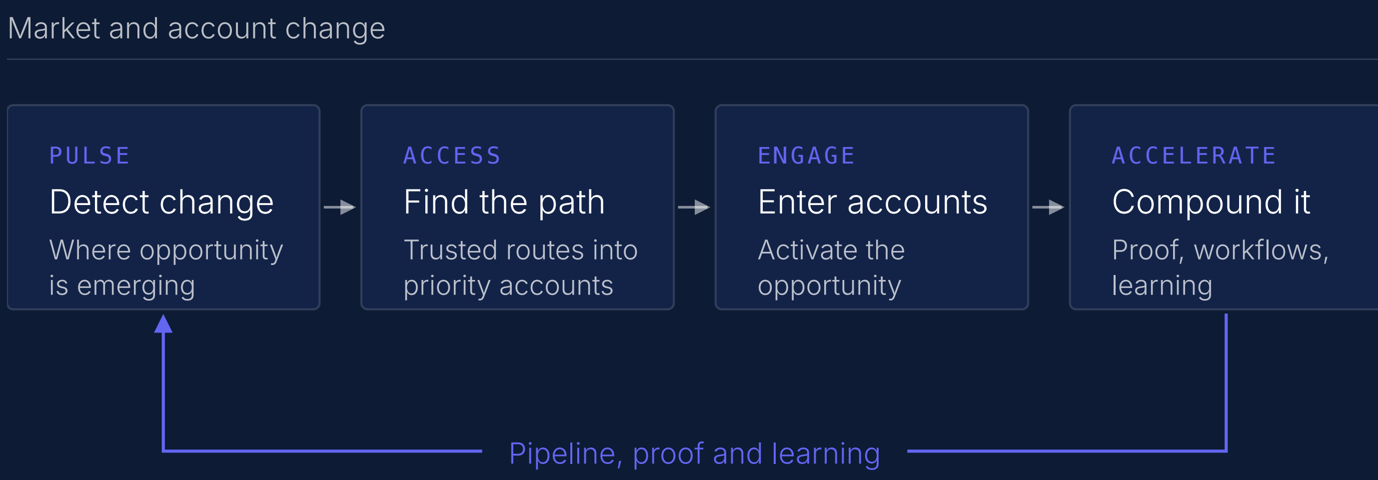

Each completed cycle returns sharper change definitions, a wider relationship graph and proof assets that shorten the next conversation. The loop is the asset, not any single stage inside it.

HelixScale

Qualified conversations. Not just contacts.

We run the loop described in this report as one system for technical B2B teams. If you would rather build it yourself, the ninety-day playbook on page 14 is written to be used without us.

Start a pilot

helixscale.com

What the figures are, and what they are not.

Figures cited here come from publicly available analyst forecasts and survey reporting published between 2022 and 2025. Survey-based numbers reflect specific populations and sample sizes and should be read as directional rather than universal.

Forecasts are informed estimates and actual outcomes may differ. Commentary drawn from news reporting is paraphrased. No customer results, pilot outcomes or performance claims for HelixScale appear anywhere in this report.

1 Gartner forecast on AI software spending and generative AI adoption

2 Gartner research on seller time allocation and quota attainment

3 Reuters reporting on agentic AI project cancellations and vendor substance

4 Gartner survey of service leaders on conversational generative AI

5 Gartner prediction on data-driven B2B sales organizations by 2026

6 Sales productivity and quota attainment statistics, Q4 2024

7 Smartsheet study on time spent in repetitive work

8 CX Dive reporting on buyer preference for human contact

METHODOLOGICAL NOTE

This report synthesises independent research with an operating model described as a method rather than a benchmarked result. The scarcity test and the organizational model are analytical frameworks, not measured findings. Readers evaluating any figure should consult the original source for its full context and sample.