

Context in Motion .

B2B Customer Support Benchmark *Report 2026*

Foreword

In 2017, my co-founder Nitesh and I built the first version of what would become Hiver, because we couldn't find a tool that let support teams work the way they actually needed to. Today, over 10,000 businesses run their customer support on Hiver. And it's given us a front-row seat to how B2B support actually works across companies of every size and shape.

That seat has taught us that B2B support is going through more change right now than at almost any point we can remember. AI is reshaping what a support rep's day looks like. The line between support, customer success, and revenue keeps blurring further. We're shipping updates faster than ever, trying to keep pace with an industry rewriting its own playbook in real time.

Somewhere in the middle of that change, we found ourselves asking a simple question: where does this industry actually stand today? That question became the starting point for the survey of over 500 B2B support professionals, and eventually, this report. We wanted a genuine benchmark of how B2B support teams operate right now. What they know about their customers, how well that knowledge travels across their

organizations, how they're thinking about AI, and what keeps them up at night. We also wanted it to work as an early map of where the support function in such high-stakes environments is headed, built from the shared perspective of those leading it.

Some of what came back surprised us. There's a gap between how support functions get measured and what they're actually worth to a business. We're grateful to every leader who gave us candid, sometimes uncomfortable, answers. Their honesty is already shaping how we think about what we build next.

We intend to run this benchmark every year from here on. An industry changing this fast deserves to check in on itself at least that often.

We hope you find something in these pages worth taking back to your own team, just the way we did.



Niraj Ranjan Rout

Founder & CEO, Hiver

Introduction

The promise behind the *metrics*

Every metric a support team reports on is, in some sense, a promise. It means track this number, and it will tell you if the team is performing well.


After nearly a decade of working with high-stakes support teams, we know that most B2B teams measure the same things: response speed, resolution efficiency, and how much AI can take on. Useful numbers, no doubt. But we wanted to go further, to know what the numbers may be leaving out.

And that meant asking a newer set of questions. What does resolving a B2B issue actually require? What does a support rep see about an account the moment a ticket lands? Where does that understanding go once an issue crosses into

customer success, finance, sales, engineering? How does an organization prove what support metrics are worth chasing? And where is AI actually being asked to go next?

For answers, we went directly to the people living with these questions every day. **Over 500 B2B support professionals, from senior support managers to VPs of Customer Support, took part in our survey, drawn from companies of every size.**

The results reveal an industry that is not short of information. **Only 9% of teams say they have zero visibility into an account from within the ticket itself, and have to tool hop to find relevant information.** The other 91% say they have ready access to information.



Context, in other words, is rarely missing in B2B support. It's almost always in *motion*.

This might sound like a moot point, but it also raises a more interesting question: **what happens to that information as the work moves across systems that were built to track the ticket, and not necessarily the context and understanding that hold the ticket together?**

This report follows that movement across six areas: what resolution means beyond deflection, the account picture available at the moment of decision, what survives a handoff, how support proves its value, where AI is headed next, and, more importantly, what all of this asks of the people actually carrying it.

Chapter 1

Beyond *deflection*



What does resolution
actually mean in B2B?


The industry's most *tracked* number

In B2B support, deflection rate has of late become a loud load-bearing metric. Teams look at deflection as a proxy of efficiency; a number that tells them how fast, and how many tickets were closed without bandwidth drain.

Deflection rate rose alongside the wave of AI chatbots and self-serve tools that support teams adopted over the past few years. As ticket volumes grew and budgets tightened, deflection became the number that justified the investment in expensive AI and automation tools.

That's why 67.5% of B2B support teams track deflection rate in some form today. But the same data shows a split. **Only 24% of respondents call deflection their most important support metric. Another 43.5% track it alongside other measures.** Most teams have already, informally, demoted the metric, without formally retiring it.

How important is AI deflection rate as a KPI for your support team?



A question with no *single* answer

The clearest evidence that deflection alone isn't enough shows up in another question we asked:

In your view, is a deflected ticket the same as a resolved ticket?

33%

Yes - if customer didn't escalate, issue was resolved



37.5%

No - deflection just means customer stopped asking



24.7%

It depends on the query type



4.8%

Haven't thought about this distinction before



The picture that emerges is of a support function in the middle of separating two ideas that used to travel as one. For years, a ticket that didn't come back was treated as a ticket that had been handled well. The assumption held reasonably well in transactional, high-volume support, the kind built around password resets and shipping updates, where a silent customer usually meant a satisfied one.


B2B support leaders are now telling us that the assumption doesn't hold as cleanly for them. A customer who stops asking might have found their answer. They might also have given up, escalated through a different channel entirely, or decided the relationship wasn't worth the friction of asking twice.

Neither outcome shows up differently on a deflection dashboard; both count as a win, which skews support metrics in high-volume operations.

Why the calculation is *different* in B2B

To understand why deflection changes nature in the complex world of B2B support, you only need to look at where a support issue actually goes once it's raised. And this is what we found.

Which internal teams does support collaborate with most frequently to resolve tickets?



In order to resolve complex issues, **59.1% of respondents say their most frequent internal collaborator is customer success or account management, the function that typically owns the broader customer relationship.** That finding alone points to an important difference in B2B support: resolution often involves people beyond the team that receives the ticket.

A billing dispute may need inputs from Finance. A question about a platform feature may require Product or Engineering help. A support interaction may surface an account concern that Customer Success or Account Management is already managing. The ticket may begin in support, but the situation it represents can extend well beyond it.

This finding has a *direct consequence* for deflection as a metric

Deflection only measures whether the conversation reached a human agent or not. It has no way of seeing what happens after that conversation ends - if a case quietly moved to account management, or if the same concern surfaced two weeks later in a renewal call. It cannot tell you if the "resolved" ticket was actually the first thread of a longer, cross-functional conversation.


This does not make deflection *a wrong number* to track

All of this does not automatically mean that deflection rate is not an important signifier of how successful the automation process is. It's just that it is an incomplete signal on its own for organizations where resolution regularly depends on multiple people and processes.

Very clearly, the ticket is not the unit that matters the most in high-stakes support anymore. **The more meaningful question is what happens after a ticket is deflected.** Was the customer's question actually answered? Did the customer consider the issue resolved? Did they have to come back, reach another team, or find another channel to get help? Did the issue eventually surface as an escalation, a renewal concern or a churn signal?

None of this means abandoning deflection. It means pairing it with the outcomes that give it meaning: **resolution, satisfaction, recurrence, and retention.**

The metric is not simply whether a ticket was deflected. It is whether the customer's problem was *resolved*.



“

“Deflection tells you whether a customer reached an agent, not whether they reached an answer. In B2B support, an unresolved issue can disappear from the queue and resurface with an account manager or during a renewal conversation. The biggest risk is mistaking less visible demand for less customer effort. Deflection is useful, but it becomes a measure of success only when you can show that customers are getting their problems solved.”




Sarah Caminiti

Principle Consultant,
Epochal Operations

Chapter 2


The *account* behind the ticket



What a support rep *actually* sees

To understand how context travels, it is important to know what information arrives with a ticket and what needs to be hunted for in other tools. We asked support teams what a rep can see the moment a ticket lands, before they start to tool hop. **Only 9% said the rep sees nothing — everyone else can see at least some context immediately.** The other 91% see different combinations of customer context the instant a conversation begins.

When an agent picks up a ticket, what customer context is immediately visible without opening another tool?



Multi-select question. Percentages are the share of respondents selecting each option and do not sum to 100%.

Two figures land close together here. 39.7% of support reps can see open issues or escalations on the account, and 39% can see full conversation history across channels. That tracks with what helpdesk tools are built to do natively — log conversations and track issue status.


Move from conversation details to account-level facts, and the numbers drop. 31.4% of reps can see account health or sentiment from past interactions. 30.9% can see CRM data such as account owner, tier, or renewal date. 25.4% can see revenue or contract value.

The rep almost always knows what was said in a ticket. They far less often know what it means for the account: whether this customer is a flight risk, whether their contract renews in three weeks, whether this is their fourth escalation this quarter or their first. That's the information that would tell a rep how much a ticket actually matters, commercially and relationally — and it's also the information least likely to be sitting in front of them.

The cost of *hunting* for customer intelligence

We asked how many tools, on average, a support rep needs to access to resolve a complex ticket.

On average, how many tools does an agent need to access to resolve a complex ticket?



Put together, *two out of three* complex tickets (66.8%), require a support rep to leave the primary tool and look in at least three different places before they can act with full information.

The full account picture is scattered, unevenly, across unintegrated systems. And when a support rep on a live call has only seconds to decide how much of that scattered picture is worth chasing down before responding, this lack of integration can end up causing much harm.

Where the *real opportunity* lies

The data points to a clear priority. Conversation history and open issues are already visible to roughly 40% of reps — pushing that number higher has limited upside from here. The bigger gap is in account health, CRM data, and contract value, all still stuck between 25% and 31%. This is the information that actually determines how a ticket should be handled: whether it needs urgency, escalation, or a different tone entirely.

But that account picture is only useful while it's attached to the ticket. The next question is what survives when the ticket has to move to whoever needs to weigh in next.

"Of all the things a customer service agent can be, there may be nothing more powerful than the role of a 'proactive guide'. So often the customer does not know the questions to ask...or what they will experience in the future. The service agent likely does, as they are knowledge experts when it comes to navigating specific challenges. By equipping agents with the right data and context, they can do so much more than resolve a ticket, they can guide a customer to success and earn lasting loyalty."

“



Nate Brown

Co-founder and Executive Director

Chapter 3

How context travels when a ticket crosses *teams*


What happens to the account
picture once it moves?

How the handoff actually *happens*

A ticket that needs another team's input needs all the account information gathered up to that point to travel with it, intact, to someone seeing the situation for the first time. How that actually happens, in practice, turns out to matter as much as whether it happens at all.

We asked how coordination typically happens when a fast answer is needed from another internal team.

When you need a fast answer from another internal team, how does coordination typically happen?



For cross-functional coordination, Slack and email edge out the helpdesk. Add phone calls and conversations in person, and **58.2%, or more than half, of all teams coordinate outside the helpdesk.**


Where the hand-carrying breaks down


We asked how often context gets lost when a ticket moves across teams. **Just 3.1% of B2B support teams said they never lose context during a handoff — which means 96.9% experience some context loss at some point in their day-to-day operations.** More strikingly, 37.6% said it happens frequently or almost always.

Context loss during a handoff isn't a rare failure mode in B2B support. For nearly all teams, it's a regular occurrence — and for more than a third, it's close to routine. Traditional helpdesk tools, built to carry a conversation from a customer to an agent, were never designed to carry that same context from one internal team to another.

This is a core finding: context that's visible inside a ticket routinely fails to survive the ticket leaving support. For nearly all teams, that loss happens at some point. For more than a third, it happens almost every time.

How often does context get lost during a handoff or escalation?





The two findings taken together do not mean that context loss should be treated as a training issue. It is not human error alone that can be solved by asking reps to be more diligent about writing complete notes. The issue isn't one of discipline, but rather of design.

The systems support teams use to coordinate with the rest of the business were not built for the complexity of moving context across teams, tools, and channels.

Chapter 4

What support *knows*, and what the business can *see*



When customer intelligence
becomes a business metric


Close to the customer, *distant from the decision*

Support sees account risk first. A complaint pattern, a shift in tone, a customer going quiet — all of it shows up in support tickets, often weeks before it reaches a renewal conversation or a churn report. That much isn't new.

What's less obvious is whether that knowledge ever leaves support. So we asked whether organizations have a formal way to connect support outcomes to commercial outcomes like churn, renewal, or expansion revenue.

79.1% of B2B support organizations either already do this or are actively building toward it. **Support isn't struggling to prove its relevance — most of the groundwork already exists.**


Does your company have a formal way to connect support outcomes to commercial outcomes like churn, renewal, or expansion revenue?



Where the measurement still *lags* in reality

That said, connecting support to revenue in principle is different from measuring the specific intelligence support actually collects.


Which metrics does your team actively track and report on?



The pattern is consistent. The metrics with the highest adoption are the ones that were easiest to measure before AI or account intelligence entered the conversation: speed, volume, compliance. The metrics closest to what actually matters in B2B - whether AI is giving reliable answers, whether a particular account's ticket volume signals a bigger problem often go unnoticed. This is in no way evidence that leaders don't care about the right things; it could mean the tooling to measure them is newer and hasn't quite caught up with how quickly the work itself has changed.

So what does *leadership* actually ask about?

What is the one metric your leadership team asks about most often?



Revenue impact of support, the metric that ties support directly to commercial outcomes, is what leadership asks about least, at just 9%.

A further 3.6% say leadership doesn't ask about specific support metrics at all.

Set that against the earlier finding that 79.1% of organisations already connect support to revenue, or are building toward it, and a gap appears between what's being built and what's being asked.

Revenue impact is one of the hardest support metrics to produce on demand as it requires linking ticket-level data to CRM records, sentiment to renewal outcomes, and individual conversations to account-level risk. Response time and cost per ticket, by contrast, come straight out of a helpdesk with no extra work. Leaders are asking about what's easy to hand them an answer to, and not necessarily what matters most.

Why this matters more than it looks

If account health, sentiment, and relationship risk are the pieces of context hardest for a support rep to see in the moment, and if that same context is the first thing to get lost the moment a ticket crosses a team boundary, then it follows naturally that this is also the intelligence least likely to survive all the way to a leadership dashboard. **It's now also a measurement gap, compounding at every stage the information has to travel through.**

Making *customer* *intelligence* and insights available to every team

Most organizations are already building the connective tissue to prove that support has a bigger role to play in the organization's revenue and growth than just resolving customer queries.


The next phase of support measurement isn't only about adding another number to the dashboard. It's about recognizing the depth of customer intelligence that support already sits on, and making sure that intelligence travels, to product, to data, to leadership, to whichever team needs it to build something better.

Support sees the account first: the complaint pattern before product prioritises a fix, the frustration before a renewal is at risk, the shift in tone before it becomes a churn number. Right now, most of that intelligence stays where it's found, useful in the moment, but rarely carried to the room where the business decision it belongs to actually gets made.

Closing that distance is what lets support move from solving what customers are asking today to becoming the earliest, and clearest source of business transformation.

Chapter 5

AI beyond *deflection*

A glowing sphere with a rainbow-colored ring around it, set against a black background with colorful light trails.

What AI needs to
know to handle more?


Deflection measured AI's first job in support: answer the question, keep the ticket from reaching a human. The next job is harder to measure, because it isn't about whether AI can answer a ticket — it's about how much judgment to hand it, and where a human needs to stay in the loop.

There is no *single* model for AI

We asked support leaders which statement best described their current view of AI's role in support operations. No single position holds a majority. Some leaders are keeping humans firmly in the loop for every decision. Others are moving toward AI handling complex cases with minimal review. Most sit somewhere between the two — AI handling routine work, escalating anything that looks like it needs a person.


The industry isn't converging on one model. It's still working out where to draw the line between human and AI judgment, case by case. And the more judgment AI is asked to exercise, the more the information available to it starts to matter — an AI making a call about an account needs the same account context a support rep would need to make that call well.

Which statement best describes your current view on AI's role in support operations?



AI is not the only *priority*

What is your top priority for support operations over the next 12 months?



For most support leaders, expanding AI usage ranks third. The two priorities ahead of it — better visibility into account health, and fewer tools to piece together before acting — point at the same underlying problem this report has traced earlier. Account health and CRM context are the hardest information for a rep to see in the moment, because the systems that hold that information were never built to bring it together in one place.

That ordering says something. **Leaders aren't asking for more AI first. They're asking for AI, and every rep, to be able to see the account clearly before either one acts.**

How AI quality is (and isn't) being *tracked*

Only 2.9% of respondents formally track AI accuracy or hallucination rate (see Chapter 4). That doesn't mean teams don't trust their AI systems, or that they aren't measuring quality some other way. It does mean formal tracking of this specific measure is still rare — worth watching as AI takes on more complex, higher-stakes work.

The first generation of AI metrics measured how much work AI could absorb. The next generation needs to measure something different: whether AI made the right decision, whether it had the right context to make that decision, and whether it recognized when a case needed a human.

Each of those questions matters more as AI moves from answering routine questions to exercising judgment on harder ones.

The next AI *benchmark*

The findings do not point to one correct model for AI in B2B support. Some leaders are keeping humans firmly in the loop. Others are moving towards greater autonomy. Many are positioning AI between the two, handling routine work while escalating complexity. And it reveals that **more AI and better AI are not necessarily the same thing**.

Expanding AI usage is one priority. Improving account visibility and reducing the number of tools required to assemble that picture are priorities too. They are not separate problems.

An AI system can only make a decision with the context available to it. As the industry asks AI to handle more complex work, the quality of that context becomes part of the AI strategy itself.

The question is no longer simply how much support AI can handle. It is how much of the customer it can understand when it does.



“

“Most of the conversation right now is about how much AI should be allowed to decide. I think that's the wrong question. The right one is how much AI can actually see. Once it has real account context instead of general documentation, the right amount of autonomy becomes obvious, you don't have to negotiate it. That's where I expect the next year or two of progress to actually happen, not in bigger AI decisions, but in better AI eyesight.”




Christian Sokolowski

VP, Customer Support

Chapter 6

The people behind the *resolution*



What it takes to hold
all of this together?


The previous chapters have followed what happens to a complex customer issue as it moves through a B2B support organization: from deflection to resolution, from one team to another, from operational data to business reporting, and increasingly, from people to AI.

But there is one part of the system that does not appear in any workflow diagram: **the people making the decisions.**

So we asked support leaders how their role is changing, what their biggest challenges in people management are, and how confident they feel about what lies ahead.

A role that is *changing*

How is your own role evolving within the organization?




The responses point to a support leadership role that is expanding in different directions.

The near-even split between "becoming more strategic" and "doing more without more influence" suggests a role that is being asked to grow, sit closer to revenue conversations, absorb AI oversight, and hold account context together across teams, without every organization yet deciding to recognize that growth with the influence it should carry.

The experience that is hardest to *replace*

We asked support leaders about the biggest culture challenge facing their teams. Retention tops the list.

What is the biggest culture challenge in your support team right now?




As support work becomes more complex, experience becomes more than familiarity with a product or process. Experienced people build judgment over time. They learn which issues need escalation, which customers need a different approach, which patterns matter, and when a seemingly routine interaction deserves a closer look.

Some of that knowledge can be documented. Some can be surfaced through systems, while some still depend on the judgment of the person handling the situation.

More than a workforce metric, retention then becomes a question of how an organization preserves the knowledge its people have accumulated and also a sign of how much weight is resting on the people who have learned to carry it.

Confidence about what comes *next*

How confident are you in your current support setup to handle the next 2 years?



Support leaders are looking at the changes brought about by AI with more caution than certainty. **When asked how confident they are in their current setup to handle the next two years, only 28.3% said they were very confident.** The remaining 71.7% land elsewhere on the spectrum: 46.1% somewhat confident, 20.4% not very confident, and 5.2% not confident at all.


Fewer than three in ten support leaders say they're fully certain about where their function stands over the next two years.

That tracks with everything this report has already found. A function that names retention as its hardest problem, is still figuring out how AI upskilling should work, and is still learning to hold quality steady while it scales has little reason to feel settled about what comes next. The uncertainty isn't a sign that leaders have lost the plot — it's a sign they can see clearly how much is still in motion at once: AI, skills, quality, retention, all changing at the same time.

The question worth asking next isn't whether support can handle that change. It's what support actually needs in order to handle it well.

What people are still *holding together*

Across the findings in this report, one pattern keeps returning. Support teams have customer context, but it is distributed across different places. They coordinate across functions, but much of that coordination happens outside the system where the ticket began. They are building connections between support activity and commercial outcomes, but those connections have not yet become part of every leadership conversation. They are adopting AI, while still working out how much autonomy to give it and what information it needs to make better decisions.




And through all of this, *people*
remain responsible for connecting
the pieces.

The findings suggest that some of that connective work still depends on human judgment and experience. People fill gaps between systems, interpret information that does not arrive in one place, and make decisions when the available context is incomplete. And that reflects not a weakness, but how complex B2B support has become.

Chapter 7

Putting *context* to work



Where Hiver fits into
this picture?

The findings in this report point to a specific challenge. B2B support teams have context, but it rarely stays in one place as customer issues move across teams, tools, and channels. Pieces of an account visible in one tool, missing in the next, carried by hand across a handoff, and often lost somewhere along the way. This is the exact problem Hiver solves.

Hiver is an agentic customer service platform that brings teams, tools and customer conversations in one unified workspace. It gives support teams a shared view of the conversation rather than requiring context to be reconstructed across separate systems.

The same principle extends beyond the conversation itself. Hiver's Customer Intelligence brings together signals such as account health and sentiment from customer interactions, CRM data, calls, and product usage, so teams can understand the full picture behind the account and the issue they are handling.

When a case needs input from another function, teams can collaborate seamlessly without separating context from work. Internal notes, Slack conversations, and linked tasks stay connected to the underlying customer conversation across escalations and handoffs. The same applies to AI.

Hiver's AI capabilities are designed to move beyond simple deflection, with AI Agents able to work from a company's knowledge, tools, and customer context to handle more complex, multi-step requests, while guardrails define when a human needs to take over. AI Quality Assurance provides another layer of visibility by evaluating AI-assisted and human-handled conversations against defined quality standards.

These capabilities reflect a simple premise: **the more complex the work, the more important the context around the work becomes.**

Conclusion

The next *benchmark*


A mature B2B support organization in 2026 is not defined by how many tickets it closes, how much work it deflects, or how quickly it responds.

Those measures still matter. But taken together, the findings in this benchmark point to something broader: how well an organization understands its customers, carries that understanding through the business, and turns it into better decisions.

Five capabilities stand out.

- **Knowing the difference between deflection and resolution.** Deflection can show how much work AI absorbs without human intervention. It does not tell an organisation whether the customer actually got the help they needed. Mature support operations will increasingly need to measure not just how many conversations AI can keep out of the queue, but how many customer problems it can genuinely resolve.
- **Visibility into the account, not just the interaction.** A ticket contains the immediate conversation. The account contains the history, health, sentiment and commercial context that can determine what that conversation actually means. The more complex the decision, the more important that
- **Keeping context intact as work moves.** Customer issues rarely remain within one team. When support works with Customer Success, Finance, Sales, Product or other functions, the account picture and conversation history need to move with the work. The less an organization relies on people to reconstruct that picture at every handoff, the more consistently it can act on what it knows.
- **Turning support intelligence into business intelligence.** Support generates a close-up view of the customer that can inform decisions well beyond the support function. The opportunity is to connect that intelligence more clearly to the commercial outcomes leadership cares about, and make that connection part of the regular business conversation.
- **Giving AI the right context if you need it to do more.** The next phase of AI will not be defined by autonomy alone. It will depend on whether AI has access to the information required to make sound decisions, and whether the organization has the right boundaries for when human judgment should take over.

wider picture becomes.



The same movement sits beneath all of them: what support learns, and the intelligence tools hold about the customer, has to reach the people deciding what the business does next.

Resolution time, first response time, ticket volume: these are the metrics most support systems are built to measure. But the customer relationship is what support is actually managing.

How well an organization *carries that relationship's context from one interaction to the next* is likely to be the benchmark that matters most from here — and, increasingly, the thing that sets one support function apart from another.


Methodology

About the *research*


The B2B Customer Support Benchmark Report 2026 is based on two survey instruments fielded by Hiver between May and June 2026. The combined benchmark dataset represents over 500 qualified B2B support respondents. The survey explored how support teams handle complex customer issues, share context across functions, measure performance, use AI, and prepare for the future.


The core benchmark findings combine responses from both survey instruments.

Respondent job profile



Organisation size





How the data was analysed

The main benchmark questions combine the two survey datasets. Percentages were recalculated from the combined response counts. For questions allowing multiple answers, percentages represent the share of respondents selecting each option and therefore do not necessarily add up to 100%.



