AI agents for the work around the journey and the stay

Agents that read your policies, your fare and rate rules and your booking records, prepare the work, and stop for a named person. On infrastructure you control.

THE PROBLEM

The hotel templates we opened send replies and compensation offers to guests with approval only an optional extra, and none of the vendor pages we read offers on-premises or customer-cloud hosting. Of 3,942 operations in 102 public interface descriptions from three publishers, 2,286 are sent as changes, so tools are chosen one by one.

HOW IT WORKS, IN THREE STEPS

STEP 01

Connect what the agent may read.

Fare rules, conditions of carriage, cancellation policies and brand standards become knowledge sources, with each reader's permissions kept. Bookings, guest profiles, room status and flight records are reached through the description of an interface you already run; audit and feedback data through a reporting database, read only.

STEP 02

Decide where it stops.

Every tool call passes your policy. A tool that changes a booking, a reply or a refund waits for a named person.

STEP 03

Run it where the data lives.

In your data centre or your private cloud, on a model you serve.

THE RULE EVERY AGENT FOLLOWS

It prepares. A named person decides.

Thirty-five agents in nine areas

Guest and passenger service	Guest Enquiry Response. Policy Q&A. Complaint Response. Special Assistance Request Check
Reservations and bookings	Booking Change Preparation. Group Booking Check. Booking Discrepancy. Refund Eligibility Check
Disruption and compensation	Disruption Notice Drafting. Compensation Claim Check. Rebooking Options. Disruption Review
Sales, groups and events	Group and Event Proposal. Corporate Rate Request Response. Event Order Check. Partner Contract Summary
Revenue and distribution	Rate Parity Report. Pickup and Pace Report. Channel Content Check
Hotel and property operations	Procedures Q&A. Shift Handover. Maintenance Request Triage. Brand Audit Preparation
Airport and ground operations	Operations Log Summary. Turnaround Delay Report. Baggage Claim Check. Ground Handling Compliance Check
Loyalty, marketing and reviews	Guest Feedback Analysis. Review Reply Drafting. Loyalty Query Response. Campaign Brief
Finance and people	Commission Statement Check. Invoice Matching. Night Audit Exception. Training Content

Also in the pack: the full list of what each agent reads and where it stops, and the white paper “AI agents inside the airline and the hotel”, with sources.

Seven questions to ask before an agent touches a booking or a refund

01 Where does it run, and where is the model?

02 Does it keep each person’s permissions?

03 Which tools change a booking, a reply or a refund, and does each wait for a person?

04 Where is the step at which a named person approves?

05 Can it answer a guest unseen, and what does it do when it does not know?

06 What is recorded, and for how long?

07 Can we stop one agent, one tool, or all of them, at once?

WHERE THE PLATFORM RUNS TODAY

In production at a top-tier bank in Asia, air-gapped, inside the bank’s own perimeter.

MODON built its AI Centre of Excellence on Katonic inside the Kingdom of Saudi Arabia.

ePLDT, the ICT arm of the Philippines’ largest telco, selected Katonic to power Pilipinas AI, a sovereign AI cloud hosted in the Philippines.

Where to start

One property, one station, invented data first. A policy Q&A agent can run on the first day, because no system of record is touched.

KATONIC.AI/SOLUTIONS/TRAVEL