MYDAS

CASE STUDY · MYDAS / MYRIA

The AI could talk. The hard part was knowing when to stop.

A Romanian phone agent that handles the conversation while prices, commitments and alternative appointments stay under human control.

ROMANIAN VOICE AGENT LIVE TELEGRAM APPROVAL TWILIO + ELEVENLABS
TESTED ON REAL CALLS

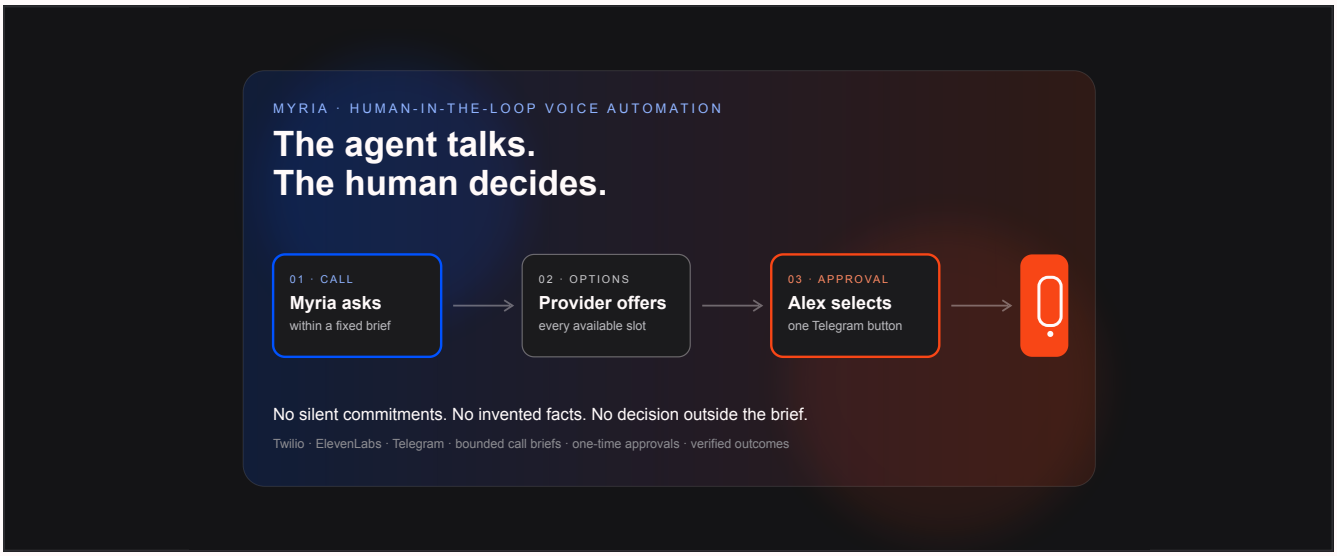
THE DESIGN PROBLEM

Natural conversation does not equal authorized action.

A provider may reject the requested time, offer two alternatives, misunderstand a name or propose work above the approved budget. A model can respond naturally to all of that. The difficult part is deciding what it may commit to.

Myria lets AI handle language and ambiguity. Deterministic workflow rules hold the authority boundary.

> The agent may negotiate the conversation. It may not expand the brief.



100%
FINAL BOUNDED-SERVICE SIMULATION

88s
LIVE SAME-CALL HUMAN APPROVAL TEST

28
APPROVAL BROKER TESTS PASSED

30
BRIDGE TESTS PASSED AFTER POST-CALL DELIVERY

WHAT BROKE

The failures became the specification.

01 **The refusal was too soft**

An unauthorized 800 RON repair was not rejected clearly enough. The protocol changed to require an explicit spoken refusal. The next simulation confirmed only the approved 250 RON inspection.

02 **18:00 became 16:00**

A timezone conversion changed the appointment. The exact local schedule became a structured value that the model is prohibited from converting.

03 **The agent kept arguing**

When 18:00 was unavailable, Myria kept pushing for it. She now accepts confirmed unavailability, collects the alternatives once and stops.

04 **The option list was never confirmed**

Myria repeated two alternatives correctly, but did not ask whether the list was complete. Provider confirmation is now mandatory before Telegram sees the choices.

05 **Success sounded robotic**

The workflow worked, but repeated fillers and formal read-backs made the call unpleasant. Forced phrases were removed after listening to the live calls.

Structured facts stay exact

Dates, local times, prices and selected options move through explicit fields instead of conversational memory.

Silence means no

A timeout creates no selection and no commitment. The system does not guess what the human would have chosen.

LIVE PROOF

Simulation first. Then the actual phone.

Same-call approval

The call stayed connected while Telegram presented the alternatives. Alex selected one button. The exact value returned to the conversation, Myria asked for confirmation and closed the call.

Bounded callback

A separate live flow ended the first call, accepted one Telegram choice and made exactly one callback using that selected option.

Conversation latency

One reviewed call measured about 0.49–0.68 seconds to the first model byte and 0.08–0.10 seconds to the first speech byte.

Privacy-bounded outcome

The post-call service reports completion and follow-up actions without names, phone numbers, addresses, diagnoses or raw transcripts.

TWILIO IE1 / DUBLIN | ELEVENLABS | TELEGRAM | PYTHON BRIDGE | AUTHENTICATED WEBHOOKS | ONE-TIME STATE | POST-CALL ANALYSIS

No invented ROI.

The technical evidence is verified. Operating volume, manual handling time and monetary savings have not been measured yet, so this case study does not turn tests into a productivity claim.

> AI handles the messy language. The workflow keeps the material decision under human control.