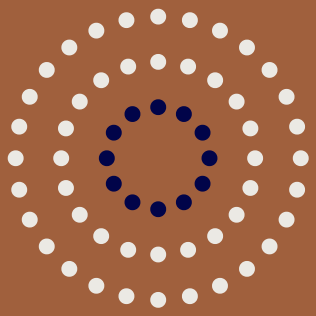


CONVERSATIONAL CAPITAL

Do customers come back after your assistant says it's done?

How the Conversational Capital Rating (CCR) counts the customers who take a finished conversation with an assistant to a person, and what that costs.

Conversation Design Institute

TOPICS

Conversational Capital Rating · Volume Substitution Rate · Trust · Repeat contact · ROI · AI assistants

An assistant tells a customer their policy change has been made. The conversation ends, and the assistant's dashboard records it as handled. Two days later the same customer rings the contact centre to check the change went through. On the assistant's figures, that was a success. On the business's costs, it was two contacts instead of one.

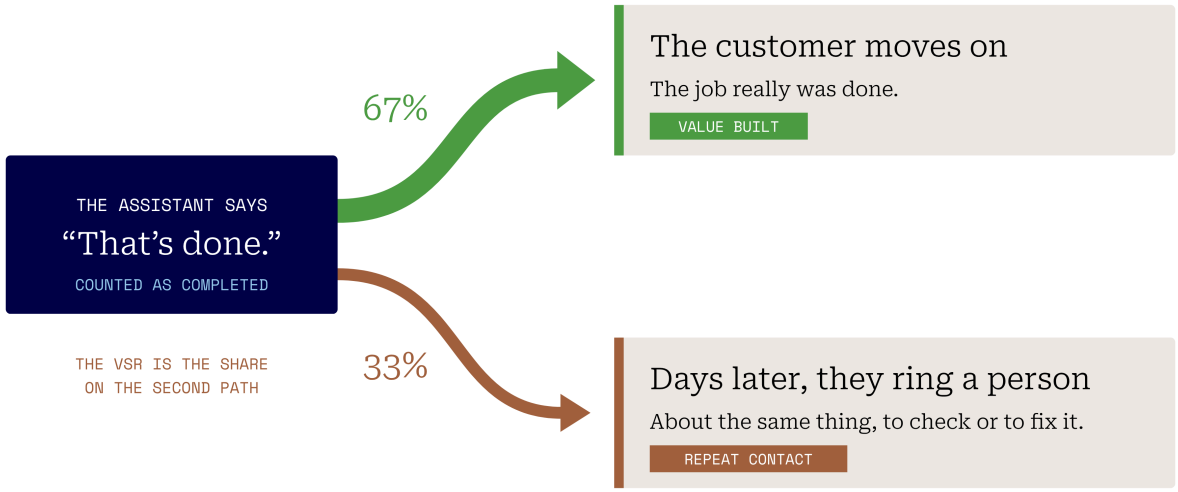
This is one of the most expensive failures an assistant can have, one that grows as assistants do more for customers rather than just answering them, and often the hardest to see. It does not happen in the conversation, so no review of conversations can find it. It happens afterwards, in another channel. The Volume Substitution Rate (VSR) is how the CCR counts it. Repeat calls are the first of three costs an assistant can hide, alongside lost customers and lost trust, and the one an organisation can count soonest.

01 Summary

WHAT IT IS	The share of conversations the assistant counted as completed whose customer then contacted a person about the same thing.
WHY IT MATTERS	Every one of those is a repeat contact: a saving the assistant's figures reported that never happened. It is the clearest sign that customers have stopped relying on the assistant, and where most of its financial cost lies.
HOW IT WORKS	It is built from counts the contact centre already keeps, by topic and month. No customer records leave the organisation.
HOW WE KNOW IT WORKS	It is counted, not judged. Where an organisation cannot link its records, the figure is estimated, and the report says so and shows how wide the estimate is.

02 How it works

The VSR asks one question of every conversation the assistant counted as completed: did the customer then contact a person about the same thing within a set number of days?



Every conversation the assistant counts as completed takes one of two paths.

It is read topic by topic. An assistant may answer opening-hours questions perfectly while a third of its policy changes are checked by phone. One overall figure would hide that; by topic, it shows exactly where the assistant is not trusted.

The assistant's own figures cannot see this. Containment counts conversations that ended without a person, so a conversation that ended and was followed by a phone call still counts as contained.

When trust failure climbs past a point set for the industry, the rating is held down however capable the assistant is, because customers going round an assistant no longer benefit from what it can do.

Take an insurer whose assistant completes 50,000 policy changes in six months. Its dashboard shows 50,000 contacts saved. But the contact centre's counts show that a third of those customers rang about the same policy within a week, so the VSR for policy changes is 33%: about 16,700 repeat contacts, each a call the business believed it had saved.

03 Where it sits among the six measures

The CCR weighs two sides of an assistant. Three measures show the value it creates and three show the trust it spends. The VSR is one of the three on the trust side.

- Craft (value): how well the assistant handles what the customer came for.
- Stake Coverage (value): whether it handles the conversations that matter, not just the easy ones.
- Velocity (value): whether conversations move at the right pace for the need.
- VSR (trust): whether customers then go round the assistant to a person.
- Friction (trust): where conversations visibly break down.
- Sentiment Decay (trust): whether customers leave feeling worse than they arrived.

An assistant cannot make up for poor handling by being fast. The trust side is combined with Friction and Sentiment Decay, because a customer who goes round the assistant has shown, not just felt, that they do not trust it. As trust failure grows, each failure costs more.

04 What it is worth

Because the VSR counts contacts, it converts directly into financial terms using one figure every contact centre knows: what a contact with a person costs. The report gives three amounts:

- Value built. Conversations the assistant completed that customers did not repeat to a person. Each is a contact saved, if it would otherwise have reached a person.
- Repeat contacts. Follow-up calls on problems the assistant said it fixed: conversations it completed that customers took to a person anyway. Each is a saving the containment figures reported that never happened: the business pays for the conversation, and again for the call that follows.
- Customers at risk. Where customers can leave, some of those who had to come back will go at renewal. Using the organisation's own churn rate and yearly revenue per customer, the report gives this as a floor, not a forecast.

For the insurer, at about €7 a call, in line with industry benchmarks, the 33,300 policy changes that stuck saved about €230,000 in six months, and the 16,700 repeat contacts cost about €120,000. That second figure is what fixing policy changes saves. The customers at risk are usually the larger figure: the revenue that stays when fewer of them have to come back.

05 What it means for clients

VSR also shows what a fix returned. After a change to the assistant, the repeat contacts on that topic should fall within a month or two, before complaints or churn

move. With continuous measurement, that fall is the best measure of the fix's saving in your own costs, and it can be set directly against what the fix cost. Later, as fewer customers leave, the fall in customers at risk is the revenue the fix protected.

If the counts show customers going round the assistant on a topic, the order of work is clear: make the assistant finish that job reliably, and confirm it only once it is done, before adding anything new. New features land on customers who have already stopped trusting it.

Conversation Design Institute

About Conversational Capital and the CCR

Conversational Capital is the trust, goodwill and willingness to engage that an organisation builds up when people feel respected, understood and helped in their conversations with it, and loses when they do not.

The Conversational Capital Rating (CCR) is how we measure it. It weighs two sides of every conversation with an assistant: the value it creates, by handling what the customer came for well and at the right pace, and the trust it spends, where the conversation breaks down or the customer leaves worse off than they arrived. The outcome is a single rating out of 100, with the evidence behind it, and a priority list of how to raise it, and so build the trust your brand depends on. It is built for assistants that use large language models and increasingly do things for the customer, not just answer their questions.

The CCR turns the quality of your AI conversations into a measure you can track and act on. It shows which kinds of conversation are building trust and which are spending it, and where you should invest first. It also shows what the assistant really saves, and what it is really costing you in repeat calls, lost customers and trust, months before those costs reach your accounts, so that further investment can be judged on its return.

