SBI | Grow Intelligently

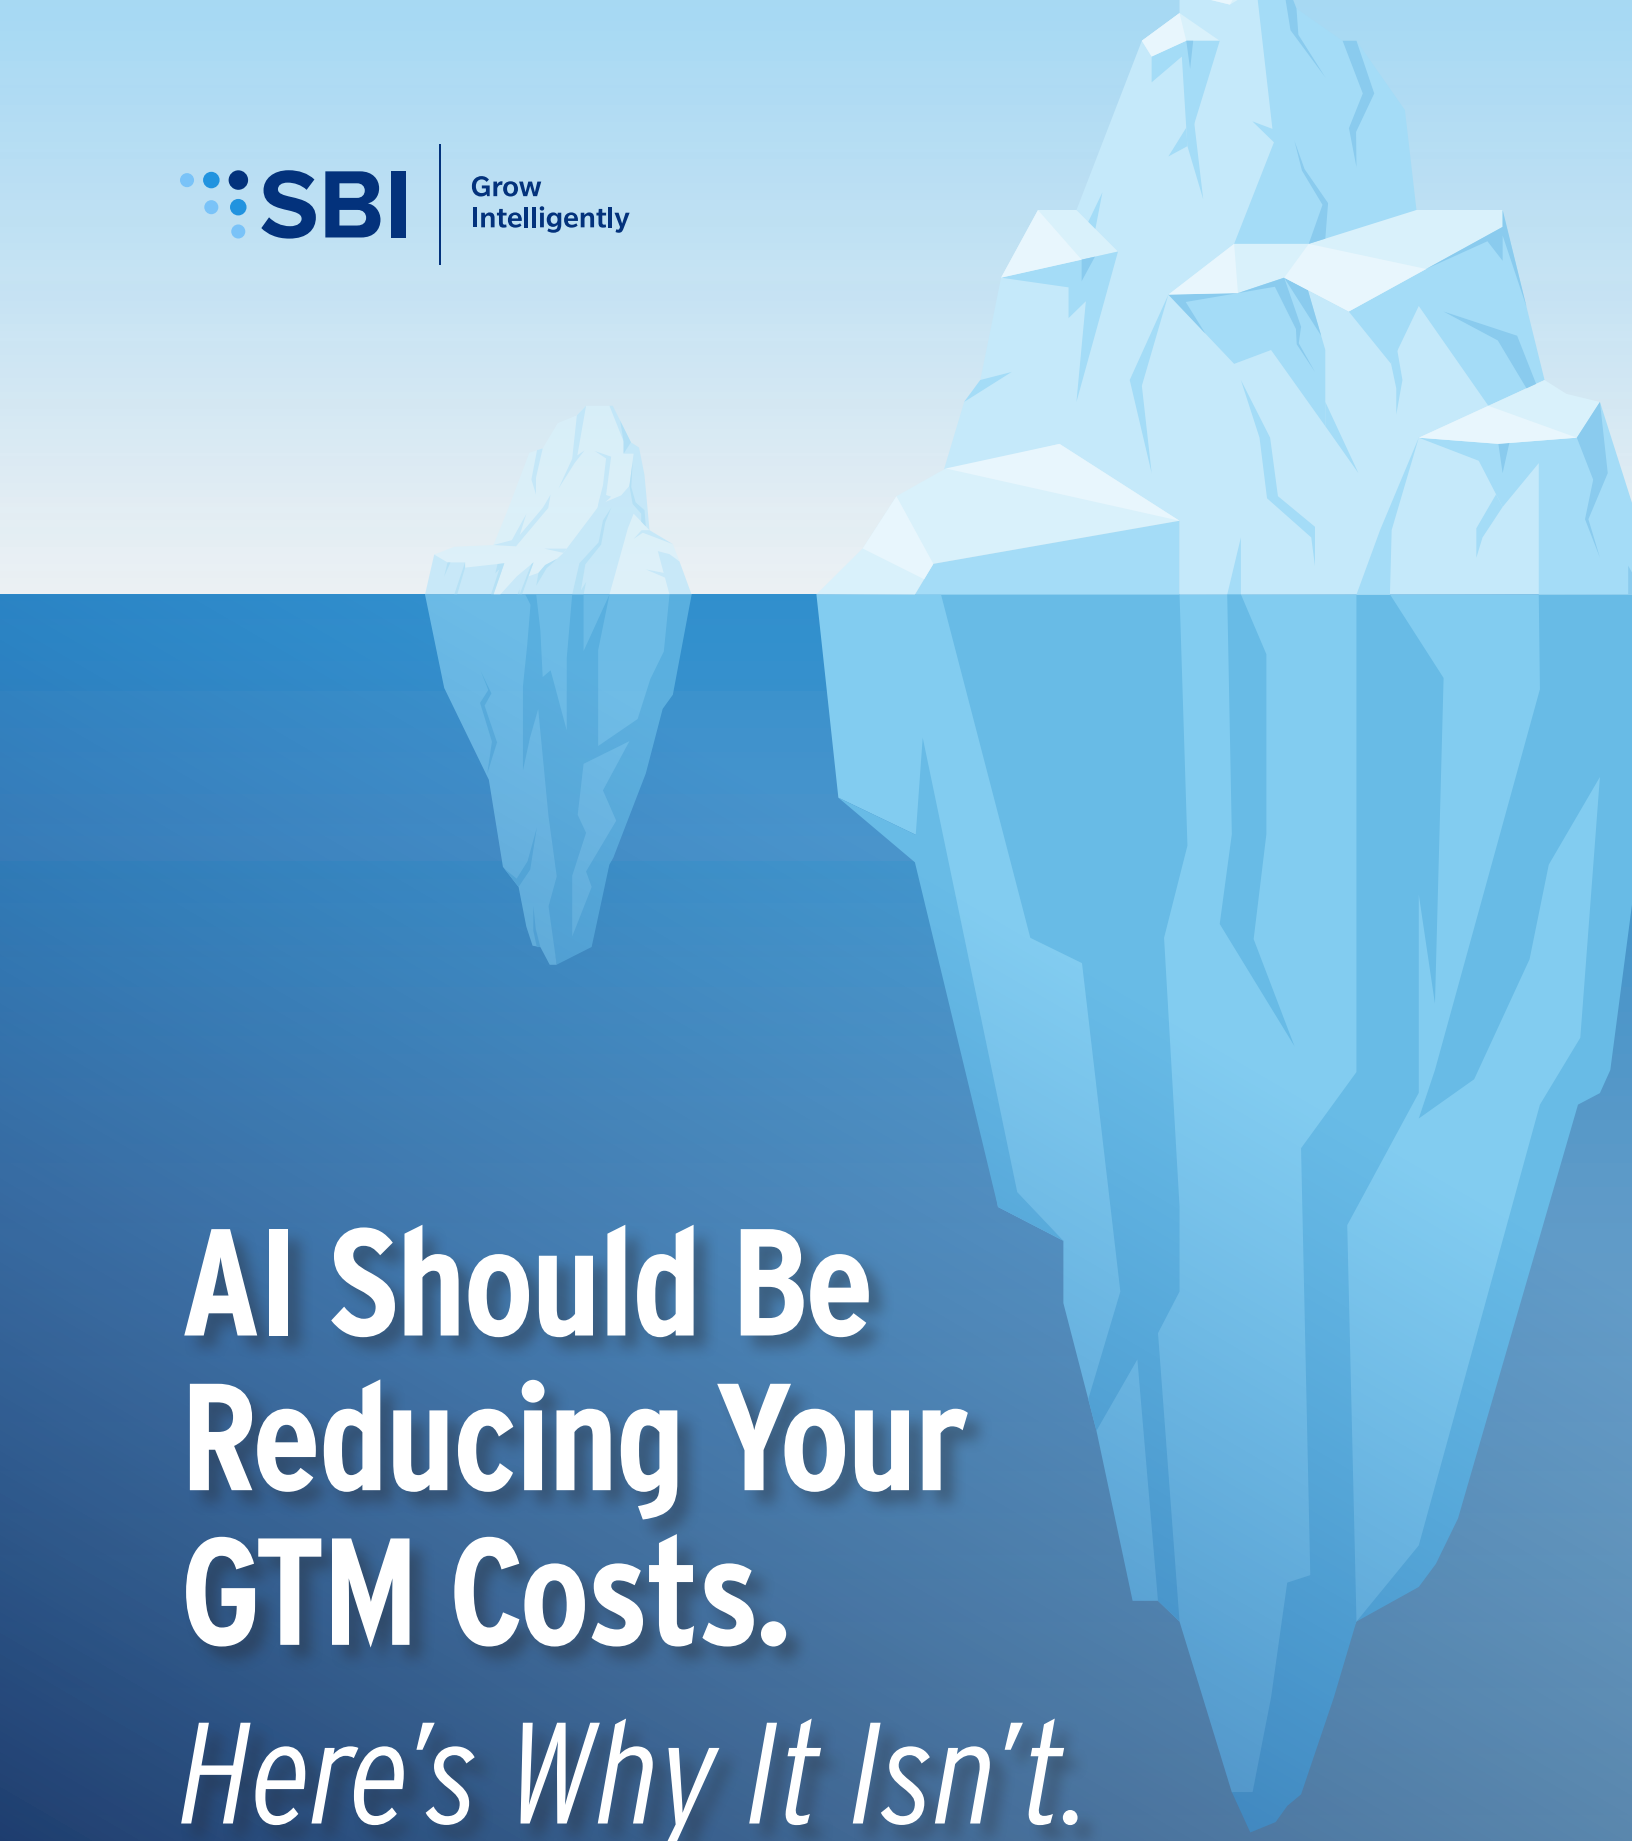

AI Should Be Reducing Your GTM Costs.

Here's Why It Isn't.

The Cost Structure AI Was Supposed to Fix

AI doesn't fix a broken GTM workflow. It speeds up the damage. That is what most SaaS companies have done, and it explains why AI spend is up, tool counts are rising, and GTM as a percentage of revenue has barely moved. Sixty-five percent of CEOs have staked their 2026 growth plans on AI-driven efficiency gains, yet only 18% of sales organizations report high ROI from AI, second to last across all business functions. The tools are being bought. They're not changing the cost structure.

CEOs reporting high/very high ROI from AI by function

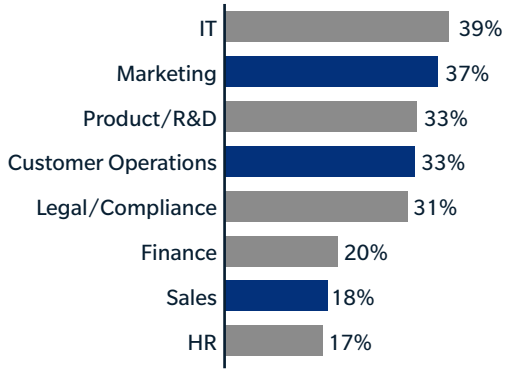

N = 118
Source: SBI Q4 2025 CEO Survey

Percent of CEOs saying customer-facing teams need AI efficiency gains to achieve 2026 goals

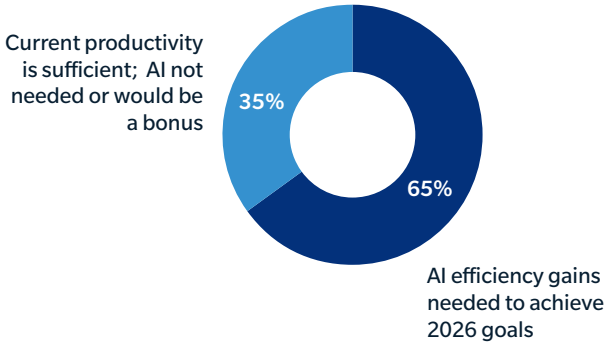

N = 118
Source: SBI Q4 2025 CEO Survey

The reason is that most companies are deploying AI to make existing activities faster without questioning whether those activities should exist at all. For PE-backed SaaS firms in the $50M to $200M ARR range, sales and marketing spend can run as high as half of every revenue dollar. AI isn't solving the efficiency problem; it's just adding a subscription to it.

GTM cost structure is one dimension of a durability problem that also runs through pricing, retention, and how AI is disrupting the commercial model itself. This piece addresses the GTM cost side.

The reason is that most companies are deploying AI to make existing activities faster without questioning whether those activities should exist at all.

AI isn't solving the efficiency problem; it's just adding a subscription to it.

The Manual Research Layer That Shouldn't Exist

Most SaaS GTM roles carry selling in their job titles, but a striking share of each workday never gets close to a buyer. SDRs assemble account context before writing a single outbound email. AEs wait for briefings before discovery calls. CSMs manually monitor health scores across spreadsheets trying to spot expansion signals before a renewal. Call it the GTM research layer. It's large, expensive, and nearly invisible to the P&L because it hides inside job titles that officially include selling.

Buying an AI tool to speed up that research doesn't solve the problem. It just makes the expensive part run faster. The research layer, headcount, and cost structure stay intact.

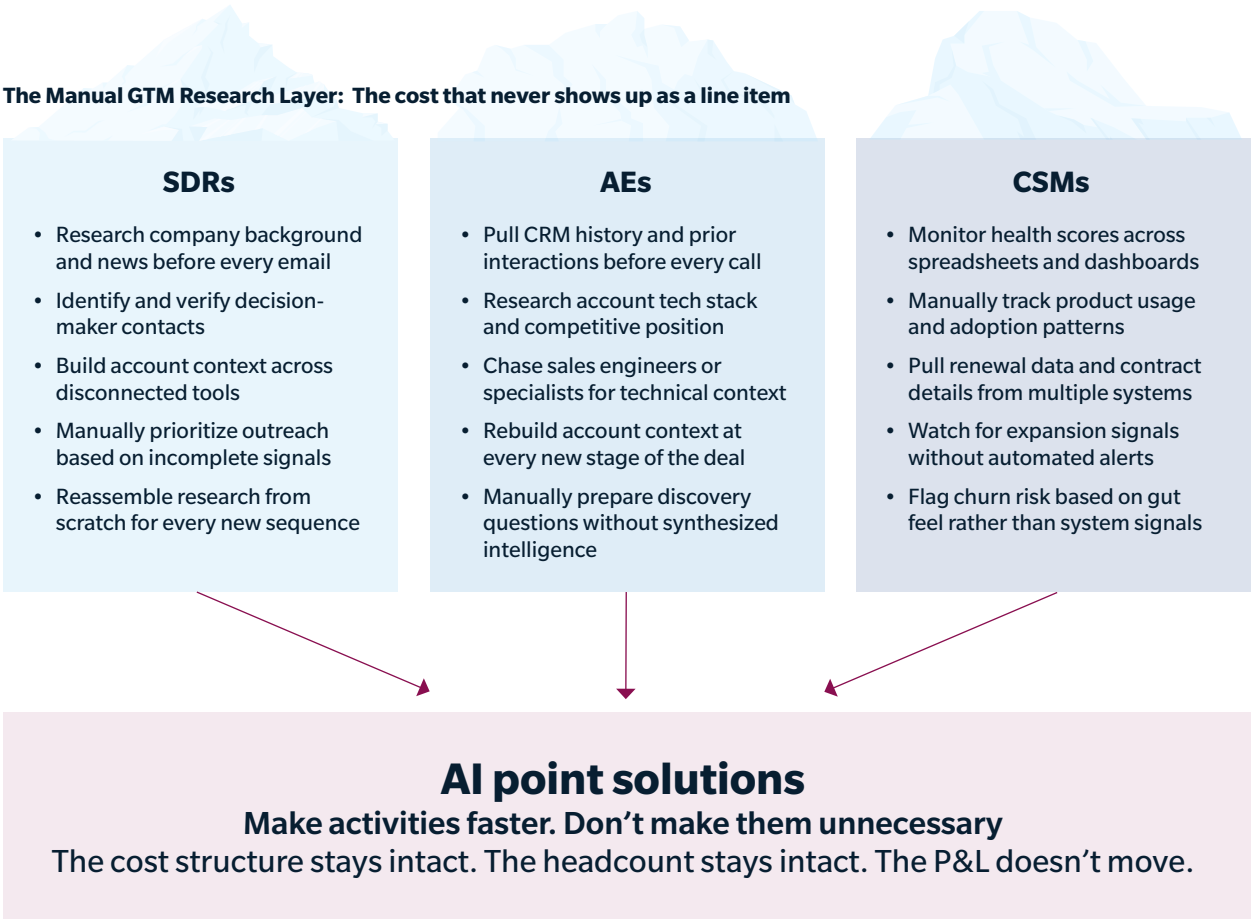

The companies that move the cost needle don't speed up the research layer. They get rid of it, replacing it with a system that does the synthesizing automatically before a seller ever opens a browser tab. Most companies don't get there because it means disrupting how teams work today, and that is a harder commitment than buying another tool. It is also the only commitment that actually changes the P&L.

Three Moves That Build a Compounding GTM Intelligence Asset

Too many companies skip straight to the part that looks like AI and wonder why the returns don't follow. The reason is that seeing meaningful AI cost reduction requires three moves that need to be executed in order. Done in order, you get a system that gets smarter over time. Done out of order, you get another set of tools that looked better in the demo than they perform in the field.

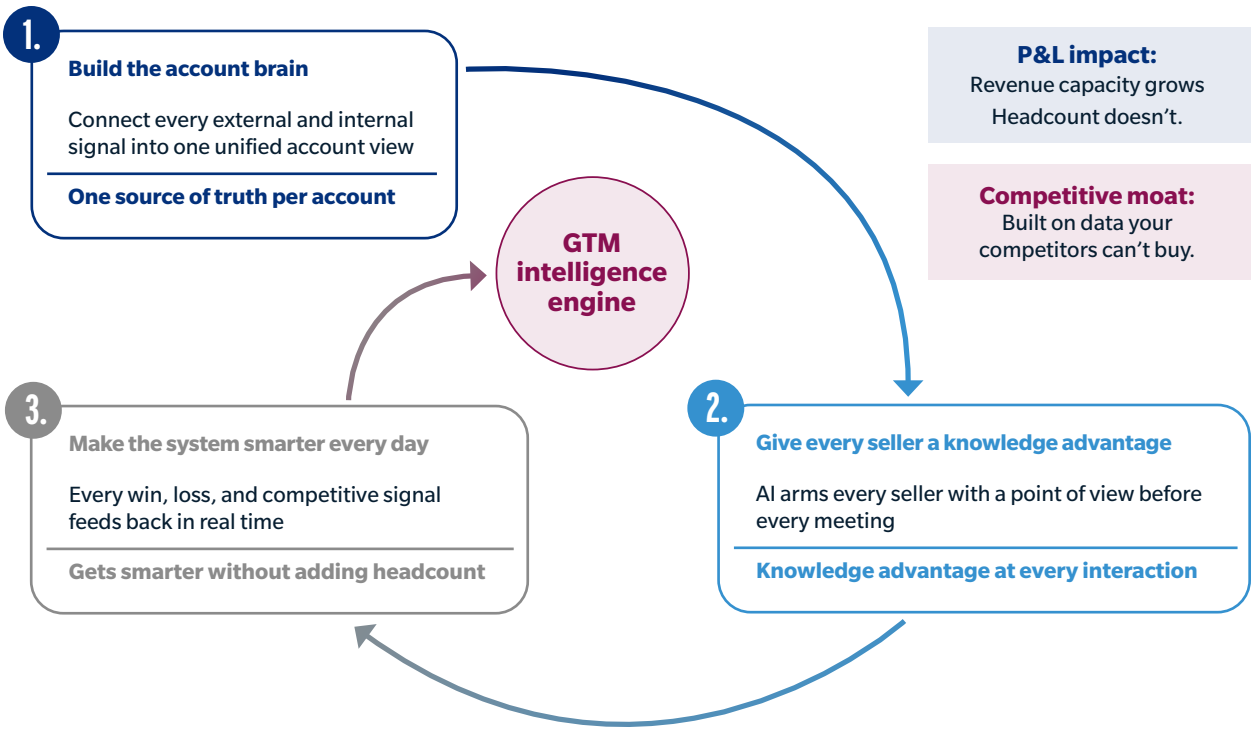

Move 1: Build the Account Brain

Most AI deployments in GTM fail before they start because the data feeding them is fragmented and incomplete. Signals sitting in six disconnected systems produce outputs that reflect that fragmentation. It's no coincidence that data quality and integration issues are the top barrier CEOs cite when asked why AI adoption in GTM isn't working, at 30% in SBI's Q4 2025 research. The tools get blamed, but the data architecture is almost always the real problem.

The first move is about getting everything the company knows about an account into one account brain: external signals from the market alongside the internal record of every interaction that account has ever had with the business. Most of this data already exists. It just sits in disconnected systems that have never been wired together. AI agents can now do that wiring in ways that would have required a dedicated analyst a year ago. The difference between a company that gets this right and one that doesn't isn't the AI model they're using. It's whether they've done the work to feed it something worth reading.

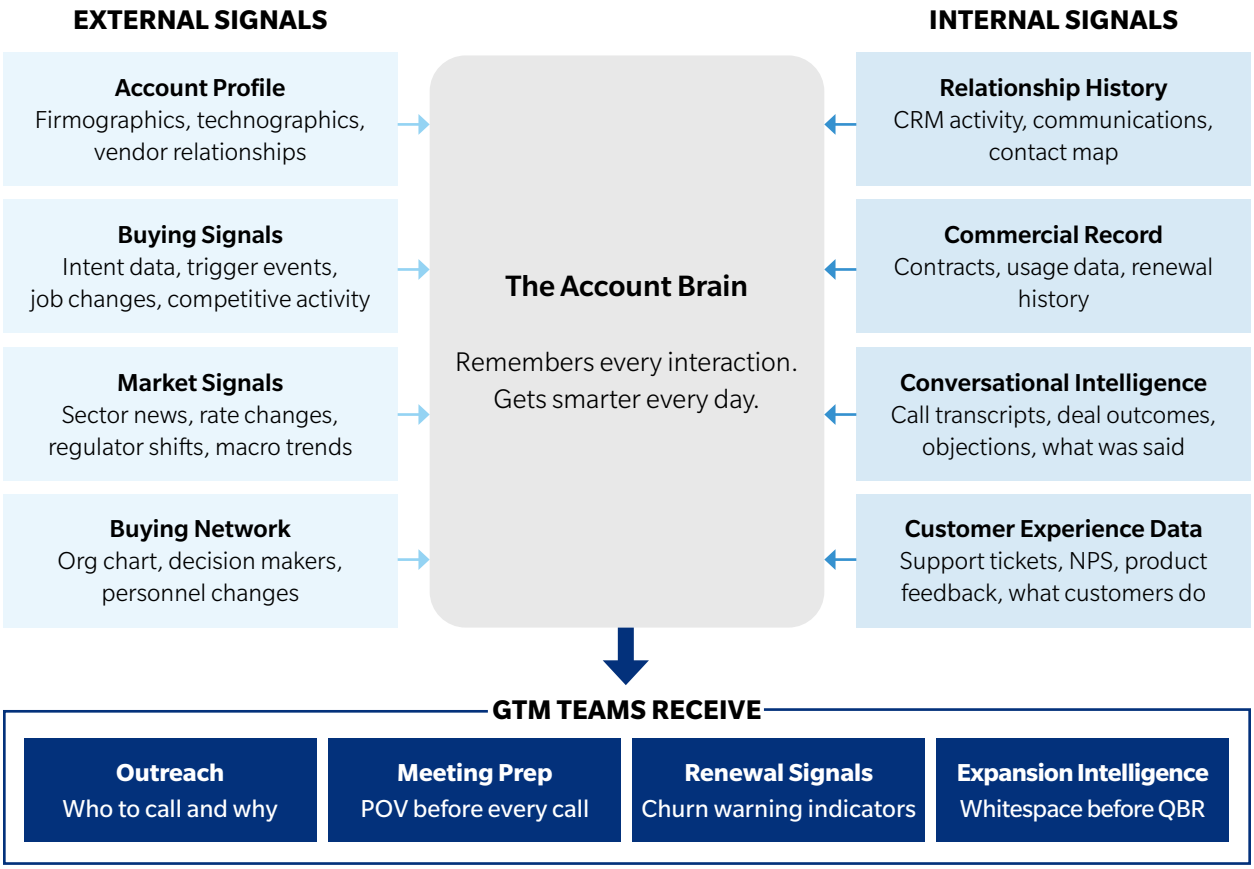

Move 2: Give Every Seller a Knowledge Advantage

A unified data pipeline is worth nothing if someone still has to spend hours reading it before a seller can act. That is just a more expensive version of the same problem.

The second move takes the account intelligence the pipeline has assembled and turns it into something a seller can walk into a meeting with: a point of view on what the account is dealing with, where the whitespace is, and who in the buying committee is likely to move. The seller's job shifts from gathering information to applying judgment, which is where the commercial value actually sits.

The practical difference shows up fast. SDR outreach goes from high-volume and generic to tightly mapped to the account's actual situation. A CSM heading into a renewal already knows whether the account has been showing churn signals for the past 60 days, so the conversation starts at a different place. An AM covering an expansion motion has the whitespace analysis before the call rather than after it. In each case the seller arrives better prepared than their counterpart at a competitor still assembling the brief manually. Whether that advantage translates to a leaner team or a more productive one at the same size depends on what the business needs. Either way, the cost of generating a dollar of revenue goes down.

Move 3: Make the System Smarter Every Day

This is what separates an ecosystem from a tool stack. Most AI tools produce intelligence that disappears the moment the session ends. The next rep, the next call, the next quarter all start from the same baseline. The system never gets smarter because nothing feeds back into a shared pool. Every tool you buy today is as good as it will ever be on the day you buy it. An ecosystem built on your own data compounds every time someone uses it.

The third move feeds what happens in the field back into the system in real time. An AE loses a deal because of a pricing objection, so that outcome updates how the system approaches every similar account still in the pipeline. A CS call surfaces a customer evaluating a competitor, so the system flags it before the renewal conversation happens and routes it to the right person. Instead of resetting at the start of each quarter, the intelligence carries forward and compounds.

Six months in, the system knows your accounts, your buyers, and your competitive landscape better than any individual on the team. Twelve months in, that advantage is genuinely hard to catch. Your competitors can buy the same AI models, the same intent data, the same email tools. They cannot buy your history.

What This Looks Like in Practice

Outbound	Retention	Expansion	Headcount Implications
BEFORE			
Large SDR team running high-volume sequences. Volume compensates for low conversion. High CAC.	CSMs monitor health manually. Warning signs missed before they become committed churn.	Overlay specialists, manual whitespace analysis. Cross-sell opportunities identified late or missed entirely. Every new product motion requires dedicated headcount to source and qualify.	GTM headcount grows linearly with revenue. Research and synthesis work hides inside every role. Adding capacity means adding people.
AFTER			
System flags top 15–20% of accounts in buying window. Smaller senior team converts at higher rate, lower CAC.	Ecosystem flags churn signals before renewals. Sentiment shifts surface automatically. CSMs cover more accounts.	System scans install base, identifies fit accounts, generates briefings automatically. Overlay headcount drops significantly.	Smaller senior teams. Research and synthesis roles redeployed. New roles, like forward deployed engineers emerge who build AI into commercial workflows.
Precision over volume	**Early detection, smarter coverage**	**System-sourced pipeline**	**Grow revenue, Not headcount**

The three highest-stakes GTM motions each look different when the intelligence system is in place. In outbound, the system identifies the 15 to 20 percent of accounts showing buying signals and builds the full account context before a seller engages, shifting the model from sending more emails to a large team toward sending the right ones with a smaller, more senior team. In retention, the system surfaces churn signals before they reach the renewal conversation so CSMs are intervening when it still matters. In post-acquisition cross-sell, the system scans the combined customer base, identifies the best-fit accounts for the new product, and generates the briefing automatically.

The EBITDA Case for PE Investors

The question for investors managing SaaS platforms is no longer whether AI can make sellers more productive. The question is whether the GTM cost structure can be fundamentally repriced, and the answer depends entirely on whether the company stops buying AI tools and starts building AI architecture. One produces marginal productivity gains that show up in activity metrics. The other produces structural cost reduction that shows up in the P&L.

A SaaS company running $400M ARR with S&M at 55% of revenue is spending $220M to go to market. Well-run PE-backed SaaS companies often operate at an S&M spend of 33% of revenue, which is $132M. The $88M gap between those two numbers is not a market condition. It is an operating choice, and it is made every quarter that the company keeps adding headcount to compensate for workflows that were never designed to scale. In SBI's Q4 research, Revenue Outperformers, the 23% of companies that exceeded their 2025 revenue growth targets, are nearly twice as likely as the rest of the market to name AI adoption as their top GTM talent priority. The financial case is straightforward: CAC payback periods compress, NRR improves because the retention model is smarter, and the Rule of 40 math changes in ways that matter at exit.

One thing that doesn't show up in deal models but consistently determines whether this works: technology is the easy part. Getting managers to reinforce new behaviors rather than defaulting to the old ones is what separates real cost reduction from cost theater.

SBI's Growth Durability Diagnostic evaluates where a company stands across four dimensions:

- Churn exposure and NRR risk assessment
- Pricing leverage analysis against key SaaS archetypes
- AI cost leverage and commercial productivity scoring
- Total growth durability scoring and a prioritized action plan, including immediate no-regrets moves

For CEOs it answers the question boards are already asking. For PE investors it provides a consistent framework for evaluating that risk across an entire portfolio.

Where does your commercial model sit?

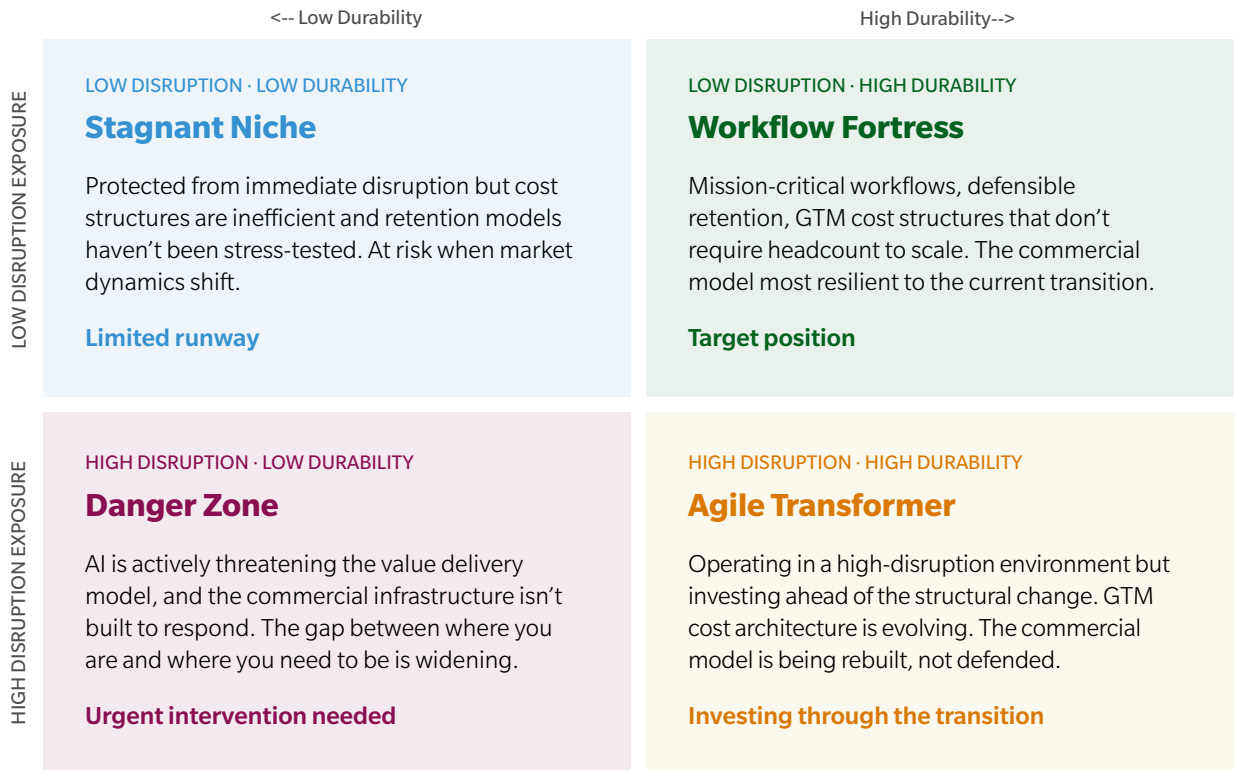

The starting point is a harder question than which AI tool to buy: whether the business is ready to rebuild the workflows that have always been there, always been expensive, and never been questioned.

SBI | Grow Intelligently

550 Reserve Street
Suite 190
Southlake, Texas, 76092
www.sbigrowth.com