


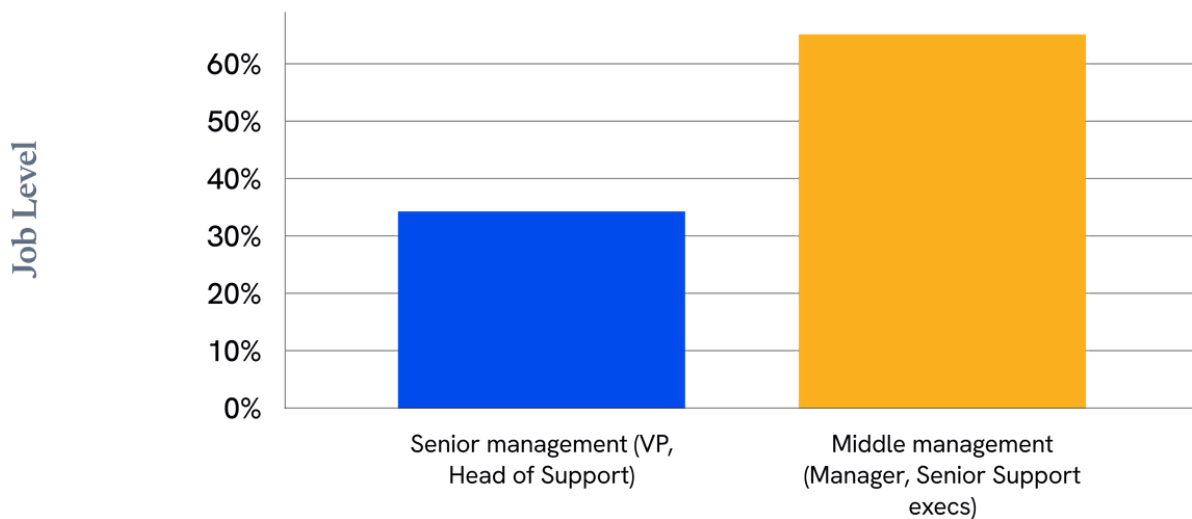

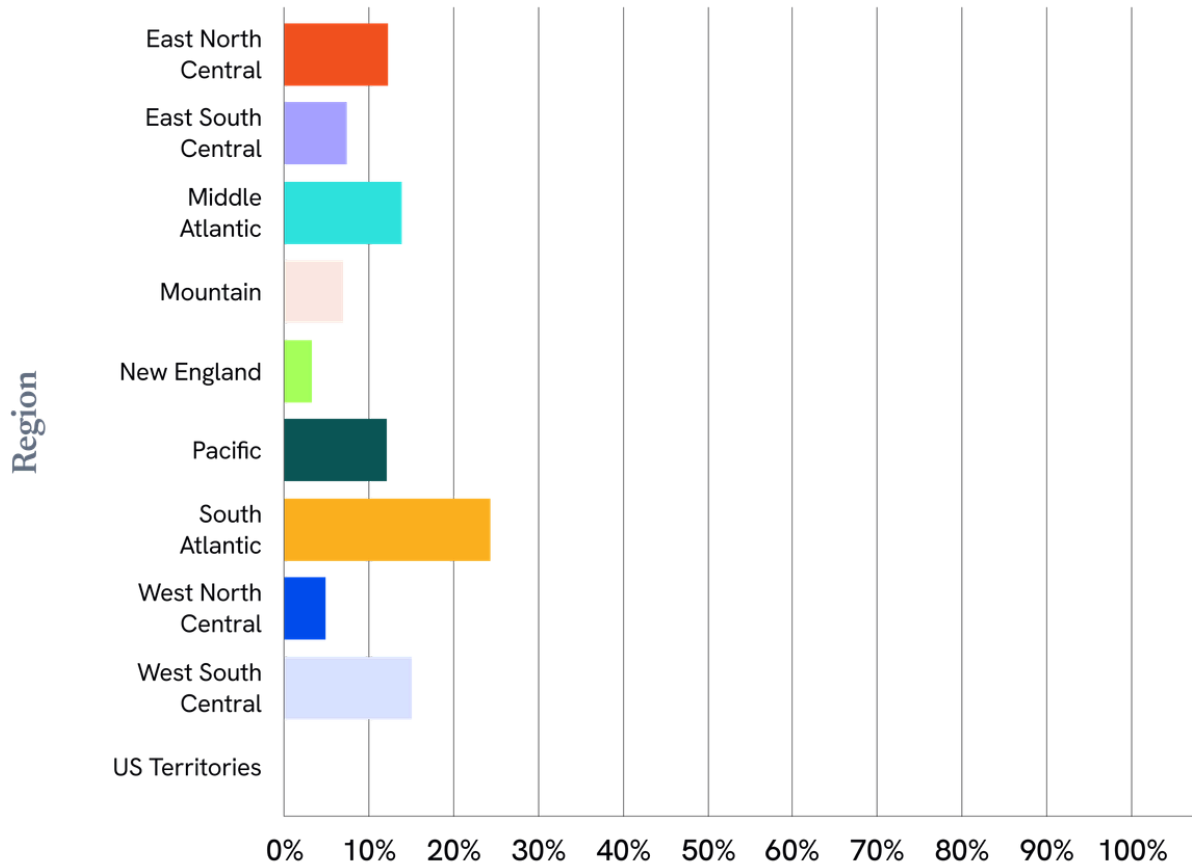


State of AI Customer Support in 2026

Based on responses from 700+ support leaders



Respondent Breakdown by Region, Age, and Role



Foreword

Across customer support, most teams are already using AI in some form. It might be for routing tickets, suggesting replies, or handling basic questions through chat.

What is still unclear, and keeps coming up in conversations with support leaders, is this:

- Is AI improving resolution times in a meaningful way?
- Is it reducing costs, or just shifting effort around?
- Can teams fully trust AI to do its job or is human oversight compulsory?

To move beyond assumptions, we surveyed more than 700 support leaders to understand not just whether AI is being used, but how deeply it's embedded into day-to-day operations; where it's delivering real value, and where it's struggling to scale.

The findings show an industry in transition. Even though AI adoption is widespread, confidence in its effectiveness is uneven.


And the gap between "AI is present" and "AI is operationally mature" is becoming the defining challenge of the next phase.

This report is designed to help you benchmark where you stand, and identify the gaps forming across the industry so you pressure-test your own approach, prioritize the right investments, and move faster than teams still stuck at the surface.

Explore the top three AI use cases that matter

Understand the hesitation around letting AI speak for the brand

Where support leaders are seeing ROI from AI



01

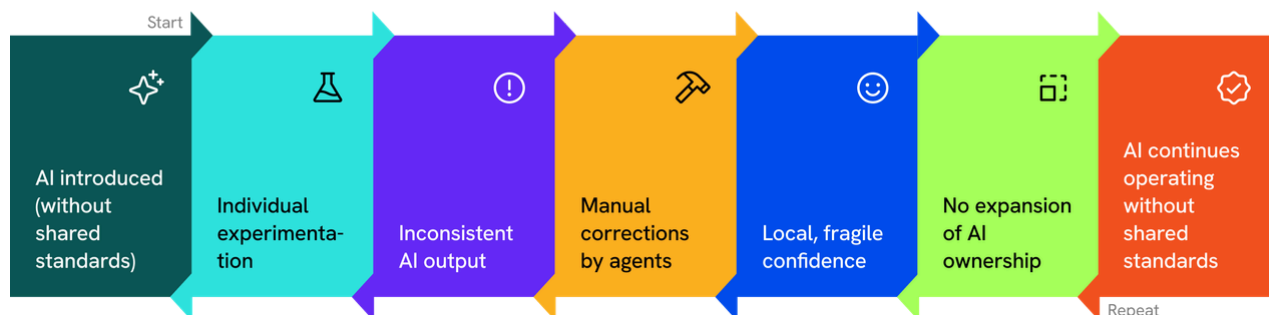
48% of Support Leaders Say
Their Teams Are Only
Moderately Confident with AI



For most teams, AI wasn't introduced with a playbook.

Instead, they are figuring things out in real time: editing AI replies, experimenting with prompts, adjusting workflows, and learning from mistakes.

How support teams experience early adoption



This loop explains why “moderate confidence” dominates. Teams are using AI daily, but without shared rules, confidence plateaus instead of compounding.

So how do teams move past “moderate confidence”?

From what we saw in our data, and heard directly from support leaders, confidence doesn’t improve by using more AI. It improves when the behavior of the AI platform or tool becomes predictable.

High-confidence teams don’t rely on trial-and-error learning. They remove uncertainty by standardizing how AI is expected to behave and how humans should respond when it doesn’t.

Two factors make the difference, every time.

- **Training that creates shared judgment (not just tool knowledge)**

Most teams introduce AI without defining what ‘success’ looks like, or when humans should intervene. As a result, every agent makes their own calls, confidence isn’t shared or standardized, and trust never compounds.

High-confidence teams do the opposite. They turn AI usage into a shared system of judgment.

Type of training	How it helps	How teams implement it
 Ai response standards	Removes tone and accuracy guesswork	Define what a “good” AI-assisted reply looks like for a small set of high-frequency scenarios, including tone, required information, and resolution steps
 Trust vs override rules	Reduces hesitation and over-editing	Define simple review checklists for policy, timelines, and commitments
 Core AI workflows	Builds confidence through repetition	Train teams on 3-4 repeatable workflows (drafts, follow-ups, QA checks)
 Edge cases & failures	Improves judgment, not blind trust	Review real tickets where AI’s input needed correction and discuss why
 Quality benchmarks	Sets a clear bar for “good” AI output	Tie AI-assisted replies to existing QA or CSAT standards

Karen Lam

Customer Support Director, Top Hat

“We talk about AI often, we share learnings, news articles, and ideas. The important part is being open, even when the conversations are difficult. While governance is important, AI is an exercise in change management so listening trumps any other strategy, unless that strategy is to spark innovation through showing potential and possibilities.”



- **Tools that reduce cognitive load**

The second confidence driver isn't training at all. It's how much mental overhead the tool itself introduces.

When AI is layered onto complex platforms with new interfaces, configurations, and workflows; teams need heavy training just to operate it. That slows learning and reinforces caution.

This skepticism is common among experienced support leaders.

Jordan Hooker

Customer Support Manager, Axuall

"I hear from countless companies every week claiming their AI will solve all my problems. The majority of them are snake oil. They're built on shaky foundations and don't actually do what they promise."



Tools that fit naturally into existing workflows do the opposite:

- Fewer concepts to learn
- Faster early wins
- Clearer mental models of how AI behaves

The simplicity of setup and workflow reduces training pressure, which allows confidence to build through usage instead of instruction.

The second confidence driver is just how easy and intuitive your customer support tool is.

Tools that require agents to learn new workflows or navigate complex configurations increase the need for training. Tools that fit into existing customer support workflows, on the other hand, are effortless to use on a daily basis.

One thing we've consistently heard from Hiver users is that the setup is extremely straightforward, and this enables them to effectively use the platform's AI features without any friction.

Nelson Jay Dela Cruz

Deputy General Manager, Miki Travel

"Every feature in Hiver has made a difference—from omnichannel ticketing and AI-powered live chat to workflows, AI Tagging, and analytics. Together, they've enabled us to scale seamlessly while delivering consistent, high-quality service to our customers."



02

Top 3 AI Use Cases in Customer Support: AI-Suggested Replies, Chatbots, and AI Quality Checks

39% Ai-suggested replies

35% Chatbots (First line)

30% Ai quality checks

Christian Sokolowski

Vice President of Customer Support, Rebuy

"We use AI for content suggestions — identifying repetitive questions we aren't addressing well, then creating documentation for them."



This approach helps reduce repeat inquiries over time by strengthening self-service and ensuring documentation evolves in step with customer needs.

What each use case is really solving

AI-Suggested Replies

The most straightforward use case of AI is drafting responses. Support teams use AI to handle common requests like billing questions, cancellations, delivery updates, access issues, and setup clarifications.

Instead of writing replies from scratch every single time, agents start with a pre-written draft that already reflects the right context, tone, and policy. They review, edit if needed, and send.

Chatbots (first-line)

Chatbots are great for answering repetitive queries. For cases that are too complicated and require human intervention, the bot gathers context and hands it off to a human agent.

This approach is already delivering measurable impact for support teams using chatbots as a true first line of defense.

Kacy Schwartz

Client Support Director, Splitit

“We built our own AI chatbot in-house. It handles the most common cases as the first layer of support, freeing up the team to focus on more complex and sensitive issues.”



Kacy Schwartz is the Client Support Director at Splitit, where she leads a global 24/7 support team serving over 450 fintech merchants. Her team built an in-house AI chatbot to handle the most common customer queries as the first layer of support.

Today, Splitit’s chatbot deflects **400+ tickets every month**, helping the team improve first-response and resolution times without compromising support quality.

AI Quality Checks

AI reviews responses the same way managers would, but a lot faster. It flags missing policy details, risky phrasing, incomplete steps, or timelines that don't match historical SLAs.

Where AI falls short:

But what's surprising is the scope of AI usage.

Only 9% of support leaders report that AI assists with most of their ticket volume.


For everyone else, AI's role thins out as tickets move from first response to resolution, follow-ups, and closure.

That said, a small number of teams are already doing something different. Instead of limiting AI to drafting or routing, they use it across their support workflow.

Giulio Castiglioni

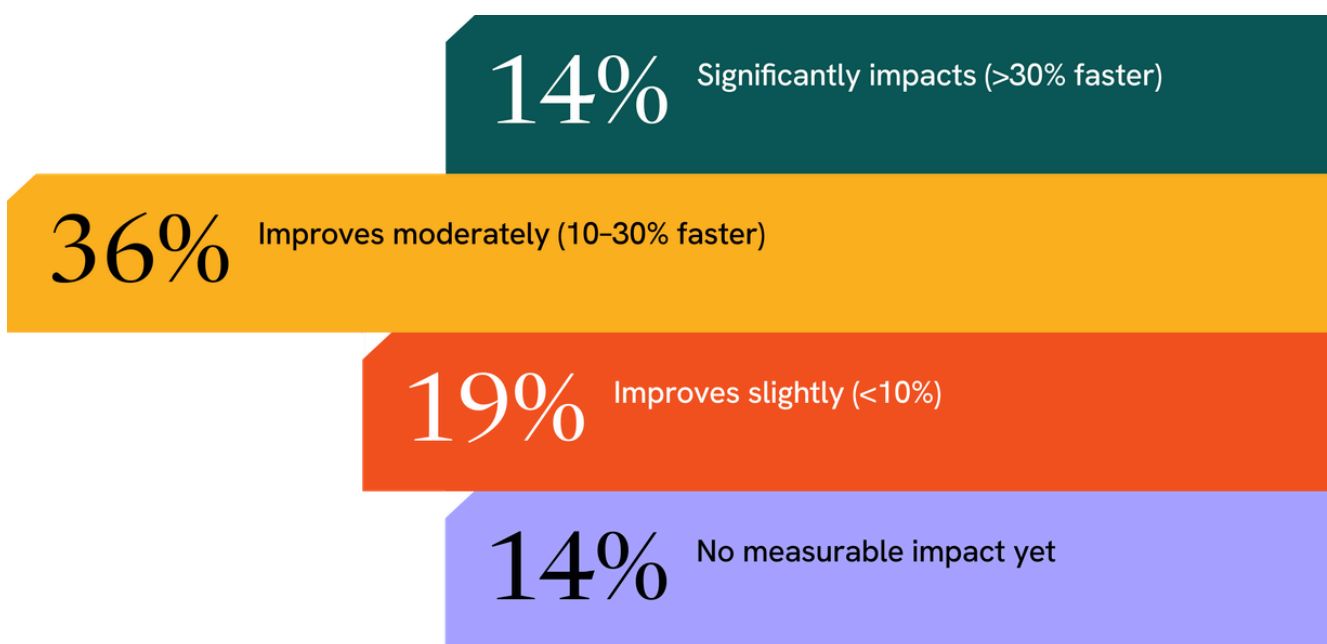
Customer Care Director, Playtomic

"Today, AI touches the majority of our inbound support volume, either by fully resolving requests or by shaping how tickets are classified, routed, and handled by humans. Its impact goes far beyond deflection and is embedded across the entire support lifecycle."



03

AI Improves Resolution Times for 55% of Teams, Mostly in Small to Moderate Ways



Resolution time isn't a single lever. It's the outcome of two very different forces working together.

One part of resolution time is repeatable work. The other is structural delay. AI is extremely good at the first. It has limited influence over the second.

Resolution time = Work the agent must repeat + Time the agent cannot control

Work the agent must repeat

A significant portion of time goes into work that is necessary but repetitive:

- Writing the same type of response repeatedly across tickets
- Searching multiple tools or documents for information that already exists
- Re-reading long ticket threads after a handoff or customer reply
- Correcting tickets that were misrouted or misclassified

This is where AI has the most impact. It helps agents regain context faster, draft responses more quickly, and avoid unnecessary rework.

Time the agent cannot control

This is the portion of resolution time that sits outside customer support's direct influence. It includes:

- Waiting for approvals
- Dependencies on other teams
- Refunds or exceptions processed elsewhere
- Product or engineering issues that stall progress.

AI doesn't eliminate these delays. It can surface context and reduce confusion, but it can't override organizational constraints.

There are two major ways in which AI impacts resolution times, according to customer service and CX expert, Shep Hyken.

Shep Hyken

CX Expert

"First, an AI-fueled self-service support system gets answers to questions or easily resolved issues quickly and efficiently.


Second, when AI supports an agent working with the customer, the agent has faster access to a more robust knowledge base, providing the information the customer needs, once again, quickly and more efficiently."



04

50% AI Scaling Failures Come Down to Trust and Context, Not Technology

Key barriers to AI adoption



We found the biggest barriers to scaling AI aren't missing features or limited budgets. It's **missing trust, missing context, and unanswered human concerns.**

Reported barrier	What this actually signals	Why it slows AI adoption
25% Resistance to change	Fear of disruption, role uncertainty, loss of control	Agents hesitate to rely on AI, managers limit rollout to “safe” use cases
24% AI responses lack context	AI isn’t grounded in customer history or prior decisions	Agents spend time correcting AI output, which reduces trust
17% No clear way to measure ROI	Leaders can’t connect AI usage to metrics they own	Hard to justify expanding licenses or use cases
11% Poor data quality or weak integrations	AI lacks access to CRM or historical ticket data	Responses feel generic or incomplete
11% Compliance, privacy, governance risks	Unclear rules around data access and AI behavior	Leaders fear reputational or legal risk
11% Budget constraints	Value hasn’t been proven at scale yet	Value hasn’t been proven at scale yet

When AI doesn’t understand customer history, prior decisions, or the nuance of a conversation, its responses feel generic. That forces agents to correct, rewrite, or second-guess what AI produces.

This extra effort doesn’t just slow work but instead, it undermines confidence. As confidence drops, resistance rises. Agents hesitate to rely on AI. Managers restrict usage to “safe” scenarios. AI becomes something teams tolerate rather than trust.

The Unspoken Concern Behind Resistance:

But there's a deeper layer underneath resistance that teams don't always say out loud.

For many agents, the concern isn't whether AI works — it's whether AI will eventually replace them. **Without clear guardrails, training, and intent from leadership, AI adoption can feel less like enablement and more like uncertainty.**

How teams can actively remove these barriers

01

Address fear before capability

- Tie AI usage to workload reduction, not headcount efficiency
- Involve agents early in defining where AI helps and where it doesn't

02

Ground AI in real context

- Integrate AI with CRM and historical ticket data
- Train AI on past resolutions and tone guidelines

03

Make ROI visible early

- Tie AI impact to outcomes teams already care about: faster handling time, fewer QA gaps
- Share small wins consistently instead of waiting for big transformations

04

Set clear guardrails

- Establish review checkpoints for sensitive cases
- Document governance rules in plain language, not policy jargon

05

Start small, then expand

- Apply AI first to repeat-heavy, low-risk workflows
- Let confidence build through usage before scaling coverage

05

60% of Support Teams Use Both Human and AI Metrics to Assess AI Quality

24.45% No quality evaluation in place yet

30.62% Combination of human + AI quality metrics

30.22% AI-specific metrics

14.71% CSAT, First Contact Resolution, QA score

60% of support leaders combine both to get a clearer picture:




Human metrics show whether the customer experience worked.
AI metrics show whether the system is working as intended.

Teams that rely only on human metrics often miss early problems. CSAT can look fine even when agents are editing most AI responses or avoiding AI altogether.

Teams that rely only on AI metrics miss customer impact. High accuracy rates don't guarantee that responses are helpful, complete, or appropriate.





Human-focused quality metrics

These measure whether the customer experience was memorable.

Metric	What it means
 CSAT	Did the response meet customer expectations?
 First Contact Resolution (FCR)	Was the issue resolved without follow-ups?
 Core AI workflows	Did the reply meet tone, policy, and completeness standards?

AI-specific quality metrics

These measure whether AI is behaving reliably inside the workflow.

Metric	What it means
 AI accuracy rate	How often AI suggestions are correct or usable
 Override rate	How often agents need to edit or discard AI output
 Error or hallucination rate	How often AI introduces incorrect or risky information
 AI usage rate	How often agents trust and actually use AI suggestions

How teams apply these metrics in practice:

Teams that use hybrid measurement don't build new reporting layers from scratch. Instead, they:

- Review AI-assisted tickets alongside regular QA reviews, not separately
- Track override and edit rates directly from agents
- Compare CSAT and FCR on AI-assisted vs non-AI assisted tickets
- Look for patterns where AI saves time, or creates extra work

The goal isn't to score AI perfectly. It's to spot where AI improves outcomes and where it needs adjustment.

Craig Stoss

VP of Solutions, Kodif

"AI Agent responses should be evaluated in a very similar way as you would QA a human agent's work. Is the message accurate, personalized to the user and contextual to the original question, does it have the right brand voice and tone, or follow style guidelines, and, most importantly, is it complete?"


"Answers with AI should not just regurgitate a knowledge article, or redirect users to read an article, it should solve the issue within the response itself. In addition, you should evaluate if the AI is "looping" customers (ex. "I don't know the answer, please rephrase"), sending generic answers without acknowledging customer context, or asking customers to repeat themselves. All of these would be unacceptable if a human agent did it, and so it should be unacceptable for your AI Agent too."



06

Around 66% of Support Leaders Report AI has Driven Some Change in Team Structure or Roles

When we asked support leaders whether AI has changed their team structure or roles:



That means two-thirds of teams are already seeing structural impact, even if the changes are still evolving.

Support teams are reorganizing and restructuring around AI, in two ways:

✓ **By creating AI-specific roles.** These roles stay embedded in day-to-day support operations, but their focus shifts from resolving tickets to training models, tuning responses, monitoring performance, and setting guardrails for how AI is used.

✓ **By building AI accountability into existing roles.** Agents and team leads are now expected to understand how AI behaves in real conversations, where it can be trusted, and when human judgment must step in.

The table below reflects this shift. It shows roles that are now appearing inside support organizations. Not as replacements for agents, but as extensions of the function.

Role title	Core responsibility	Key skills needed
AI Support / AI Trainer	Teaches and refines AI responses based on real support tickets	Prompting, ticket taxonomy, quality evaluation
AI Customer Experience Strategist	Designs and governs AI use to align with CX goals	CX strategy, data analysis, stakeholder communication
Conversational AI Designer	Builds dialogue flows and response logic for AI agents	UX writing, flow design, NLP basics
AI Data Analyst (Support Focus)	Tracks performance and error patterns in AI outputs	Data analysis, ticket tagging, trend spotting
Knowledge Manager / Content Curator	Maintains the knowledge base that feeds AI accuracy	Knowledge architecture, content management
AI Quality & Governance Specialist	Oversees safe, compliant AI behavior	Risk assessment, policy governance, compliance

Routine, repetitive ticket handling is no longer the primary growth path.

Instead, teams are investing in roles that:

- Improve how AI performs over time
- Reduce risk and inconsistency
- Ensure AI aligns with support goals and brand standards

Sarah Caminiti

Head of Business Transformation, SupportNinja

“We’re seeing the emergence of AI-literate support roles—people responsible for training, tuning, and auditing AI outputs. AI shouldn’t be eliminating roles because it’s really been raising the bar. Support teams now need fewer task-runners and more critical thinkers who know when not to trust automation.”



07

Nearly Nine Out of Ten Support Leaders Are Cautious About Letting AI Represent Their Brand

In customer support, representing the brand doesn't just mean using the right logo, tagline, or greeting.

It means making the right judgement calls that align with your brand values and principles while interacting with customers.

When AI responds to a customer, it is implicitly deciding:

- ✓ How strictly to apply policy
- ✓ How much empathy to show
- ✓ When to escalate vs resolve
- ✓ What tone is appropriate in a tense moment
- ✓ Whether to prioritize speed, accuracy, or reassurance

These decisions shape how customers experience the brand far more than any marketing message.

87% of support leaders said they are cautious about letting AI represent their brand in customer interactions.

That caution isn't about whether AI can follow rules. It's about whether AI can interpret **context and emotion**.

Cases where AI can get it wrong:



Rigid refund decisions

AI cites the policy correctly, but misses context like long customer tenure, prior issues, or repeat complaints. What should be a recovery moment can lead to churn.



Emotionally flat responses in crisis

AI responds quickly and accurately to messages about outages, billing errors, or service failures — but without providing the necessary reassurance. Speed replaces empathy where empathy matters more.



Loss of conversation context

AI answers the latest message correctly while missing earlier context. Customers are forced to repeat themselves.



Off-brand tone and voice

The AI-generated response is polite and factual, but doesn't sound like the company. Humor, warmth, or flexibility that customers expect is missing, making interactions feel generic.

Dan Gingiss

CX Keynote Speaker & Author of *Becoming The Experience Maker*

“AI is a powerful enabler - but not a replacement - for human empathy in customer support. The most effective teams use AI to eliminate friction and accelerate service, while reserving the final mile for a human. That’s where connection, loyalty, and remarkable experiences are created.”



08

AI Governance Is Human-Led Today: Manual Review (32%) Outpaces Automated Quality Controls (14%)

Top AI governance practices

36% Data privacy & security policies

32% Human review before sending responses

32% Post-response audits & quality checks

27% Training guidelines integrated into AI

As AI moves closer to live customer conversations, support leaders are being deliberate about how much autonomy they give it.

The data makes it clear.

This isn't a lack of trust in AI's capabilities.
It's a reflection of how support teams manage risk.

Data privacy & security policies — 36%

For many teams, this is the starting point. AI is only introduced into live workflows once basic safeguards are in place:

- Customer data is stored securely
- Access is restricted appropriately
- Systems comply with internal and regulatory standards

Human review before sending responses — 32%

This is the most common operational safeguard. Teams require human review for:

- Refunds and billing issues
- Complaints or escalations
- Sensitive or policy-heavy responses

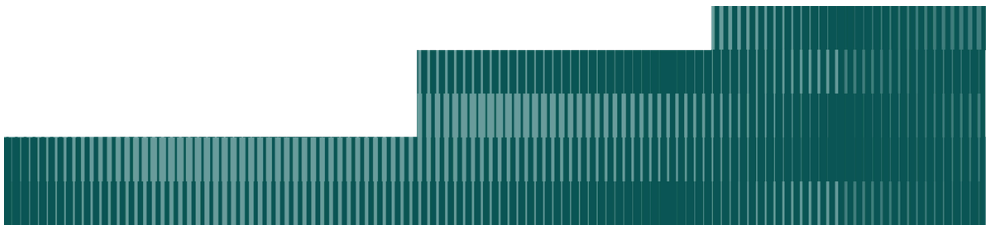
Post-response audits & quality checks — 32%

Governance doesn't stop at sending the response. Teams regularly review AI-assisted tickets to:

- Identify errors and overrides
- Spot recurring failure patterns
- Track performance trends over time

Training guidelines integrated into AI — 27%

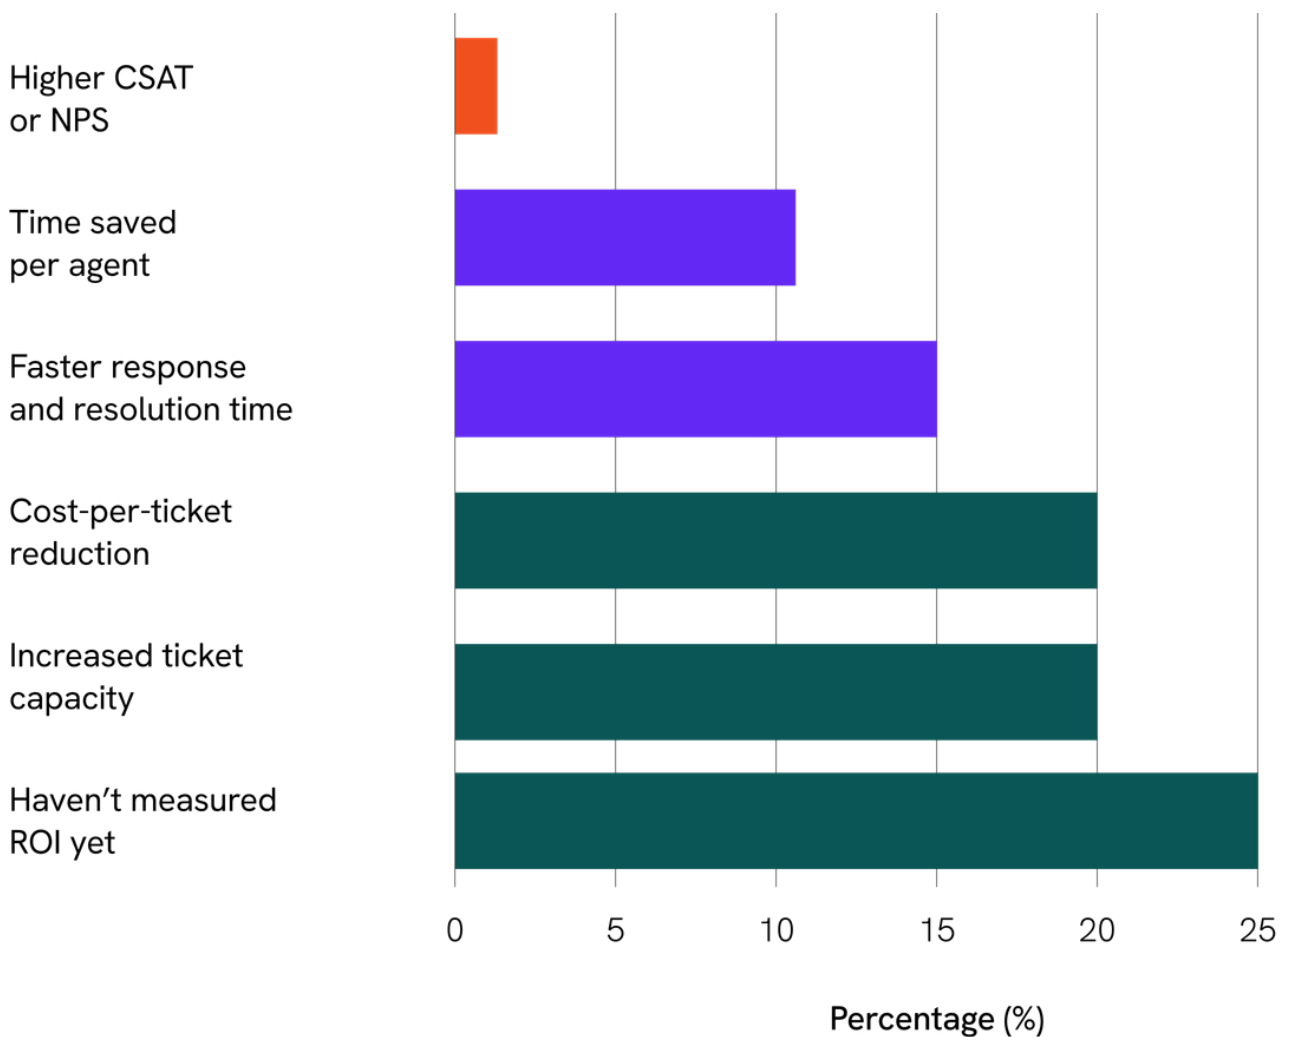
Some teams go a step further and encode expectations directly into AI behaviour:

- Approved tone and phrasing
 - Clear policy boundaries
 - Defined escalation triggers for sensitive cases
- 

09

Where Support Leaders Are Seeing ROI From AI

ROI indicators used by support leaders:



Three patterns stand out when support leaders talk about AI ROI:

1. AI is most often credited for operational gains, not customer metrics


Nearly 40% of leaders measure AI ROI through cost-per-ticket reduction or increased ticket handling capacity (both cited by ~20% each).

Another 15% point to faster response and resolution times as a core signal of value.

Christian Sokolowski

Vice President of Customer Support, Rebuy

“We use AI heavily across support to improve quality, insight, and overall customer experience. First, AI plays a big role in quality assurance for both human- and AI-handled conversations. Second, we use AI to analyze support content and identify trends in what customers need help with. Lastly, AI helps us continuously optimize our content by identifying gaps, improving clarity, and refining structure so our resources get better over time.”



2. Time savings exist, but they're not dramatic enough to stand alone as ROI

Only 12% measure ROI primarily as time saved per agent, suggesting that AI shaves effort incrementally rather than freeing up large blocks of agent time.

3. A significant group still hasn't formalized ROI at all

Despite these gains, 25% of support leaders say they haven't measured AI ROI yet — the single largest response in the dataset.

These three facts create a clear tension: AI is delivering value, but that value is distributed across workflows in ways that are hard to consolidate into one headline metric.

Why ROI feels real, but stays unofficial

Most AI gains don't trigger immediate structural changes.
Saving time per ticket doesn't automatically reduce team size.

Handling more tickets often prevents backlog rather than boosting output.

Faster resolution stabilizes SLAs without dramatically shifting customer scores.
So leaders experience AI as relief, not transformation.

What this means:


AI ROI in support is not absent. It's incremental, operational, and cumulative.

Right now, AI proves its value by:

- ✓ Absorbing pressure
- ✓ Reducing friction
- ✓ Making scale manageable

As teams mature their measurement frameworks, these signals will compound into clearer ROI stories. But at this stage, AI earns trust by keeping systems stable — not by delivering a single breakout metric.

That gap between everyday impact and formal ROI is exactly why AI adoption continues, even when measurement lags behind.





AI Adoption Was Step One. Context Is the Real Advantage

If this report points to one clear shift, it's this: AI adoption is no longer the story. AI maturity is.

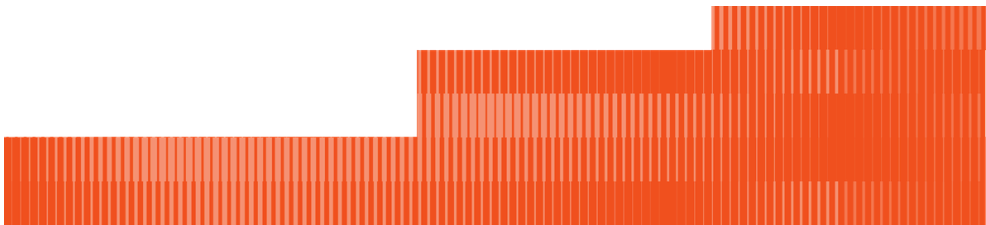
Most support teams are already seeing incremental gains from AI with faster replies, lighter workloads, modest efficiency.

But those gains plateau when AI operates without enough context to make reliable decisions. Speed improves, but judgment doesn't. And without judgment, trust never fully compounds.

The next phase of customer support will be defined by how well teams solve for context—bringing together customer history, knowledge, policy, and quality signals so AI can resolve work with confidence, not just assist with drafts.

The teams that get this right won't just use AI more. They'll run support differently.

We are already seeing this change in B2B support, where issues are often nuanced and heavily dependent on context. Chatbots alone tend to fall short in these scenarios.



How Hiver Brings Context-Rich AI Into Daily Support Work

A hybrid approach works best: AI surfaces relevant context such as past issues, documentation, and policies, while humans focus on troubleshooting, judgment, and resolving the problem.

This is why simply adding AI to isolated parts of the support stack rarely delivers meaningful results. Without the right workflows, guardrails, and consistency across channels, AI helps in pockets but doesn't fundamentally change how support teams operate.

That's where Hiver's model makes sense to 10,000+ teams globally. We bring AI into every customer service touchpoint, and not just the front end. Here's what that looks like in practice:

- **AI Copilot** helps agents write and refine replies right from the inbox.
- **AI Agents** automate repeat tasks and simple resolutions without extra setup.
- **AI Insights** surface patterns, volume spikes, and opportunities for improvement without running separate reports.
- **AI Quality Assurance (QA)** automatically reviews every customer conversation to identify patterns in tone, accuracy, and resolution quality, and flags specific areas where teams can improve.

The screenshot displays the Hiver support interface. On the left is a navigation sidebar with options like Compose, Inbox, Sent, Drafts, and various support channels. The main area shows an email thread from Sarah Thompson asking for a discount. An AI-generated draft response is visible, offering a 10% discount and providing a claim ID. On the right, a metadata panel shows the ticket is assigned to John, tagged as Urgent and Billing, and includes a claim ID field.