The Post-Sale Operating Model Is Broken.

Here Is How to Rebuild It

Omid Razavi · April 23, 2026

The Post-Sale Model No Longer Fits the Market

Let me be direct: The post-sale model many SaaS companies still run is no longer fit for purpose.

It was built for a different era — one in which new-logo growth could conceal retention weakness, a QBR could stand in for proof of value, and "digital customer success" too often meant under-serving customers outside the high-touch model.

That era is over.

The model is now misaligned with what customers expect, what CFOs demand, what boards require, and what AI makes possible. Net Revenue Retention has become one of the clearest signals of company health, and the gap between top and bottom performers is no longer easy to ignore. The strongest companies grow through the installed base. The weaker ones are stuck on a treadmill, spending to replace revenue that should never have been at risk in the first place.

The leaders I speak with across the SuccessLab community already feel this. The real question is not whether change is needed. It is what must change, in what order, and how much urgency leadership is willing to bring to the redesign.

This article reflects what I am hearing from CCOs, post-sale leaders, and operators who are not trying to tune the old model. They are trying to replace it.

Three Forces Reshaping the Post-Sale Model

One pattern shows up repeatedly in my conversations with leaders across the SuccessLab community: post-sale teams are working harder, carrying broader books of business, and operating under tighter scrutiny, yet many still struggle to prove their value in terms the CFO cares about.

This is not a talent problem. It is not an effort problem. It is a design problem.

The assumptions underneath the traditional post-sale model no longer hold. Renewal decisions were expected to flow through champions who already believed in the product. Regular engagement was often treated as evidence of value. And AI was seen as something to pilot on the side, not a force that would reshape both what customers expect from vendors and what they may choose to build for themselves.

Three forces are now reshaping the model at once, and incremental improvement will not be enough.

First, the scrutiny standard has shifted

One Customer Success leader said to me recently, “We have plenty of customer conversations. What we do not always have is a shared and credible way to prove progress.” That observation captures a fundamental change in the renewal conversation.

The center of gravity has shifted. Renewals are increasingly shaped by the CFO, not just the champion.

The question behind most renewals has changed. It is not “Do they like us?” It is “Can they justify us?” And the person asking it is increasingly the CFO.

In many companies, SaaS renewals that once moved forward on the strength of sponsor confidence are now facing much greater scrutiny. Spend-management tools are surfacing underutilized software. Pilot-to-proof expectations are replacing renewals that are mainly grounded in trust and momentum. More customers are asking a harder question: what would we actually lose if we stopped paying for this?

Vendors that cannot answer that question with specific, quantified evidence of business value are losing renewals they once would have won. Not because the product has become weaker, but because the standard for justification has become much tougher.

Second, the build capacity of customers has materially changed

This is the force I believe the market is still underestimating, and the one most likely to reshape renewal conversations over the next two years. The data already points in that direction. Retool’s

2026 Build vs. Buy report found that 35% of teams have already replaced at least one SaaS tool with a custom build, and 78% expect to build more internal tools in 2026.

That should get every SaaS leader’s attention.

Tools from OpenAI and Anthropic have materially lowered the threshold for what a customer’s engineering, AI, or operations team can build on its own. Lightweight internal applications, workflow automation, and partial substitutes for vendor functionality are now far easier to assemble than they were even a year ago.

That is why build-versus-buy is no longer an abstract strategy discussion. It is a live question in the renewal room.

In several conversations I have had with customer leaders, the real threat was not a direct competitor. It was the customer’s own team, a director of AI, a product engineering lead, a RevOps architect, deciding that part of the vendor’s functionality could be built internally with enough control and at lower cost than the renewal price.

But that is only the first chapter.

A prototype built quickly with AI is not the same as a durable enterprise capability. Governance, security, testing, versioning, integration, maintenance, compliance, and operational continuity do not disappear once something works. They accumulate, and they usually become visible later — when the tool has to scale, survive audit, integrate with live systems, or be maintained by people other than its original builder.

This is why the issue belongs squarely in the renewal conversation. In an increasing number of accounts, the real alternative is not another vendor. It is the internal substitution. A customer does not need to replace the entire product to put the renewal at risk. They only need to believe they can build enough of it to question the price, narrow the scope, or leave. By the time that idea surfaces formally, the vendor is already behind.

Third, agentic AI is changing the economics of coverage — and raising the bar

For most of Customer Success history, one constraint shaped everything else: capacity. Teams had to decide where to direct human attention and which customers would receive meaningful coverage.

That constraint is starting to ease.

Agentic AI is changing what a post-sale team can reasonably cover. Systems can now monitor signals, summarize context, recommend actions, and execute parts of workflows that once required significant human time. The result is not that humans disappear. It is that strong teams can cover more accounts, more consistently, and with better timing.

And that changes the standard.

Once AI-enabled teams can provide faster, more continuous, and more tailored engagement across the customer base, high-quality support no longer needs to be reserved for a narrow slice of accounts. Teams still relying on periodic reviews, manual health score updates, and calendar-driven interventions are not just operating less efficiently. They are becoming less competitive by design.

Five Principles of the New Operating Model

What does the new model actually look like in practice?

Across the SuccessLab community, I see a growing divide between organizations that are redesigning post-sale to meet new realities and those still trying to squeeze more performance out of a model that is no longer fit for purpose.

The companies getting ahead of this shift tend to share five principles.

First, Value Proof Is an Operational Discipline, Not a Slide Deck

The best Customer Success organizations have stopped treating value demonstration as something that happens in a QBR. They have made it part of the underlying infrastructure of the relationship.

In practice, that means connecting product activity to business outcomes the customer actually cares about, in a form that both sides can access, revisit, and use throughout the lifecycle. The goal is not to assemble a persuasive story at renewal. It is to build evidence steadily over time.

That requires a more disciplined understanding of value. It is not enough to know what matters to the champion. The more important question is what the economic buyer or executive sponsor will need to defend internally. There is a meaningful difference between showing that 80% of users are active and showing that the platform reduced the customer's close cycle by three weeks. One is an activity report. The other is a business case.

The organizations doing this well start much earlier. They define the success plan before the contract is signed. They align on metrics before onboarding begins. And by the time the renewal conversation starts, the evidence should already be there.

Second, Signal Intelligence Replaces the Periodic Review

Too many post-sale organizations still depend on scheduled reviews to understand customer health. Teams wait for the QBR, the renewal checkpoint, or the executive escalation to realize that something has gone off track. By then, the room to recover is much smaller.

The better model is continuous signal reading. Product telemetry, support interactions, billing events, engagement patterns, stakeholder changes, and communication sentiment are brought together into a more complete picture of account health that updates as conditions change.

That matters because timing matters.

The gap between a 90-day renewal flag and a nine-month early warning is not incremental. It is the difference between an intervention that has a real chance to work and a conversation that mostly manages the exit.

It also changes how human effort should be deployed. If you know which accounts are drifting early, you can bring your strongest CSMs and leaders to the right accounts while the relationship is still recoverable, rather than after it has become fragile.

Third, Digital is a designed motion, not a fallback

Low-touch has too often been a polite way of saying under-served. In many organizations, digital Customer Success was built on a narrow premise: reduce human contact, lower the cost to serve, and accept a thinner customer experience as the tradeoff.

That premise no longer holds.

The strongest digital CS motions are not cheaper versions of high-touch. They are deliberately designed customer experiences, personalized by behavioral signals, triggered by usage events rather than calendar dates, and coordinated across channels.

The goal is not fewer CSM calls. The goal is a customer who has what they need to succeed, receives guidance when it matters, and experiences no meaningful gap in care, regardless of contract size.

What is changing now is that AI is making this standard more achievable. Agentic platforms are beginning to extend tailored support, answer routine questions, and help manage moments like adoption milestones, business reviews, and renewal preparation across a much broader portion of the customer base.

That is why digital can no longer be treated as a secondary motion. It needs to be designed as a core part of post-sale coverage.

Fourth, The CSM role needs a sharper definition

One of the most consistent observations I hear from CS leaders is that too much CSM time is still consumed by coordination work: preparing reports, updating CRM records, scheduling meetings, pulling data from disconnected systems, and writing summaries that AI can now generate almost instantly.

That is not Customer Success work. It is administrative overhead wrapped around it.

The new model requires a clearer answer to a question many organizations are still avoiding: where does a CSM add unique value, and where should AI or automation do the work instead?

The strongest answer is narrower and more strategic than many teams have historically designed for. The CSM's highest-value contribution sits in a relatively small set of moments: navigating complex stakeholder dynamics, leading strategic conversations with important accounts, helping recover relationships that are beginning to fray, and identifying expansion opportunities that require real judgment about the customer's business and direction.

Bain captured the tradeoff well in its 2025 research: "The real cost isn't just the hours spent on manual tasks; it's the opportunity lost." When CSMs are buried in coordination work, the real loss is not only productivity. It is the relationship, judgment, and commercial work that never gets done.

Fifth, Renewal Is a Lifecycle Outcome, Not a Commercial Event

The organizations getting ahead of renewal pressure are not winning in the contract's final 90 days. They are winning much earlier — when adoption patterns are being set, business outcomes are either taking shape or stalling, and executive relationships are being built or allowed to drift.

That is the right way to think about renewal.

By the time the formal renewal cycle begins, most of the real work should already be done. The customer should have a shared, visible understanding of what has been delivered. The executive sponsor should understand the return in the language of the business. And the internal champion should be able to make the case without needing the vendor to reconstruct it from scratch.

When a renewal conversation requires significant effort to prove value, the problem is usually not confined to the final commercial cycle. More often, it points to an earlier failure in the relationship — in onboarding design, adoption architecture, value proof, or stakeholder alignment.

That is where most post-sale organizations need to focus their improvement energy.

The Build-vs-Buy Conversation Needs a Better Answer

Build vs. buy is no longer a side discussion. It is entering renewal conversations across a growing number of SaaS categories, faster than many vendor CS organizations have recognized.

That means vendors need a better answer than most have today.

A credible response starts with four questions:

What can this customer actually build? AI capability varies widely. Some customers have the talent, infrastructure, and executive backing to build meaningful internal tools. Others have the interest, but not the ability, to take them into production in a durable way.

What would it really cost them to build and maintain it? Prototypes are cheap. Enterprise-grade capabilities are not. Governance, security, testing, integration, compliance, and support do not disappear once something works. They accumulate.

What does the product do that an internal build will not? This is not a feature-count argument. It is a capabilities argument: integration depth, data breadth, security posture, pace of innovation, ecosystem support, and the leverage of not having to maintain what you did not build.

Where is the internal build most likely to break down? Internal builds often look strongest at the prototype stage and weakest 12 to 18 months later, when scale, auditability, support, and maintenance become real.

Vendors that have seen this pattern across customers have a legitimate perspective to offer. But it has to be specific. Defensive positioning will not carry this conversation. A grounded, commercially credible point of view might.

Why customer AI capability matters before the renewal

One of the most important things a CS organization can do in a build-vs-buy environment is to understand, before renewal, which customers have real internal build capability and which have interest but lack the infrastructure to support it.

Those are very different situations, and they require different responses.

- Customers with real build capability usually show recognizable signals: technical teams already using AI coding tools at scale, experience taking internal software into production, executive sponsorship behind internal build efforts, and a use case that maps directly to a meaningful part of the vendor's functionality. These accounts need a serious renewal argument, not a feature tour.
- Customers with interest but limited infrastructure are different. They may build something impressive quickly, but they are far more likely to run into the production gap: the distance between a convincing prototype and a durable enterprise capability. That gap includes integration with live systems, role-based access controls, security review, governance, monitoring, and the practical question of who maintains the solution when the original builder moves on.

Knowing which kind of customer you are dealing with is not a nice-to-have. It shapes the renewal strategy itself.

Four dimensions of a credible build-vs-buy response

Most vendor responses to build-vs-buy are still too generic. They lean on feature breadth and integration depth without really engaging the customer's situation.

A credible response needs to do four things:

- Address full lifecycle cost, not just build cost. Building a prototype is cheap. Governing, securing, integrating, scaling, and maintaining an enterprise capability is not. The conversation has to move from prototype cost to total cost of ownership.
- Explain what the product does that an internal build will not. This is not a feature-list argument. It is a capabilities argument: data breadth, integration depth, security posture, pace of innovation, ecosystem support, and the operating leverage of not having to maintain what you did not build.
- Clarify the risk profile of internal builds in that category. Every SaaS category has its own failure modes. Workflow-heavy, configurable products tend to look easier early and become much harder later, when edge cases, integration drift, audit requirements, and support burdens begin to accumulate.
- Address the governance risk already forming inside the customer environment. What begins as a skunkworks project can quickly become a production dependency before the right controls are in place. That risk is now part of the conversation.

That is why build-vs-buy has become a more demanding renewal discussion. The vendor's case cannot be generic. It has to be specific, commercial, and grounded in the customer's real operating risk.

How to position the product credibly

The vendors that handle this conversation well are not the ones that dismiss the build option. They are the ones that engage it directly, understand the customer's technical reality, and make a specific case for why the product should remain a vendor relationship.

That case has two parts:

- The rational argument: total cost of ownership, capability depth, integration stability, governance infrastructure, and pace of innovation. This needs to be prepared in advance, grounded in the customer's actual AI maturity, and build appetite.
- The relational argument: the vendor's track record, roadmap investment in the customer's use cases, and a shared interest in making the product more valuable over time, not more substitutable. The relationship still matters, but as a complement to the rational case, not a substitute for it.

The bar for purchased software has gone up. The answer is not to argue for a lower bar. It is to meet it with a sharper case — commercial, specific, and credible in the face of what AI now makes possible.

Cross-Functional Coordination Is a Revenue Issue

The most consistent organizational gap I see in post-sale organizations is not a tools problem or a talent problem. It is a coordination problem.

The signals that matter most to renewal and expansion are spread across every function that touches the customer — Customer Success, Support, Product, Services, Sales, RevOps, and Finance — yet the ability to see those signals together and act on them cohesively still rarely exists.

That is not a minor operating flaw. It is a revenue problem.

A support pattern indicating a product issue never reaches the CSM before renewal. Product usage data that should reshape the commercial conversation sits in a separate analytics environment. An executive sponsor leaves, and the account team learns about it only when a meeting quietly disappears from the calendar.

This is no longer primarily a technology constraint. It is an organizational choice.

The new model requires three structural shifts:

1. A shared account intelligence layer Create a unified operational view that brings together CRM data, support signals, product telemetry, billing events, and stakeholder changes in a form every function can access and act on.

2. Explicit ownership of the renewal signal chain Someone has to own the twelve-month view of renewal health and be accountable for connecting the dots across functions. The title matters less than the visibility and authority to act.

3. Commercial alignment between CS and Sales In too many organizations, expansion is still treated as a handoff. The stronger model treats it as a coordinated motion: informed by Customer Success, carried commercially by Sales, and managed as part of a shared post-sale strategy.

EXECUTIVE ACTION AGENDA

What Leaders Need to Do Now

These are the decisions that separate organizations redesigning for what is next from those still operating on a model built for a different era.

Redefine the mission

Action required: Shift the post-sale mandate from relationship management to measurable value and renewal readiness.

Why it matters now: CFOs and boards are applying ROI discipline to every renewal. Relationships alone no longer carry the case.

Rework the coverage model

Action required: Rework segmentation. Decide where high-touch is warranted, where pooled coverage fits, and what digital should own.

Why it matters now: Many current models are too expensive, too broad, and too difficult to defend at scale.

Build a real signal architecture

Action required: Unify product, support, billing, CRM, sentiment, and stakeholder signals into one operating view of account health.

Why it matters now: Teams cannot intervene early on what they cannot see. Late signals produce late responses.

Sharpen the CSM role

Action required: Define where CSMs add differentiated value and move coordination-heavy work to AI and automation.

Why it matters now: When CSMs are buried in low-value work, the strategic and commercial work does not get done.

Design digital deliberately

Action required: Build digital CS as a purposeful engagement model, not a low-cost fallback.

Why it matters now: Customers now expect timely, relevant guidance regardless of segment or contract size.

Prepare for build-vs-buy

Action required: Assess each customer's AI maturity and internal build tendency, and prepare a specific response well before renewal.

Why it matters now: Internal substitution is now a real renewal risk, not just a product strategy issue.

Move renewal upstream

Action required: Treat renewal readiness as a lifecycle discipline built through onboarding, adoption, support, and value proof.

Why it matters now: By the time the formal renewal cycle begins, most of the outcome has already been shaped.

Lead the redesign personally

Action required: Treat the new operating model as an executive agenda, not a tooling or CS Ops project.

Why it matters now: Incremental fixes will not redesign a model that is structurally out of date.

QUESTIONS FOR THE EXECUTIVE TEAM

What Leadership Should Be Asking Now

These are the questions I would bring into any leadership discussion about whether the current post-sale model is built for the next three years or the last ten.

On value

- Can we prove value, or are we mostly reporting activity?
- Which accounts would accept a 20% price increase today?
- Where is our value story strongest — and where is it thin?

On coverage

- Where are we over-serving the customer base?
- What is our true cost to renew, by segment?
- If AI expands capacity, what should no longer require human coverage?

On signals

- How early do we truly see renewal risk?
- Are our critical signals connected, or still fragmented?
- Does our stack improve judgment and action, or mostly generate reports?

On build-vs-buy

- Which customers can credibly build around us?
- Are we addressing internal substitution early enough?
- What costs of building are customers still underestimating?

On the operating model

- What in our model was built for a pre-AI world?
- What needs redesign, not optimization?
- Are we treating this as a tooling issue when it is really an operating model issue?

CLOSING PERSPECTIVE

The future of Customer Success remains human

Let me close with a conviction I hold strongly.

The AI conversation in Customer Success too often swings between overstatement and resistance. Neither is especially useful.

The future of Customer Success is still human. The moments that matter most — recovering trust, resetting executive alignment, leading a consequential renewal discussion, identifying expansion in the context of the customer's business — still depend on judgment, credibility, and relationships. AI does not replace that.

It should reduce the administrative burden that keeps strong operators from doing their best work. It should improve the signals that surface risk earlier. It should strengthen the digital engagement that supports customers between human touchpoints. And it should make cross-functional coordination far less fragile than it is today.

That is the real opportunity.

Customer Success and Renewals do not need more activity. They need a better operating model — one that is clearer about where human judgment belongs, more disciplined about proving value, more intelligent in how it reads and acts on signals, and more honest about what AI and automation should own.

The companies that build that model deliberately will be in a much stronger position to retain customers, expand with them, and compete in a market that is becoming far less forgiving of weak post-sale design.

That work starts with leadership.

It starts with asking harder questions, making clearer choices, and redesigning the model with greater urgency than the market currently allows.

I hope this article contributes to that conversation.

Omid Razavi Founder, SuccessLab CCO Perspectives | April 2026

About SuccessLab SuccessLab is an executive community and advisory firm at the intersection of AI, customer leadership, and enterprise service strategy. Founded by Omid Razavi, SuccessLab organizes peer forums, roundtables, and the CCO Perspectives series for senior executives navigating the post-sale transformation underway in enterprise SaaS.

More perspectives and conversations: successlab.us