

AI Customer Service Chatbot for a Credit Union, Answering
Member Questions Across Three Separately Grounded Assistants

The first assistant worked, so the credit union built a second, and then a third

Three assistants added one at a time, each grounded separately, all three still answering members on the last day recorded, and 97.8% of the questions their material covers answered well.

48,372

Questions asked, raw count

97.8%

Answered well, of what the
credit union's material covers

198

Revisions the credit union has
made to how the three answer

3

Assistants in production, added
one at a time

THE CHALLENGE

A credit union answers to its members, and the members ask on Saturday

A member working out on a Saturday morning whether a payment posted before the weekend, and what happens if it did not. A contact center representative who has answered that same question nine times before lunch and will answer it again after. A digital banking manager who can see in the analytics that the questions arriving by phone and the questions arriving on the website are not the same questions, and who has one website to answer both with.

Member ownership raises the stake on a wrong answer. A member who is told the wrong thing about a fee, a hold or an eligibility rule is an owner who was told the wrong thing, and the recourse is a phone call that costs more attention than the original question ever did. Scale makes that worse rather than better. At tens of thousands of questions, a small rate of confident error is a large number of members.

WHAT GOT BUILT

One assistant first.

A single grounded deployment on the credit union's published material, which is still the highest-volume of the three.

A third, some months later.

Narrower again, the smallest by volume and the highest by share of questions classified.

A second, within weeks.

Added alongside the first rather than replacing it, with its own grounding and its own actions.

Answer quality recorded.

This deployment classifies whether a question was answered, which is what makes the rate on page three publishable at all.

THE REQUIREMENTS

Five things a member-owned institution would not compromise on

A credit union is answerable to the people asking the questions. That changes what an acceptable answer looks like, and it changes what an acceptable unknown looks like.

- 1 Bounded by published material.
- 2 Room to add, without starting over.
- 3 Narrow scopes stay narrow.
- 4 Whether it answered had to be recorded.
- 5 No member data in scope.

SWORN TO SECRECY

A member-owned credit union in the southeastern United States serving several counties, with roughly three hundred staff. Shared anonymously. Every figure comes directly from the credit union's own deployment and is scoped to its three production assistants.

AT A GLANCE

INDUSTRY

Credit union. A member-owned institution in the southeastern United States serving several counties.

DEPLOYMENT

Three production assistants, added one at a time over fifteen months, all three still answering.

ACCESS

Public, on the credit union's website, serving members and prospective members alike.

HOW IT WAS BUILT

One at a time, each addition judged on what the last one was doing

Ship one and watch it.

The first assistant ran on its own long enough to produce a record rather than an impression.

Add the third where the first was thin.

Narrower scope, narrower corpus, and the highest share of questions classified of the three.

Add the second alongside.

Within weeks of the first, with its own grounding, and without taking the first out of service for a day.

Keep the record running.

Answered and unanswered counts accumulate per assistant, so the deployment can be read assistant by assistant rather than in aggregate only.

ADDING AN ASSISTANT, THE CAPABILITY THAT DECIDED IT

The second assistant did not cost the first one a day of service

A new assistant here is an addition rather than a migration. The credit union put one in production, let it run long enough to produce a record, then added a second with its own grounding while the first kept answering, and later a third narrower again. Platforms that make you widen one assistant to cover a new scope force a choice between coverage and precision. Separate grounding means a narrow question can be answered from narrow material, and the numbers below show the three behaving differently because they are asked different things.

THE RECORD

Three assistants, read one at a time

WHAT THE NUMBERS SAY

Read the rows separately, because they are separate assistants answering separate questions. The second carries a quarter of the volume of the first and has been revised more often, which is the signal you would expect from a narrower front door taking harder questions. The third is the smallest and the most worked on of all. Across the deployment, questions arrive on 446 of the 462 days it has existed, which is the shape of a service rather than a pilot.

WHY CUSTOMGPT.AI CAME OUT ON TOP

The deployment could grow. Two assistants were added without interrupting the one already in service.

Scopes stayed separate. Each assistant carries its own grounding and its own tuning, 198 revisions between the three of them.

The record is per assistant. Answered and unanswered counts accrue separately, and two of the three run advanced reasoning over their own material.

WHAT THIS MEANS FOR YOU

Ask what happens when one assistant stops being enough

Test the second one before you buy the first. The question worth asking is what adding a scope costs you later, because the answer decides whether you end up widening one assistant until it is vague.

Insist on a per-assistant record. Aggregate numbers hide the assistant that is struggling, and a deployment you cannot read row by row is one you cannot fix row by row.

The member checking on a Saturday whether a payment posted gets the credit union's own published answer at the hour she asks. The contact center representative answers the ninth question of the morning once instead of ten times. The digital banking manager can finally see which questions arrive where, because each assistant keeps its own record.

ABOUT CUSTOMGPT.AI

CustomGPT.ai is a no-code platform for building AI agents that answer only from an organization's own content: documents, websites and knowledge bases. Every answer is grounded in verified sources with citations, and the platform is SOC 2 Type 2 and GDPR compliant. Organizations use it for customer support, internal knowledge search, website chat and more, across help centers, apps and the API.