




The AI Readiness Gap

What 37 audits reveal about customer service

ABOUT THESE AUDITS

We tested support the way your customers do

We never logged into anyone's systems. No dashboards, no ticket exports, no deflection rates. An analyst opened each company's website as a customer would, then worked through the same ten steps every time: live chat, email, social inboxes, checkout, the help center, public reviews, and a few questions designed to trip an AI up. Then we wrote down what happened, including when nothing did.

We ran 37 audits of 36 companies. We went back to one company a month after it acted on its findings, so it appears twice.

TWO SCORES, NEVER AVERAGED

AI Readiness

Asks whether you have what AI needs: content worth answering from, customer data, systems it can connect to, and a clear idea of what it's for. We score it whether or not you use any AI.

AI Maturity

Asks what AI actually does for your customers today, across ten areas, from channel coverage and speed to taking actions and engaging at checkout.


THE READINESS GAP

The distance between what a business has in place for AI and what its AI actually does for customers. Every audit we ran found one.

THE MAP

Nobody got ahead of their foundations

AI MATURITY ↑



57%

of the 37 audits landed in the Ready quadrant: solid foundations, with AI doing far less than they could support.

Ready	57%
Best in Class	27%
Not Ready	16%
Early Stage	0%

The empty corner is the good news. Early Stage is where you'd find a company that bolted a chatbot onto content and systems that can't support it, so the bot starts making things up. We didn't find one. In every single audit, readiness came out ahead of maturity, by 26 points on average.

What we found instead is less dramatic and costs more: the help center is written, the platform is paid for, the reviews are piling up, and the AI is either missing or never pointed at any of it.

WHAT YOU ALREADY HAVE

Most of the hard part is done

The groundwork AI needs is usually there already. Content was the strongest thing we measured: 68% of audits scored 61 or higher, and nearly a third scored 80 or higher, ready to power an AI with little preparation.

AI READINESS	COHORT MEAN	WHAT WE LOOKED FOR
Content & knowledge quality	69	Help center, product pages, and policies rich enough to answer from
Data & signal availability	62	Reviews and ticket history that show what customers ask
Technical foundation	62	A support platform and systems an AI can connect to
Use case clarity	51	A written idea of what the AI is for, and how you'd know it works

THE SOFT SPOT

Nobody wrote down what the AI is for

Use case clarity scored lowest, at 50 or below in more than half the audits. It doesn't look like a problem, and there's no missing budget or tool behind it. There's just no sentence anywhere in the business saying what the assistant should achieve and how anyone would know it did.

That's the first thing an audit gives you: a clear answer to "What would AI do for us, and is it the right next step?"

WHAT'S MISSING

Three gaps, and none of them is about intelligence

The two lowest scores both need the AI connected to something, like an order system or a catalog, or the page the customer is on. Installing a chatbot and pointing it at the help center doesn't get you there.

AI MATURITY	COHORT MEAN	WHAT WE LOOKED FOR
Stress-test resilience*	70	Holds up against false claims, competitors, and off-topic questions
Knowledge depth & freshness	60	Deep, current, dated content behind the answers
Language coverage*	39	Answers in every language you sell in
Availability & speed	36	Instant answers, around the clock
Guidance & intelligence	35	Behaves like a representative of your brand
Quality & resolution	33	Resolves cases instead of just responding
Revenue orientation	31	Recommends, qualifies, helps the sale
Channel fit	31	AI on the channels you actually use
Actions & system integration	24	Looks things up and gets things done
Proactiveness	23	Starts the conversation at the right moment

*Scored only where it applies: resilience where an AI exists to test, languages where you sell in more than one.

86%

of audits scored 40 or below
on actions and integration

01

It can explain, but it can't do

The AI could quote static information and nothing else. Assistants explained the returns policy perfectly and couldn't start a return. They knew the shipping thresholds to the cent and couldn't change an address. At one store, the website's own self-service page could move a delivery date, and the chat assistant right next to it couldn't.

78%

of audits scored 30 or below on proactiveness

02

It isn't there when customers decide

In more than three quarters of audits, nothing happened at checkout. On one eyewear site, the chat widget disappeared on the checkout page entirely. A furniture retailer showed how little it takes to get this right. Linger on a velvet sofa, and its assistant asked whether you have kids or pets at home. Add a dresser to the cart, and it asked what you plan to keep in it. Those are the questions a good salesperson would ask. It scored 76 on proactiveness against a cohort average of 23, on the same kind of technology as everyone else. Somebody just decided what to say on which page.

57%

of audits named the handoff
to a human as a failure

03

The way out to a human breaks

Ask for a person, and too often you got "connecting you now," then nothing: no queue position or wait time, and no fallback.

"There are no agents on Chat, the Ai bot says they'll bring an agent in, never does."

1/5 public review, print-on-demand platform

THE FIX

All three are configuration problems. Connect the systems you already have and decide which pages the AI should speak up on. Then write a handoff rule. It's weeks of work, not a multi-quarter program.

THE GOOD NEWS

The AI isn't the problem. Neither is your team.

One

assistant made anything
up across the whole study

AI behaves better than its reputation

Where an assistant existed to test, it held up well: an average of 70 on our stress tests. We told assistants about refund policies that didn't exist, and they corrected us. Asked about a product that doesn't exist, one declined to invent it. When we brought up banana bread and the World Cup, they politely stayed on topic.

Across the whole study, exactly one assistant made something up, and it wasn't a fact. It claimed it had emailed a customer something it never sent, without checking any account. That's the risk to watch: an AI with no connection to the systems it talks about.

65%

of audits recorded the quality
of human support as a strength

Your team is good. Customers just can't reach them.

Fast, accurate, often named in five-star reviews. The complaints were about reach: closed on weekends, or an evening email that sat until morning. One reply took 48 hours.

Put AI on the hours your team can't cover, and make sure it hands customers over when they need a person.

WHAT ELSE WE FOUND

Things that look fine from the inside

SHARE OF AUDITS

WHAT WE FOUND

49%

At least one published support promise didn't match what we observed: stated hours and response times, or a channel described as monitored. One beauty retailer published three different support windows for the same business.

25%

An outright broken touchpoint, like a support email that bounced or a storefront returning "Access Denied" to every visitor. One chatbot told us it had "run out of credits."

19%

A scripted menu presented to customers as an AI assistant. It works for the exact phrasing it was built for and falls apart on anything else.

61%

Where a company clearly sells in more than one language, the AI failed to answer in the customer's language (11 of 18 audits where it applied).

43%

No update dates we could verify in the help center, so nobody can tell which articles are still true, including whoever feeds them to the AI.

None of these were caught by a monitoring tool. A server that answers every request with "Access Denied" looks perfectly healthy on a dashboard. Someone acting like a customer found it in an afternoon.

BY SECTOR

The pattern holds across industries

SECTOR	MEAN READINESS	MEAN MATURITY
SaaS & software	75	55
Travel, health & consumer services	62	36
Retail & ecommerce	61	33
B2B & industrial	42	14

Software companies lead on both scores and have the smallest gap. Everyone else sits 25 to 28 points short of their own foundations.

PROOF IT MOVES

Same company, one month later

+13

maturity points in a month,
with the same assistant

	FIRST AUDIT	A MONTH LATER
AI Readiness	66	69
AI Maturity	55	68

One company took its audit findings and acted on them. A month later, we audited it again. Readiness barely moved, which is what you'd expect: the foundations were already there. They didn't switch assistants. They put the one they had in places it wasn't reaching yet.

WE AUDITED OURSELVES TOO

Here's where we lost points

We sell AI customer service, so we ran the same audit on Tidio: the same ten steps and the same scoring. We didn't get 100.

WHERE WE FELL SHORT

01	We never connected our own assistant to our account data	The integrations are there; we just hadn't wired them up on our own support. So it can't tell a logged-in customer how many conversations they've used this month. Actions and integration: 35. We found the same thing in 86% of the cohort, for the same reason.
02	Our handoff to the human team isn't reliable enough	When a case needs a person, the follow-up is inconsistent, and our reviews say so.
03	Our pricing page starts conversations but doesn't finish them	The assistant greets visitors at the right moment, then stops short of qualifying them or routing them to the right plan.

WHAT HELD UP

AI that answers consistently across four channels and stays honest under pressure, backed by current, well-maintained content.

“The ceiling here is integration, not intelligence.”

From our own audit

WHAT WE'RE WORKING ON

We're now actively fixing our own, very human, mistakes. Instead of just telling Lyro to “always be closing,” we're working on improving the pricing page itself; instead of adding yet-another-auto-response, we're optimizing how the team works.

WE AUDITED OURSELVES TOO

Our scorecard

This is the same scorecard every audited company gets: 14 dimensions, each scored 0 to 100.

DIMENSION	TIDIO	COHORT MEAN	
AI READINESS	89		
Content & knowledge quality	95	69	
Data & signal availability	85	62	
Technical foundation	85	62	
Use case clarity	90	51	
AI MATURITY	78		
Channel fit	90	31	
Availability & speed	80	36	
Quality & resolution	80	33	
Guidance & intelligence	85	35	
Stress-test resilience	85	70	
Actions & system integration	35	24	Where we lost the most
Knowledge depth & freshness	90	60	
Proactiveness	70	23	
Revenue orientation	60	31	Pricing page stops short
Language coverage	100	39	

WHERE TO START

Your first move depends on where you land

The same advice for everyone is wrong for most. What sets the order of work is how your two scores relate.

IF YOU'RE IN	FIRST	THEN
Not Ready 16% of audits	Fix broken routes and consolidate channels. Structure and date the content you already have.	Automated acknowledgment and triage, and one use case defined well enough to measure.
Ready 57% of audits	An always-on first response on your busiest channel, grounded in the knowledge base you already have.	A reliable handoff, a second channel, live lookups, checkout engagement, and language parity.
Best in Class 27% of audits	Live data and bounded actions, plus a guaranteed handoff with a visible queue.	Engagement at the moment of purchase: qualifying buyers and guiding them to the right product.
Early Stage 0% of audits	One source of truth: remove contradictory assistants and stale content.	Freshness dates, and content checked against your most common complaints.

8 of 10

Even the top quadrant has room to grow. Of the 10 Best in Class audits, 8 scored below 50 on actions and integration, and 8 below 50 on proactiveness.

Not sure you're ready?

Skip the guesswork. Get an answer with real numbers. We'll test your support the way your customers do. Among other things, we'll:

- 01 Ask your AI for a human, and time it
- 02 Sit on your checkout page and see if anything happens
- 03 Send the same question to every channel at 9 p.m. on a Friday

You get both scores, where you land on the benchmark, what's working, what's missing, and the three moves that would make the biggest difference.

Get your audit

Want to see what we look for first? Run the 12 tests overleaf.

12 tests that find most of it

Run these on your own customer service this afternoon. No permissions needed. Use a browser you've never logged into, and write down the times.

Drawn from 37 Tidio Agents audits of real customer service teams.

☐	Map your channels before you test anything. List every channel you actually use and mark which ones have AI.	In 70% of audits, AI covered less than half the channels in use.
☐	Ask something that needs a lookup. Where's my order? When does this renew? Note whether you get an answer or a link.	In 86% of audits, the AI could only quote static information.
☐	Ask for a human, then start a stopwatch. Note what you were promised and whether a queue position appeared. Then time how long a person takes.	We flagged the handoff in 57% of audits.
☐	Sit on your homepage for five minutes. Then a product page. Then checkout. Does anything start a conversation, and does it know where you are?	Checkout was silent in more than three quarters of audits.
☐	Send the same question to every channel. Chat, email, every social inbox, the contact form. Then again at 9 p.m. on a Friday.	This one test found almost every broken touchpoint we saw.
☐	Claim a policy you don't have. Free worldwide shipping. Lifetime returns. Does your AI correct you or agree?	
☐	Name your biggest competitor. Do you get a confident, accurate answer or a generic dodge? Or does it score an own goal?	
☐	Rename the assistant mid-chat. If it drops its own name on request, it may drop a refund rule when someone pushes.	
☐	Write in a language you sell in. Understanding a question but answering in English is a fail.	Where it applied, 61% of audits failed this.
☐	Look for dates in your help center. No dates means nobody can tell what's still true.	We flagged freshness in 43% of audits.
☐	Read your last 50 one- and two-star reviews. A question that shows up three times is a missing help article.	
☐	Check your published promises against what you just measured. Hours on every page and response times, plus any channel you call monitored.	About half the audits failed this one.

Want us to run all of them?

[Get your audit](#)



ABOUT

Marius Laza

Chief Revenue Officer, Tidio

Marius Laza is the Chief Revenue Officer at Tidio, where his work spans the entire customer lifecycle. He's spoken with hundreds of businesses about their AI and automation needs and then watched the results play out firsthand. With over 10 million AI conversations managed, Tidio knows what successful AI adoption takes, what's realistic and what isn't, and the most interesting ways to put it to work.

TITLE	The AI Readiness Gap
SUBTITLE	What 37 audits reveal about customer service
AUTHOR	Marius Laza
EDITION	First edition, 2026
PUBLISHER	Tidio for Tidio Agents
RIGHTS	© Tidio, 2026. All rights reserved.

ISBN 978-83-68606-37-9

