

Customer service: where AI meets your culture first

How to identify dead company culture in the age of AI and agentic work

DEMO EBOOK

GetSorted · getsorted.valuebull.com.au/ident/customer-service/

Published 19 Oct 2026

Short answer. Customer service is where the gap between what a business promises and what it actually does becomes visible first, because customers ask the questions nobody wrote down. The signs: a written policy that differs from the real one, exceptions that depend on who answers, escalation that's a person rather than a rule, and values that promise more than the script allows. The test is simple: give your team and an AI agent the same case and the same written information. Where they disagree, you've found an unwritten rule.

Forty days

Bec runs customer service for an outdoor furniture business in Perth. Seventy people, a showroom, a warehouse in Canning Vale and a busy website. In September they switched on an AI agent for web chat after hours.

On the second Saturday a customer wrote in. Her new dining table had a scratch across the top. It had been delivered 40 days earlier, but she'd only just taken the plastic off.

The agent checked the returns policy. Thirty days. It politely declined, offered her a care guide, and closed the chat.

By Monday there was a one-star review. Bec read the transcript at seven in the morning with a coffee she didn't enjoy.

Here's the thing. Her team would have replaced that tabletop without blinking. Scratches on a wrapped table almost always come from freight, and the business claims the cost back from the carrier. The team has handled it that way for years, out to about 90 days. It just isn't written anywhere. It's in Bec's head, and in the heads of the three people she trained.

The agent did exactly what the policy said. The policy just wasn't what the business does.

The warning signs

Customer service is usually where culture on paper meets a real person first. In most businesses we see, these signs show up.

1. Everyone knows the real policy except the policy. The document says 30 days. The team says "depends". New starters learn the difference by getting it wrong.

2. The answer depends on who picks up. One person waives the restocking fee for regulars. Another never does. Both are acting in good faith.

3. Escalation is a name, not a rule. "Check with Bec" is the process. When Bec's on leave, the process is on leave too.

4. The values promise more than the script allows. The website says "We make it right". The script says "Unfortunately, that falls outside our policy."

5. Your best people answer fastest because they know the workarounds. That's a skill worth having. It's also a sign the workarounds are where the real policy lives.

6. Customers ask for someone by name. Lovely for that person. It also means the customer has worked out the answer depends on who they get.

The gap

What the company says: "We make it right."

What's written down: A 30 day returns policy from 2022 and a three page refund procedure.

What really happens: Freight damage replaced out to 90 days. Trade customers get credit, not refunds. Restocking fee waived for repeat customers, if you get the right person.

Before you go and rewrite the policy, look at who's right. Usually it's the team. They've learned what keeps customers and what doesn't, one complaint at a time. The friction isn't that people bend the rules. It's that the bent version, the one that works, never made it back onto paper.

The AI test

This is the department where AI tools tend to turn up first, so it's where the gap gets found first.

A new person on the phones would have asked someone about the table. An AI agent can't ask. It has your written policy, your FAQ and whatever else you've connected, and it acts on those. It has no Bec.

Bec's team could have caught this before switching on. Here's the version we'd suggest. Take 50 recent tickets. Give the agent only your written policies and ask what it would do. Then compare with what your team actually did. When one business ran this, the agent and the team agreed on 31 and disagreed on 19. In 14 of those 19, the team was right and the written policy was wrong. **DEMO**

That's not a failed pilot. That's 14 unwritten rules found in an afternoon, before a customer had to find them for you.

What an agent needs is the same thing a new starter needs, just written down:

- which exceptions are allowed, and up to what amount
- who can approve what
- when to stop and hand the customer to a person

If you give it those three things, it will usually do better than a new starter on day one. If you don't, it'll do exactly what Bec's agent did.

What good looks like

Picture Bec's Saturday again, with one extra page in place. The agent reads that freight damage on a wrapped item gets replaced up to 90 days after delivery. It reads that a replacement under \$800 needs no sign-off **DEMO**. And it reads that anything over that, or any customer who is clearly upset, goes to a person first thing Monday with a note attached.

The customer gets a new tabletop, or at worst a message saying Bec will call her on Monday morning. No one-star review. Same team, same products, same agent. The only difference is that the rule Bec has been carrying for years is now written where everyone, and everything, can find it.

The connection

Customer service mostly catches problems that started somewhere else. That makes it a good place to trace back from.

- **Sales** sets the promise. If a rep said "anything goes wrong, we'll sort it", customer service inherits that line with no policy behind it. That's week 5.
- **Marketing** writes "hassle-free returns" on the website. The agent reads the policy, not the website. That's week 6.
- **Operations** owns the freight damage process and the carrier claims. If that process isn't written down, customer service can't promise anything with confidence. That's week 7.
- **Finance** decides what refunds and replacements cost and who signs them off. If the dollar limits only exist in someone's head, so does the authority.
- **The leadership team** wrote "We make it right" without saying what "right" costs. Back to last week.

The ident check

Yes or no. Run it with your customer service lead.

1. Could everyone on your service team name the three things the business most wants from a customer conversation, without notes?
2. Could each of them say what they're allowed to approve on their own, and up to what amount?
3. Could each of them name one old policy or exception that has stopped?
4. Is every exception your team regularly makes written down somewhere an AI tool could find it?
5. Would two people on your team give the same answer to the same complaint?
6. Is there a written rule for when a customer gets handed to a person?
7. Does what your website promises match what your written policy says?

Questions 1 to 3 are the recall test, adapted for the front line. If question 2 is a no, that's usually the most useful place to start.

Cheap first

1. **Pull ten recent tickets.** Costs nothing, takes an afternoon. For each one, ask: what did we actually do, and does the policy say that? Mark the ones where it doesn't.
2. **Write down what your most experienced person knows.** One hour with your Bec and a notepad. Just a list of the exceptions they make and why. Don't judge them yet.
3. **Run the same-decision test.** Give five real cases to a team member and to an AI tool that only has your written policy. Compare the answers side by side.
4. **Decide out loud.** For each exception: it becomes policy, or it stops. Write the keepers in. Tell the team which ones are gone.
5. **Then the harder fix.** Agree who can approve what, up to which dollar amount, and write the handoff rule. That one page does more for your customers, and for any AI agent you add later, than a new values poster ever will.

The Monday close

This week, pull ten tickets and check them against your written policy. Count how many match.

Next week we start Part 2: how culture travels. First stop is the team leaders, the translation layer. They turn leadership's intentions into what happens on the floor, and they're usually doing it on their own, in their own words.

Questions people ask

Why do AI customer service agents get answers wrong?

Usually because they follow the written policy exactly, and the written policy isn't what the business actually does. The exceptions experienced staff make every day were never written down, so the agent can't know them.

What does an AI agent need before it answers customers?

The same thing a new starter needs, written down: which exceptions are allowed and up to what amount, who can approve what, and when to hand the customer to a person.

How do you test an AI agent before switching it on?

Take 50 recent tickets, give the agent only your written policies, and compare what it would do with what your team actually did. Every disagreement is either a policy to fix or an unwritten rule to write down.

The full series, episodes and downloads: getsorted.valuebull.com.au/ident/. Demo content: items marked Demo are illustrative.