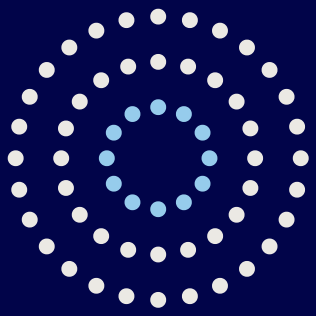


CONVERSATIONAL CAPITAL

What is your AI assistant really worth?

How the Conversational Capital Rating (CCR) measures what an AI assistant is building with customers and what it is spending, and shows what they are worth in financial terms.

Conversation Design Institute

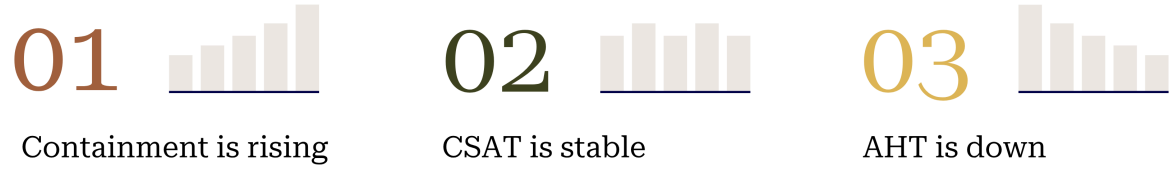
TOPICS

Conversational Capital Rating · Trust · Value · ROI · Evaluation · AI assistants

Most organisations can say what their AI assistant cost. Few can say what it returns. The usual figure, containment, counts the conversations the assistant ended without a person. It does not count the customers who then phoned, queued or complained about the same thing, or the ones who quietly stopped using the assistant at all.

So if you ask where the return is, nobody can answer with confidence. The assistant may be doing well. It may be costing more than it saves. Containment cannot tell the two apart.

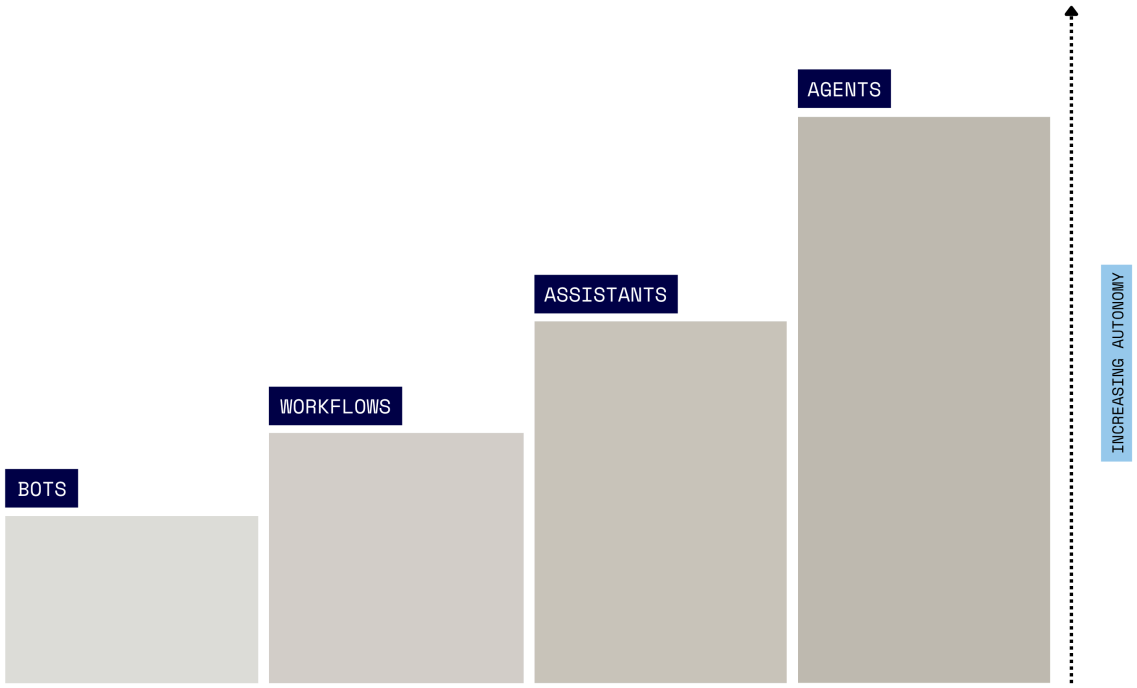
Put plainly, the question is what your assistant is really costing you: in repeat calls, in lost customers and in trust. Lost trust shows up in your conversations months before it reaches your accounts, and that is where the CCR looks for it.



The usual figures can all look healthy while customers go round the assistant.

The **Conversational Capital Rating (CCR)** from Conversation Design Institute is built to answer that question. It analyses how the assistant handles customers and returns a rating out of 100, the evidence behind it, what the assistant is worth in financial terms, and the fixes that pay back first.

Conversation design and agentic experience design are how an assistant is built well; the CCR is how you know it is working. It gives an independent, repeatable measure of the trust your assistant creates or spends. It is built for assistants that use large language models and increasingly do things for the customer, not just answer their questions, and it can rate the assistants you run today.



From bots to agents: as assistants do more for customers, more rests on the trust they keep.

01 Summary

WHAT IT IS	A rating out of 100 for an AI assistant, built from its real conversations. It weighs the value the assistant creates against the trust it spends.
WHY IT MATTERS	It turns the quality of an assistant into figures you can use: what the assistant saves, what its failures cost, and which fixes would pay back first.
HOW IT WORKS	Six measures, each read from evidence an organisation already holds, combine into one rating. When customers have stopped relying on the assistant, the rating is held down however capable the assistant is.
HOW WE KNOW IT WORKS	Every reading is made against written rules, every failure it reports points to the exact words in the conversation, and each measure states how consistent it is.

02 How it works

Every conversation with an assistant does two things at once. It creates **value**, by handling what the customer came for, well and at the right pace. And it spends **trust**, whenever it breaks down, leaves the customer worse off, or sends them round the assistant to a person.

The CCR measures both sides. Three measures make up the value side and three make up the trust side. The value side shows what the assistant is capable of. The trust side shows how much of that capability customers can actually rely on.

A capable assistant with broken trust does not score well, and neither does a trusted assistant that handles little. Small amounts of lost trust cost a little; as failures mount, each one costs more. The result is one rating out of 100, shown with its six parts so the reason for the score is always visible.

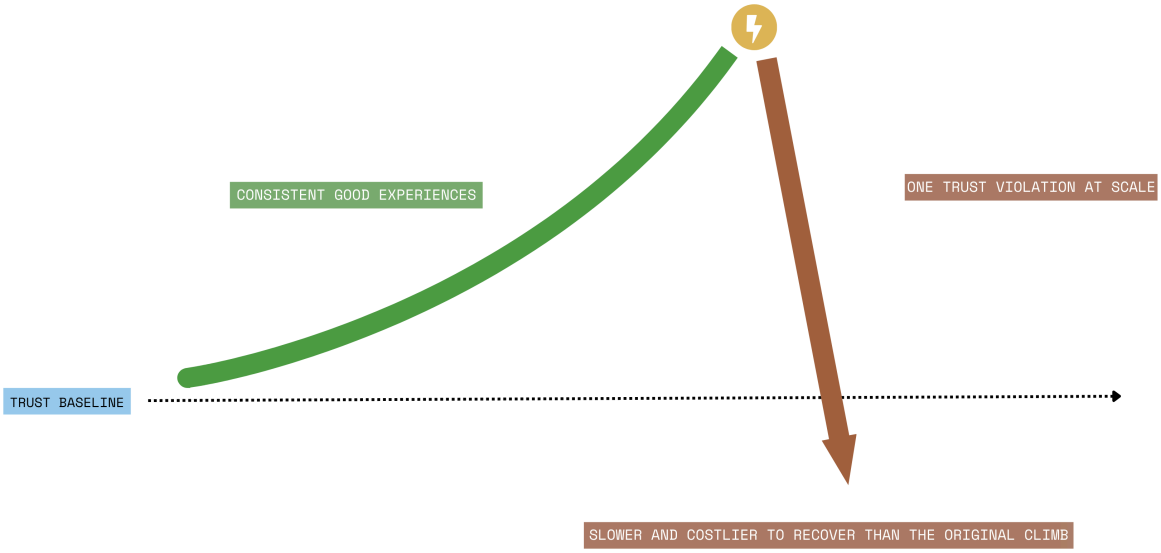
03 The six measures

MEASURE	SIDE	THE QUESTION IT ANSWERS
Craft	Value	Did the assistant handle the conversation well?
Stake Coverage	Value	Is it handling the conversations that matter, or only the easy ones?
Velocity	Value	Did the conversation move at the right pace for what the customer needed?
Volume Substitution Rate (VSR)	Trust	Do customers the assistant counted as helped then contact a person about the same thing?
Friction	Trust	Where do conversations visibly break down?
Sentiment Decay	Trust	Does the conversation leave the customer feeling worse than when they started?

04 When trust breaks

Customers do not lose trust gradually for ever. They put up with a failing assistant for a while, then stop trying and go round it. From then on, whatever the assistant can do no longer reaches them. Once trust failure passes a set point, the CCR is held below a limit that no amount of new capability can lift until customers come back. To fix this, the steps are: repair trust first, then add features, because new features only

pay off with customers who still rely on the assistant.



Trust builds slowly and can be lost at once; recovering it is slower and costlier than the original climb.

05 What it is worth

The CCR puts the assistant's performance in financial terms using one figure the organisation already knows: what a contact with a person costs.

- Value built. Work the assistant finished that customers did not have to repeat to a person.
- Repeat contacts. Follow-up calls on problems the assistant said it fixed: work it counted as done that customers then took to a person anyway. This is what the VSR counts. These repeat contacts are savings 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. It is usually the largest of the three figures.

Take an insurer, whose assistant completes 50,000 policy changes in six months, and a third of those customers ring afterwards to check the change went through. At about €7 a call, in line with industry benchmarks, about €230,000 of contacts were saved, if each would otherwise have reached a person, and about €120,000 went on repeat contacts. The second figure is what fixing the failure saves. The third is usually larger: the revenue that stays when fewer customers have to come back.

Together, the three figures correct the number most dashboards report. Containment would credit this assistant with all 50,000 policy changes. The CCR takes away the customers who called anyway and puts a floor under the customers at risk; what is left is what the assistant really returns.

This changes the conversation about ROI. You can see what the assistant returns now, set the next investment against it, and, with continuous measurement, see what each fix returned: the fall in repeat contacts after a fix is the best measure of that fix's saving, and the fall in customers at risk is the revenue it protected.

06 What is the CCR worth to you?

Customers going round an assistant are expensive because a contact with a person costs many times more than one the assistant completes: industry figures put the median cost of an assisted contact at about seven times that of self-service.

Take an assistant that completes a million conversations a year. If one in a hundred of those customers then contacts a person about the same thing, that is 10,000 repeat contacts, or about €70,000. If one in five does (20%), that is about 200,000 repeat contacts, around €1.4 million a year.

Fixes that bring that share down from 20% to 15% would recover about €350,000 a year in contact costs alone, before counting a single customer kept. The CCR shows where that gap is, topic by topic, which fixes would close the most of it, and, through continuous measurement, how much each fix actually returned. That return can be set directly against the cost of the rating and of the fixes.

07 What it means for clients

A rating needs three things your organisation already has: a few hundred of your assistant's conversations per channel; monthly counts by topic of conversations the assistant handled, the ones it completed and, where you can count it, how many of those customers came back to a person; and the cost of one contact with a person.

The CCR can be used once, as a Baseline Assessment that rates the assistant at a point in time. With continuous measurement, you can track how a change in the assistant, or in how customers respond to it, shows up so you can act on it. CDI can then work with you to make those fixes, and the next rating shows what they returned.

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.

