

One Door, Three Modes: Designing a Multi-Channel Appointment Booking System Driven by Conversational Agents

Vivek Sinha
Shielva

Abstract

Small service businesses such as clinics receive booking requests through WhatsApp, Instagram and Messenger messages, phone calls and web forms, and many have no calendar system at all. We describe the design of ABS, an appointment booking system in which conversational agents take bookings across these channels for many independent businesses. The central decision is that every channel and every booking style passes through one capture operation; three booking modes (a booking link, a fully conversational booking, and an autonomous mode in which the system itself is the business's calendar) differ only in who supplies the fields and in default settings, not in separate pipelines. We describe how availability and conflicts are computed at read time on the business's own clock, why every customer-facing message is editable per business and per language and is never sent with an unresolved variable, how operators can configure agents without being able to read customers' conversations, and four defects that shaped the design, including a language model inventing appointment times and timestamps that silently shifted by the viewer's time zone. This is a design and experience report; we have no controlled evaluation and say so.

Keywords: appointment scheduling, conversational agents, multi-channel messaging, multi-tenancy, localisation, software architecture, privacy by design

1. Introduction

A clinic's front desk receives the same request ("can I come in Thursday afternoon?") as a WhatsApp message, an Instagram direct message, a phone call and a web form. Commercial scheduling tools assume a staff member types the booking into a calendar; messaging-first businesses need the booking to be taken where the conversation happens. Agents built on large language models can hold that conversation, but a booking is a commitment: the agent must not invent availability, must respect the business's hours and time zone, and must leave an auditable record.

ABS combines a scheduling core with conversational agents built on a separate agent platform. This report describes the design decisions we would repeat and the defects that motivated them.

2. One Capture Door

Every booking, whatever its channel or mode, enters through one operation that takes the business, the agent, the source channel, a reference to the source message or call, and the raw fields (customer handle, service, requested time, notes) and returns the booking plus whether it is new. Capture is idempotent on the source reference, so a retried webhook or a repeated message does not create a duplicate.

We distinguish two kinds of problem. Problems known at capture time, such as a missing service or an unparseable time, are stored on the booking. Conflicts that depend on the current state of the calendar, such as a clash with another booking or a day the business has since closed, are computed whenever the booking is read and are never stored, because a stored conflict becomes wrong the moment the calendar changes.

3. Three Modes, One Pipeline

Table 1. The three booking modes differ in who fills the fields and in defaults.

Mode	Who supplies the fields	Defaults
Booking link	The agent sends a signed link; the customer picks a free slot on a web page	Confirmation by staff; external calendar sync optional

Conversational	The agent collects service, time and details in the chat	Confirmation by staff; external calendar sync optional
Autonomous	As conversational	Auto-confirm; the system is the calendar of record

The autonomous mode exists for businesses with no other calendar, often the least digitised. It was tempting to build three subsystems. Treating the modes as front doors onto one capture operation, plus one declared setting and different defaults, kept a single place for validation, conflict checks and notifications.

A booking moves through requested, confirmed, completed, no-show or cancelled. A separate scheduled state means only that an event exists on an external calendar, so the autonomous mode rests at confirmed. Customers receive a signed self-service link in every confirmation and reminder to reschedule or cancel without contacting the business.

4. Agents and the Booking Core

Agents call the booking core through a small set of declared actions: list availability, book, send a booking link, find my bookings, move and cancel. Each action declares its inputs and the variables it returns, and a parity test runs every action and checks that what it returns matches what it declares, so flow authors see exactly the variables they can use. An agent may use these actions only if its business holds the booking feature and the agent is bound to the booking automation.

The agent never states a time the core did not return. Availability is computed from opening hours, staff schedules, service durations, closures and public holidays for the business's country, and the agent's wording around those times is author-controlled text rather than free generation.

5. Messages and Languages Belong to the Business

Early versions shipped fixed phrases for confirmations, reminders and follow-ups, and assumed a fixed pair of languages. Businesses wanted their own voice and their customers' languages. Every outgoing message is now a per-business, per-language template whose built-in text is only a default. The languages a business supports, and its default language, are chosen when the business is set up rather than written into code, and no user-interface logic contains a language code. A message whose variables cannot all be resolved is not sent. Outside a messaging platform's customer-service window, confirmations go out as pre-approved message templates.

6. Configuring Agents Without Reading Conversations

Two parties touch a business's agents: the business's own staff and the platform's operators, who design dialog flows, intents and knowledge for many businesses. We separate them by design. Businesses connect their own messaging channels and are the only ones who read their customers' conversations. Operators configure agents, see whether a channel is connected and healthy, and may test voice agents, but cannot read customer conversations, and there is no emergency override that grants such access. Channel behaviour is tested through a simulator that performs a dry run: it executes the agent and the booking actions without writing bookings or sending messages. Roles within a business (owner, manager, front desk, staff) govern who can connect channels and who sees which calendars.

7. Defects That Shaped the Design

- **Invented availability.** When a business had not enabled bookings, the availability action returned nothing and a language-model step produced plausible times ("10 AM, 2 PM or 4 PM"). Actions now return an explicit, author-visible sentence for the unavailable case and the agent repeats it rather than improvising.
- **Timestamps read in the wrong zone.** The database returned times without a time zone. They were UTC, but serialised without an offset, a browser reads them as local time; in India every "just saved" item showed as six hours old. Timestamps are now marked as UTC at the serialiser, the single owner, rather than corrected in the interface. Slot arithmetic is done on the business's own clock and converted only for display.

- **A uniqueness rule that was too strict.** A unique index meant to apply only to contacts with a phone number also applied to contacts without one, so only one such contact could be saved per business. Tests caught it; the index now applies only when the field exists.
- **Hard-coded conversation text.** Described in Section 5; the fix was structural (templates per business and language) rather than adding more phrases.

8. Scope and Limitations

- Deposits and payments are out of scope by decision; waitlists exist but are off by default.
- Reading Instagram and Messenger messages depends on the messaging platform's app review; WhatsApp was the first channel proven end to end.
- Class bookings are not yet connected to the waitlist, and package sales cannot yet be voided.
- We report no controlled measurements of booking completion, no-show reduction or staff time saved; these require a study across businesses, which we have not run.

9. Conclusion

The decisions we would repeat are few: one capture operation behind every channel and mode; conflicts computed when read, not stored; agents that may only say times the core returned; messages and languages owned by each business; and a hard separation between those who configure agents and those who may read conversations. Most came from a defect that made the alternative visibly wrong.

References

1. V. Sinha. Building a production telephony voice agent on a cascaded pipeline: An experience report. Zenodo, 2026. doi:10.5281/zenodo.23243533.
2. P. Lewis et al. Retrieval-augmented generation for knowledge-intensive NLP tasks. In *Advances in Neural Information Processing Systems* 33, 2020.
3. ISO 8601-1:2019. Date and time: Representations for information interchange, Part 1: Basic rules. International Organization for Standardization.
4. A. Cavoukian. Privacy by design: The 7 foundational principles. Information and Privacy Commissioner of Ontario, 2009.