

TRAVEL & TRANSPORTATION

The travel and transportation phone line

What 3.7 million AI-agent conversations and 240,000 analyzed human-agent calls across public transit, cruise, and vehicle rental reveal about how travelers call, how those calls end, and where conversational AI can create the most value.

ANONYMIZED INDUSTRY INSIGHT

FOUR DEPLOYMENTS

TRAILING 12 MONTHS

OCTOBER 2026



IN BRIEF

- 1 Travel and transit AI agents run in **two distinct operating models**. Self-service deployments resolve **35 to 50%** of conversations inside the agent. Front-door deployments hand off about **80% by design** and log under 3% as resolved. Blended, about 22% of conversations resolve and 65% are handed off, so resolution rate alone misleads.
- 2 Roughly **one conversation in eight is abandoned** (3 to 18% by deployment), and the leak is concentrated where callers must authenticate or supply data: those failures end **6 to 12%** of conversations in the two self-service deployments.
- 3 The human line behind the front door is **booking-led and long**. At the cruise deployment, booking, sales and add-ons are **31% of tagged reasons** and changes and trip logistics another **25%**, on calls with a median of about **8 minutes**. This is the work the front door routes today and the largest room to automate next.

IN THIS REPORT


SECTION 1	Travel AI agents run as self-service resolvers or as front-door routers	02
SECTION 2	Self-service conversations run longer and lose more callers mid-flow	03
SECTION 3	Behind the front door, booking and trip changes drive more than half of human calls	04
SECTION 4	Where to focus first: a priority map	06
SECTION 5	Two agendas, one playbook	07
APPENDIX	Methodology & scope	08

SECTION 1

Travel AI agents run as self-service resolvers or as front-door routers

Across four deployments and 3.7 million AI-agent conversations, outcomes split cleanly by operating model. The two transit deployments are built to resolve: they complete a trip or account request inside the conversation for 35 to 50 percent of callers. The cruise and rental deployments are built to triage: the agent captures intent and routes the caller to the right human team, so about four in five conversations end in a planned handoff and almost none are logged as resolved by the agent. Neither is wrong, but the two cannot be compared on resolution rate, and a blended rate hides the difference.

EXHIBIT 1
How AI-agent conversations end, by deployment



Share of AI-agent conversations by harmonized final outcome, trailing 12 months. Blended figures are equal-weighted across four deployments with deployment minimum to maximum in parentheses; bars show each deployment; right column is conversation volume. Source: Replicant unified conversations mart, AI-agent conversations only (n=3.66M).

SECTION 2

Self-service conversations run longer and lose more callers mid-flow

Roughly one in eight conversations is abandoned, and abandonment is highest, at about 18 percent, in the two self-service deployments, where callers who start a flow can leave before it completes. Authentication and data-collection failures are a visible share of that loss: in the self-service deployments, between 6 and 12 percent of all conversations end on a failed authentication, a failed data collection, or invalid user data. Front-door conversations are far shorter because the agent only identifies intent and hands off.

~13%

Of conversations abandoned, blended (3 to 18% by deployment)

6 to 12%

Of self-service conversations end on an authentication or data-collection failure

14–124 sec

Range of median conversation length across deployments

EXHIBIT 2

Length and leakage by operating model

DEPLOYMENT	MEDIAN LENGTH	ABANDONED	AUTH OR DATA FAILURE
Self-service, transit A (voice)	111 sec	18.4%	5.9%
Self-service, transit B (voice and chat)	124 sec	17.7%	11.7%
Front door, cruise (voice)	64 sec	10.6%	not separately labeled
Front door, vehicle rental (voice)	14 sec	3.4%	not separately labeled

Median duration and share of AI-agent conversations, trailing 12 months. Authentication or data failure combines authentication hang-ups and escalations, data-collection failures, and invalid-user-data transfers; shown per deployment, not blended, because only the two self-service deployments label these outcomes. Source: Replicant unified conversations mart.


SECTION 3

Behind the front door, booking and trip changes drive more than half of human calls

Conversation Intelligence analyzes the human-agent calls at the cruise deployment, which shows what the front door routes to people. Grouping the tagged reasons into six categories, booking, sales and add-ons are the largest driver, followed by changes, cancellations and trip logistics. Together they are more than half of tagged reasons. Payments and billing add another tenth. These calls are long, with a median near eight minutes, and callers are generally satisfied, but one in four is only partially resolved. This is a single-deployment spotlight and is reported separately from the blended rates above.

EXHIBIT 3

Booking and trip changes are over half of what reaches a human



Share of tagged call reasons; a call can carry more than one tag, so shares are of tags (n=441K), not of calls. Single deployment, human-agent calls analyzed by Conversation Intelligence (n=243K calls). Source: Replicant Conversation Intelligence, call reasons mart.

SECTION 3 · CONTINUED

Callers are generally satisfied, but one in four human calls is only partially resolved

66%

Of analyzed human-agent calls fully resolved; a further 25% partially resolved

7.0/10


Average caller sentiment; 7.5 on resolved calls, 4.4 on disconnected calls

~8min

Median human-agent call length (average ~12 min)

EXHIBIT 4

How analyzed human-agent calls end



Share of calls by final disposition, single deployment, human-agent calls (n=243K). Automated quality scores use a per-program rubric and are not compared across deployments. Source: Replicant Conversation Intelligence.


SECTION 4

Where to focus first: a priority map

Positioning the call categories by automation readiness against volume points to a sequence. Changes and trip logistics, payments, and promotions are structured, data-backed requests that can be resolved in the conversation. Booking and sales are the largest volume but need qualification and a warm handoff before they are ready to automate. Unclear intent is a fix-first item: it drives escalations at the front door. Complaints and requests for a named representative stay with people.

EXHIBIT 5

Automation priority map: readiness versus volume



■ Servicing intents

■ Sales and growth

□ Fix first or keep human

Bubble size ≈ share of tagged reasons

Illustrative positioning by Replicant based on reason mix, call outcomes, and data availability at the cruise deployment, cross-checked against outcome patterns in the transit and rental deployments; axes are directional, not to scale. Source: Replicant Conversation Intelligence.

SECTION 5

Two agendas, one playbook

The opportunities are consistent across deployments and fall into two agendas. Where the agent already resolves, the prize is closing the authentication and data-collection leak. Where the agent routes, the prize is converting planned handoffs into resolutions for the structured requests that make up most of the human line.



SELF-SERVICE AGENDA

Fix the entry gate. Shorten authentication and make data collection forgiving, since 6 to 12% of conversations end there.

Trip and account status in-conversation. Look up and answer, with a graceful fallback when the lookup fails.

After-hours deflection. Resolve or schedule when the human team is closed instead of ending the call.

Voice and chat parity. Offer the same resolved flows on both channels.



FRONT-DOOR AGENDA

Resolve before routing. Payments, confirmations, and status answered by the agent, with handoff only on exception.

Self-service changes. Modifications, waitlists, and cancellations are a quarter of human reasons.

Qualify and warm-transfer sales. Capture intent and trip basics, then pass context to a sales agent.

Sharper intent capture. Reduce escalations from unrecognized or overly broad requests.

METHODOLOGY & SCOPE

Blended outcome figures cover AI-agent conversations at four production deployments (two public transit, one cruise, one vehicle rental; ~3.66 million conversations) over the trailing 12 months, from the unified conversations mart. Each deployment contributes equally, so no single program drives a headline, and the deployment minimum-to-maximum range is shown. The rental deployment is small (~6,000 conversations) and still carries equal weight. Final-disposition labels differ by deployment and were harmonized into resolved, handed off, abandoned, and other; handoffs labeled planned handover, intent identified, or account status were treated as by design, and transfers labeled post-flow success (about 2 to 4% at the transit deployments) are counted as handoffs, not resolutions. Authentication or data failure combines authentication hang-ups and escalations, data-collection failures, and invalid-user-data transfers, and is reported per self-service deployment only. The human-call spotlight is a single cruise deployment, analyzed by Conversation Intelligence since late October 2025 (243K calls, 441K tagged reasons), reported separately and not blended; reasons are grouped into six categories, and a call can carry more than one tag. Quality and sentiment scores use per-program rubrics and are not compared across deployments. Hospitality, hotel-booking, and airline deployments were considered and excluded because the warehouse held too little usable volume, and roadside and auto-club programs are out of scope. The priority map is an illustrative synthesis; its axes are directional, not to scale. Numbers are rounded and banded. No deployment, client, or individual caller is identified, and no underlying data left the analytics environment.

NO-COST ANALYSIS

Skip the demo. Start building your AI roadmap.

Get these insights for your contact center with a no-cost Conversation Intelligence analysis.

[LEARN MORE →](#)