



Grow
Intelligently

Artificial activity:

Your 2027 plan bets on AI leverage
your GTM model can't deliver

AI use across GTM exploded over the past eighteen months; 93% of B2B SaaS companies now ship live AI functionality, creating a proliferation of point solutions.¹ CEOs are staking growth plans on the AI surge, with 65% saying customer-facing teams will need AI-driven efficiency gains to hit their goals.² Instead of the needed returns, they're getting shadow IT, siloed and non-impactful use cases, pockets of activity, perhaps some AI heroics, but ultimately limited business outcomes, with only 18% reporting high ROI from AI in their commercial organizations.²

The gap between AI expectations and reality is now a P&L problem. Having lost patience with pilots and experiments, boards are asking a simple question: spend is up, AI activity is up, so where are the returns? Leaders who can show how they combined AI with process and talent changes to drive operating leverage will set the agenda; everyone else will be explaining failed pilots.

Answering it requires a 2027 annual plan that can stand up to three direct board questions:



1. Where does the operating leverage from AI lie?



2. What is it worth when AI is fully integrated into the GTM model?



3. Most pressing: What changes will make next year different?

The window for real answers is the plan being built right now, and the CRO seat, already among the shortest-tenured in the C-suite,³ rarely gets a second planning cycle. AI is unlikely to take the commercial leader's job. Failing to turn it into operating leverage will.


| The gap between AI expectations and reality is now a P&L problem.



AI is creating more activity than ever but less intelligence than you need

Look inside your own GTM organization. AI has led to an explosion of new capabilities: marketing ships several times the content it did two years ago, sequences run around the clock, every GTM rep has a meeting summarizer, a deal coach, and every vendor is pushing leaders to use their latest AI feature.

More activity than ever. Less return than you need.



This has created what we call artificial activity:

AI-generated motion that grows faster than the organization can absorb it, pushing return on GTM spend toward diminishing or even negative returns.

AI is helping GTM teams create more activity than ever. Converting that activity into enterprise value is where AI struggles.

The reason is structural: AI produces two kinds of gains, and only one of them turns into advantage.



1. Speed is the first: drafting emails, summarizing calls, building decks. Real hours saved, but they are collected by every competitor using similar tools, which is why they show up as efficiency and then vanish into the baseline.



2. Intelligence is the second: the pattern recognition that reads broad, varied data and spots what no person can. The accounts quietly in market, the sellers one skill away from the top tier. Those patterns exist in every organization, but are often missed because the data is spread across functions and disconnected, so AI only finds them when it sees the whole picture.

Teams chasing speed feed artificial activity and keep AI below the intelligence threshold: the point where AI sees enough connected data to read patterns instead of just producing motion. Each point solution runs below it by design, bought for speed, measured on the activity it generates, and reading one function's slice of the customer: accurate, fast, and blind to the rest. A rep juggling half a dozen point solutions gets six partial readings of the same account and no verdict.

Without meaning to, these teams are trading efficacy for activity: more motion built on less understanding. The bill lands in the two numbers every GTM leader is paid on, new bookings and net revenue retention.



Three traps keep AI below the intelligence threshold

Why does the most data-rich part of the company leave AI blind? Because go-to-market produces value as a single system with one customer and one cost of growth, while AI arrives function by function: bought, measured, and budgeted in silos. The whole picture never assembles. Three traps keep it that way:

1. Pilots measure speed, not intelligence.

Success gets counted in activity generated and siloed efficiency gains, never in financial outcomes, so more activity flows into one part of the GTM motion while the broken parts stay broken. That is what a pilot mostly exposes: the dependencies that keep automation from paying.

2. AI efforts are siloed by design.


Every functional leader under pressure to show results builds their own AI capability on their own data: four different AI handbooks and four partial views of the same customer. Each team's activity multiplies but the intelligence never compounds.

3. Nobody owns the 24-month question.

What should the GTM model look like two years from now, what do its economics become, and which teams and routes to market survive the change? Answering takes the cross-functional view no single leader holds.

The third trap breaks a fifteen-year pattern: for the first time, the transformation runs toward contraction rather than expansion, and contraction is not a euphemism for cutting people. The redesign is a triage of work: what still needs to happen at all, what a chain can absorb end to end, and what stays with humans because judgment, relationships, and accountability live there. That sorting decides which roles change, which disappear, and which become more valuable. No functional leader designs the shrinking of their own team, which is why the question sits unowned.

Fifteen years of adding structure. The first redesign that removes it.



These traps persist because, for functional leaders, acting feels riskier than waiting. Doing nothing reads as a slow quarter; acting feels like painting a target on you and your team by staking your name on productivity gains from emerging technologies or volunteering to shrink your own roster.

The boldest bets sit unowned, so the CEO becomes their arbiter by default.

GTM leaders can’t afford to “wait-and-see”: the commercial leadership role is being redefined around AI’s potential whether or not these traps get fixed. GTM leaders won the last decade with human middleware: playbooks, headcount, inspection. The next decade’s leaders will generate operating leverage by building systems that eliminate point solutions and instead drive predictive, action-oriented intelligence at the deal, account, and business performance levels.

AI chains expand operating leverage, point solutions just expand activity

Board Question #1 asks where the operating leverage lies. The answer sits in chains, not individual tools. A chain is a connected sequence of AI capabilities wired into one workflow, running from signal (the system spots something) to action (the right move is triggered) to proof (the result sharpens the next signal). Companies getting leverage from AI today stopped treating the use case as the unit of value. Activity stops being artificial when it compounds across a chain. The question becomes which chains to build first.

Four enablers determine whether a GTM model produces leverage, and none of them is new: growth strategy, the operating model, talent, and data. What is new is AI chained to move each one, in ways that break the productivity and capacity assumptions sitting in your 2027 plan.

SBI identified six AI chains across GTM: demand creation, deal conversion, forecast accuracy, retention and expansion, rep productivity, and the foundation that feeds them all. Each chain earns its place by moving an enabler, and for most companies starting out, four carry the leverage, one per enabler: demand creation rewires growth strategy, retention and expansion rewires the operating model, rep productivity rewires talent, and the foundation rewires the data layer.

Four chains carry the leverage

Enabler	Each chain’s job	The payoff
Growth strategy	Create demand efficiently	5-8% bookings gain from increased pipeline
Operating model	Keep and grow the base	+5% NRR average
Talent	Improve rep productivity	6-12% productivity per head
Data & insights	Make informed bets	40-60% of RevOps capacity freed

Each chain below follows the same arc, signal to action to proof, with the payoff called out.

Growth strategy: Demand creation.

A growth strategy answers two questions: which markets to prioritize based on right to win, meaning the evidence you beat competitors there, and which route wins in each of them. Most companies answer once a year in planning while markets move weekly. This chain creates demand you can win, at the market's pace. AI reads buying signals, builds account understanding no analyst team could, and models propensity to buy on richer variables than firmographics. Outreach gets leaner and sharper: sequences fire only on accounts the model says are in market, sellers follow modeled potential instead of the org chart, and results re-rank the list as conditions change.

Companies running this chain report 10 to 15% more client-facing time and a 5 to 8% bookings productivity lift per seller.⁵

Case in point.

A national field-services company built its growth plan on gaining share in specific high-profit markets without adding headcount, yet sellers spent roughly a fifth of their time driving territories, hunting for prospects. SBI conducted AI-driven scoring on service-demand signals, property economics, and branch proximity, cutting one branch's ~550 raw prospects to ~150 re-ranked targets. Sellers stopped scouting and started selling, and a fifth of seller capacity moved to client-facing work with no new heads.

Operating model: Retention and expansion.

An operating model is the set of interconnected processes that turn strategy into revenue: demand generation, coverage, methodology, comp and quota, renewal motions. Most of those processes run on the calendar. The cost of a fixed calendar is most apparent in motions designed to keep and grow the base: a generic playbook fires on every account 180 days out, blind to whether the account is quietly expanding or quietly leaving, and this calendar approach is part of how half of SaaS companies ended up with lower NRR than two years ago.⁴

Using data to predict churn is not the new part; scores built on support tickets and health surveys read one slice of the account, arrive late, and change little. The AI difference is reading product usage at scale.⁴ The chain converts each prediction into a play with a named owner, and retained revenue tunes the model. The payoff comes from doing less: fewer generic plays, no heroics on unsavable accounts.

Companies managing the base this way report the ability to predict churn and expansion with 90% accuracy a year out and an average 5% NRR improvement.⁴

Case in point.

One software provider entered 2026 facing nearly 28,000 renewal events worth more than half a billion dollars, gross renewal rate tracking in the mid-70s against a high-80s benchmark, and multi-year contracts meant customer decisions were solidifying a year before account teams engaged. Usage-based prediction surfaced more than \$40M of at-risk ARR inside accounts the health scores called green, and a twelve-week renewals growth office turned every signal into an assigned action. The same team now covers the largest renewal base in company history.

Talent: Rep productivity.

The biggest prize in talent is moving the middle of the roster. Top performers deliver multiples of the middle, and closing that gap takes more than coaching. This chain builds a complete picture of every seller, anchored to what top performers do differently, and runs the whole talent lifecycle as one closed loop: source the A players who fit your environment, hire them the way you would win a deal, ramp them to full productivity, develop them past your current ceiling, and feed what the system learns about who succeeds back into who you source next.

Today, ownership breaks the loop. HR sources, enablement ramps, field leaders develop, and the intelligence splits with it. Each owner sees a slice; nobody sees the pattern. The connector is a live map of skills, performance, and behaviors, hundreds of data points per person, the talent mapping SBI runs across client commercial teams. Coaching hours stop counting as progress; the map points to each rep's next move.

Modeled through, the loop is worth 6 to 12% more productivity per head, the same roster producing outsized results.⁶

From the field.

A new path to rep productivity is emerging. Instead of reading performance metrics in isolation, we score every seller across hundreds of data points: competencies assessed, behaviors read from live activity, performance tied back to both. One map, and patterns surface that no single owner of the talent lifecycle could see. We often find that the competencies HR's hiring profile screens for explain little of who actually performs, while a handful nobody screens for explain a lot. A nine-month ramp analysis misses that performance diverges by background, and the cohort that starts slowest, sellers out of business services, out-earns everyone by year three. Findings like these each rewire a different stage of the talent lifecycle: sourcing profiles, ramp expectations, development targets. The larger gain comes from connecting them: each hiring class's results flow back into the map, so sourcing, ramp, and development all get smarter every time.

Data and insights: Foundation.

This is the chain the others plug into, the one that lifts them above the intelligence threshold. Its keystone is the account brain: everything the company knows about an account combined with all external signals, the whole picture finally assembled. The same foundation absorbs the manual data work running underneath every function. Sellers stop reconciling a dozen feeds; buyers stop receiving contradictory touches. Every play the other chains run lands its outcome back in the brain, sharpening the picture with use.

Automating that data work frees 40 to 60% of RevOps capacity within two years, and the foundation pays for itself by feeding the other chains.⁶


Case in point.

A PE-backed software company believed a long data cleanup had to come before AI; its CRM was full of duplicate and conflicting records. Agentic deduplication against the existing data lake, with revenue control totals so no dollar disappeared, produced a usable account foundation in weeks rather than quarters. The sponsor is repeating it across the portfolio, and the work a cleanup would have staffed simply stopped existing.

The bets you make this quarter decide 2027

Board Question #2 is what AI is worth when it is fully integrated into the GTM model. The risk leaders face is trying to give a “safe” answer pointing to across-the-board improvements vaguely attributed to AI. When this approach is fed into the bookings model that runs through 2027, it puts the whole number at risk once specific returns fail to materialize.

Where 2027 results get decided



Performance against the 2027 bookings plan, indexed. Every path looks identical on the day the plan is signed.

That failure mode has a history. The advice of the last two years was to build leaner functions and redeploy the savings, but what most companies got was the same headcount plus AI subscription costs. The typical approach, focusing on task automation, speeds up pieces of a workload in each silo, but the work itself survives.

Headcount comes down when whole workloads disappear, and a workload disappears only when a chain absorbs it end to end.

Chains turn vague answers into ones that can be priced. Each frees money inside named functions, and the money goes one of two places: back into growth or down to margin. Every credible plan captures it one of two ways: fewer heads or outsized output from the same heads. Landing the bets takes the CEO, CFO, and CRO agreeing on which functions, how much, and which destination. The four functions that follow are priced the way finance reads them, function by function, with the bet each implies.

SDRs: The research, targeting, and outreach workload that consumed most of an SDR's day is absorbed into the demand creation chain. Today's point tools are making that workload faster; the demand creation chain augments that speed with the intelligence to prioritize efforts. Going forward, prospecting rebuilds around a smaller pod of higher-skill sellers working only in-market accounts, and a third to half of today's prospecting capacity comes free, roughly \$0.7M to \$1.2M for every \$100M of ARR,⁶ to fund client-facing coverage or drop to margin.

Enablement: Content-generation tools are making the content factory faster, but point solutions for improving rep productivity are hitting their limits. The rep productivity chain spans the whole talent loop, wiring coaching and guidance into the moments where sellers work to improve performance. This allows the team to get leaner, and the money is downstream: the selling time the chains return to reps is worth another \$0.7M to \$1.1M per \$100M of ARR,⁶ the equivalent of adding sellers without adding payroll, and the enablement function rebuilds around performance engineering.

RevOps: The function evolves into a strategic partner when the foundation chain automates the reporting, hygiene, and administration that fill most of its calendar today. Reporting copilots make the reports faster; the chain makes most of them unnecessary, which is how well-run operators can plan for a materially smaller team within two years, releasing \$0.5M to \$0.8M per \$100M of ARR.⁶ What remains rebuilds around revenue architecture, focusing on the models and the leverage math necessary to succeed.

Customer success: Health-score dashboards hand teams a score and leave the next move to them; the retention and expansion chain makes the move, triggering the right play

with a named owner while the account is still winnable. This means coverage stops being a ratio decision and becomes a deployment decision: each CSM carries meaningfully more book, and customer success cost falls from roughly 8% of ARR toward 5 to 6%, worth \$2M to \$3M per \$100M of ARR and the largest pool of the four.⁶

Add the four together, along with smaller savings inside enablement itself, and a \$400M ARR company is looking at roughly \$16M to \$25M a year of capacity to reallocate against higher-yield bets, before counting anything from marketing.⁶

That is Board Question #2 answered, and the specificity is the point: each chain investment ties to a named team.

One rule keeps the redesign honest: prove the chain before you resize the function. Teams that cut first often struggle and are forced to reverse their decisions, ultimately realizing neither the savings nor the capability.

The planning question to sit with: do your 2027 assumptions reflect the post-AI operating model you will actually run, or the pre-AI model you ran in 2023? Answering it honestly takes more than conviction; it takes a mechanism.

The Growth Office: the mechanism most plans are missing

Board Question #3 is where boards press hardest: what changes will make next year different? In most organizations, P&L gains attributable to AI have been underwhelming or hard to measure, and 2027 has to be different.

Getting the bets right starts with aligning teams inside the planning window, on the first GTM transformation in fifteen years that runs toward contraction. That takes a cross-functional mechanism most companies do not have today: one owner for the growth initiative portfolio, common data underneath it, and tight accountability across GTM teams. We call it the Growth Office; call it whatever you like, but make sure it does three things:

1. Brings market objectivity:


The hardest calls (whose team shrinks, which motion gets rebuilt) are the ones nobody inside volunteers for. An external view across dozens of companies facing the same problem shows what is working elsewhere, and turns those calls from career risk into evidence.

2. Activates the plan:

It converts planning-season assumptions into running changes, with your team, not to your team.

3. Embeds the capabilities into your systems and your team's operating rhythm, so the results keep compounding after the office stands down.

The Growth Office does three things



Chains supply the leverage. The growth office supplies the capacity to absorb the change.

Most companies mistakenly believe they already have a growth-office equivalent, having stood up a committee testing tools inside the current model, staffed by leaders who each own a function being redesigned. This approach will fall short, and it exposes the quiet reason a growth office cannot be fully internal.

The bets necessary to succeed require independent arbitration, because everyone inside the room is affected by the outcome.

Another committee, working group, or tiger team designed to integrate a new technology and drive change can be a difficult sell. The hesitation from leaders burned by past transformations is earned, which is why the growth office's mandate is narrow by design. Avoid scope creep like ICP redefinition or replatforming your stack.

The goal is to get AI wired into the model you already run, one governed initiative at a time, and scoped to a single motion so it can be time-boxed.

The office stays small, senior, and paced to your planning cycle.

Case in point: a renewals growth office made prediction pay

The chain found the money; the office made it real. Return to the software provider from the retention and expansion chain, this time from the governance side: its twelve-week renewals growth office ran one weekly cross-functional forum with four scenario plays and explicit ownership, converting every prediction into an assigned action, and the hand-back was planned from day one. The predictions existed before the office did; outcomes changed when governance forced action on every signal. The constraint sat in the organization's capacity to absorb what the models saw, and the office manufactured that capacity.

The intelligence dividend

Three questions tell you where you stand going into planning.



1. When your board asks what AI has done to your return on GTM spend, what's your answer today?



2. How many AI use cases are live across your GTM organization, and how many tie to revenue, retention, or productivity per rep?



3. Is AI a line item in your plan or a new set of operating assumptions underneath it?

If the answers sting, act on them this quarter. The leverage lives in the chains. Its worth shows up when the four functions rebuild around them. And a growth office makes next year different by building the capacity to absorb it all. Pick the four chains, give the portfolio one owner, and set 2027 assumptions for the model you will run. The trade is less activity for more intelligence, and the assumptions you sign this quarter decide who collects the dividend.

Endnotes

1. SBI 2026 State of SaaS Pricing Report, n=356. Respondents reporting AI functionality in their products.
2. SBI Q4 2025 CEO Survey, n=118.
3. Nick Toman, Bryan Kurey, and Dave Lingeback, "The High Costs of Chief Revenue Officer Turnover," *Harvard Business Review*, October 10, 2024.
4. SBI NRR dataset, 2026 update.
5. SBI client program benchmarks.
6. SBI GTM capacity model: function capacity decomposed into judgment, coordination, and research/administration layers using time-allocation observations from client engagements and SBI published research, with chain absorption rates applied to non-judgment layers and outputs validated against SBI Q4 2025 CEO Survey role-level expectations. Ranges reflect conservative and aggressive adoption cases; figures are planning estimates, not measured outcomes.



Grow
Intelligently