AI agents inside the retailer and the restaurant

What they can read, what they can change, and who decides

IN BRIEF

01 Retailers have put assistants in the hands of their own people faster than in front of their customers. One European grocer says more than 80,000 of its colleagues already use its store assistant, in more than 50,000 conversations a week. An American chain piloted a store assistant in about 400 stores and then took it to nearly 2,000.

02 In front of the customer the record is mixed. One burger chain ended its drive-through voice test after running it in more than 100 restaurants. Another measured 86 percent of orders taken without help from the crew, which leaves one order in seven for a person. A supermarket removed a script from its assistant after customers said it claimed to be human.

03 This paper is about the work between those two: the reading, checking and drafting done in head office, in the store and in the restaurant. It sets out what an agent can read, how it reaches the systems that run the business, what it may change without a person, and what the rules already say about prices, staff, allergens and chatbots. It ends with seven questions to put to any supplier, and with a place to start.

04 The position throughout is simple. An agent prepares. A named person decides. The agent runs on infrastructure the company controls.

IN THIS PAPER

01 Where retailers and restaurants are today 2
02 The work nobody has packaged 3
03 Three doors, one gate 4
04 What an agent may do 6
05 What the rules already ask 6
06 A worked example: one restaurant group 8
07 Seven questions before an agent touches a store 9
08 Where to start 9

PART 01

Where retailers and restaurants are today

Over 80,000

colleagues at one Dutch grocer use its assistant for associates

Over 50,000

conversations a week are started with that assistant

86%

of orders in one chain's drive-through pilot were handled without a crew member stepping in

Over 100

restaurants ran another chain's drive-through voice test before it ended

Staff assistants work, and the numbers are public. Albert Heijn, the Dutch grocer, reports that over 80,000 colleagues use its assistant for associates and start more than 50,000 conversations a week; an earlier test found most questions were about products: where they are, whether they are in stock, the price, the promotion [1]. Target built a store assistant in-house that answers process questions and coaches new team members, piloted it in about 400 stores and announced it for all of its nearly 2,000 [2]. Coles, in Australia, reports more than 15,000 staff questions a month to an assistant that answers on leave and rostering, and estimates more than 30,000 hours saved a year [3].

Voice ordering has been harder. McDonald's ended its drive-through voice test with its technology partner in June 2024, after running it in more than 100 restaurants; an analyst put its accuracy in the low-to-mid 80 percent range [4]. Wendy's reported that its own pilot averaged 86 percent of orders handled without a crew member stepping in, and describes the technology as "an assistant – not a replacement" for the crew [5].

Customers notice when an assistant pretends. In February 2026 Woolworths, in Australia, removed scripted small talk from its customer assistant after users reported that it claimed to be a real person [6]. In 2023 a New Zealand supermarket's recipe assistant suggested a mixture that would have produced chlorine gas [7].

A company is held to what its assistant says. In 2024 a Canadian tribunal ordered an airline to honour what its chatbot had told a customer, with the sentence a retailer's counsel will remember: "It makes no difference whether the information comes from a static page or a chatbot" [8].

What the published accounts leave out. Of the retail and restaurant deployments we read, almost none says where the assistant runs, and almost none describes a person approving what it produces.

On the suppliers' side, none of the retail agent product pages we read offers to run on the customer's own infrastructure.

PART 02

The work nobody has packaged

21 of 52

jobs done inside a retailer or a restaurant operator, about, appear in none of the template libraries

None

of the workflow templates we opened holds a step for a person

We listed 52 jobs done inside a retailer or a restaurant operator and looked for each in the template libraries of the agent and workflow platforms. About 21 appear in none.

WHERE TEMPLATES EXIST	WHERE NONE WAS FOUND
Order status and returns	**A refund or goodwill credit with an approver**
Product questions and recommendations	**Replies to public reviews**
Table booking and order taking	**Franchise brand-standards audits**
Competitor price collection	**Franchisee sales and royalty reporting**
Low-stock alerts and automatic reorder	**Delivery and invoice discrepancies**
Abandoned basket messages	**A review of a finished promotion**
	Till variance and chargebacks
	Lease and landlord documents
	Food safety records

The templates cover the shop window: the customer's question, the basket, the reorder. They thin out behind it, in finance, property, compliance and the running of a franchise.

Two more findings matter. None of the workflow templates we opened holds a step for a person: purchase orders go to suppliers and emails go to customers as soon as the workflow decides. And nearly all lean on a public cloud model, messaging or spreadsheet service, which settles where customer and sales data goes before anyone has asked.

That is where the reading is slow and the decision already belongs to a named person, which makes it the safer place to begin.

PART 03

Three doors, one gate

An agent is only as useful as what it can reach. In a retailer or a restaurant operator that means three kinds of thing.

Documents.

Store procedures, the returns policy, the allergen matrix, brand standards, the franchise manual, supplier terms, leases. These are indexed once as knowledge sources, with each reader's permissions kept, so an agent answers with the document it relied on and says when the documents are silent.

Systems of record.

Orders, stock, the catalogue, the till, the service desk. Most are reached through an interface the company already runs or its supplier publishes. The company gives the platform the description of that interface, chooses which operations become tools, and decides which of them need approval.

Data for analysis.

The data warehouse and reporting databases, read only.

All three pass one gate: every tool call is checked against the company's policy, which can allow it, record it, hold it for a person, or block it.

SIX WAYS INTO A RETAILER'S SYSTEMS

	ROUTE	WHAT IT REACHES	WHAT TO KNOW
01	**A file the user uploads**	An invoice, a visit form, a lease, a delivery note	Nothing is connected. The safest start
02	**A knowledge source**	Procedures, policies, the allergen matrix, the franchise manual	Permissions are kept per reader
03	**An imported description of an interface**	Orders, stock, catalogue, till	You choose the operations. Changes wait for approval
04	**A reporting database**	Sales, stock, labour and waste figures	Read only
05	**A server from the tool catalogue, run in your cluster**	Databases, web search, document conversion	It runs inside; check what each tool does
06	**A supplier's own tool server**	A supplier's commerce or point-of-sale system	Some exist for commerce platforms; few for restaurant systems, and some run only in the supplier's cloud

PART 03 · CONTINUED

24
interface descriptions taken from their publishers' own addresses; the importer accepted all 24

1,436
operations the importer found in them

867
of those operations change something; 569 read

995
operations are carried by three files alone

What is published. Commerce is well described. We took 24 interface descriptions from their publishers' own addresses, among them a payments and point-of-sale company, three commerce platforms, a restaurant point-of-sale company, two European marketplaces, the supply chain event standard, a food products database and India's open commerce network, and gave them to the platform's importer. It accepted all 24 and found 1,436 operations.

What that exercise taught us. Retail interfaces are built to change things. Of the 1,436 operations, 867 change something and 569 read. Three files alone carry 995 operations, and one of them holds 211 that change something. A care team answering "where is my order" needs a handful.

What is not published. Some of the largest systems give an importer nothing to take: one widely used commerce platform has moved to a different kind of interface, the big merchandising and restaurant suites keep the description on each customer's own installation, and the delivery marketplaces open theirs to approved partners only. There the company supplies its own description, or reads from its data warehouse.

Card data. A system that could affect the security of card data falls within the card industry's security standard [9]. The safe design is the one the newer checkout standards already use: the agent hands over to a hosted payment page and never holds a card number.

So choose the operations at import, start with the ones that read, and add nothing that changes a price, an order or a refund until the approval step has been seen working.

PART 04

What an agent may do

One rule runs through every agent in this paper: it prepares, and a named person decides.

01	**It reads** what the person using it could read, and no more.
02	**It drafts** the reply, the report, the claim, the order, with each statement tied to its source.
03	**It lists** what is missing, what does not match, what falls outside the rule.
04	**It stops** before anything that changes a record or reaches a customer, a supplier or a colleague. A tool that changes something waits for approval.
05	**It decides nothing about a person.** It does not issue or refuse a refund, change a price, allocate a shift, rate a colleague, or name anyone in a loss report.

Three lines are drawn tighter than the rest.

Allergens.

An agent that answers a question about a dish reads the answer from the operator's own allergen matrix and cites the row. When the dish or the ingredient is not there, it says so and sends the question to the manager. It never works an allergen out from a recipe name.

Staff.

An agent may check a draft roster against working time rules and the labour agreement and list each breach. It does not allocate shifts and it does not score people.

Prices.

An agent may recommend a price with its reasons. A merchant approves, and where the law asks for the customer to be told how a price was set, that is for the company to do.

PART 05

What the rules already ask

This part needs a lawyer's read before it is relied on. It reports what the texts say; it is not advice.

Telling the customer

An assistant must say what it is. The European Union's AI Act requires that people "are informed that they are interacting with an AI system, unless this is obvious" [10].

A personalised price must be declared. European consumer law requires a trader to tell the consumer, where it applies, "that the price was personalised on the basis of automated decision-making" [11].

A reduction must show the real earlier price. The same law requires that "any announcement of a price reduction shall indicate the prior price applied by the trader", meaning the lowest price in at least the previous 30 days [11]. An agent that drafts a promotion needs that figure, not last week's.

The whole price, and real reviews. In the United Kingdom the total price must include "any fees, taxes, charges or other payments" the customer has to pay, and a review must be "based on a person's genuine experience" [12]. An agent may draft a reply to a review. It must never write one.

PART 05 · CONTINUED

Staff

The AI Act lists as high-risk any system used "to allocate tasks based on individual behaviour or personal traits or characteristics or to monitor and evaluate the performance and behaviour of persons" at work, with duties that apply from 2 December 2027 [10] [13]. It bans, already, AI that infers the emotions of people in the workplace [10]. A company that employs people must tell them and their representatives before it uses a high-risk system at work.

Food

In the European Union, allergen information is mandatory for food sold loose, in a restaurant or over a counter [14]. In the United Kingdom, food packed on the premises for direct sale must carry the name, the ingredients and "any of the 14 allergens emphasised in the ingredients list" [15]. Whatever channel the answer travels through, the operator answers for it.

Decisions about people, and where data sits

MARKET	WHAT WE READ	WHAT IT ASKS
European Union	AI Act; data protection regulation; consumer law	Say it is an AI; a right not to be subject to a solely automated decision with significant effect; declare a personalised price
United Kingdom	Digital Markets, Competition and Consumers Act 2024	Total price shown; no fake reviews
Australia	Privacy Act amendments, from 10 December 2026, read on the regulator's page	Say in the privacy policy where a computer program makes or informs decisions that significantly affect people
Saudi Arabia	Personal data law, read in law firms' summaries; draft amendments of October 2026	Conditions on moving personal data out of the Kingdom; the draft would require it to be stored inside
India	Data protection rules 2025; dark patterns guidelines 2023	Duties phased in to mid-2027; no hidden price elements
Singapore	Guidelines on personal data in AI recommendation systems; agentic AI framework	Consent or a stated exception for recommendations; checkpoints where a person approves

We did not open the primary text for the United Arab Emirates, and for Saudi Arabia we read summaries. A company in those markets should ask its own counsel.

PART 05 · CONTINUED

What the rules have in common

01	**The customer knows** when they are dealing with an AI, and how a price shown to them was set.
02	**A person decides** what significantly affects a customer or a colleague.
03	**The company answers** for what its assistant says, about a refund or an allergen alike.
04	**The company knows where the data is,** and in some markets that is moving toward inside the country.

PART 06

A worked example: one restaurant group

Take an operator with several hundred restaurants, some franchised, a call centre, an app and a head office. Nothing here is a customer's story; it is the list in section 4 put to work.

01

In the restaurant, a crew member asks whether a dish contains sesame. The agent reads the operator's allergen matrix, answers from the row for that dish, and shows the row. Asked about a limited-time item that is not in the matrix yet, it says so and tells the crew member to ask the manager.

02

In the area office, the agent reads a month of temperature and cleaning records for twelve restaurants and lists the gaps and the readings outside the limits in the food safety plan. The food safety lead decides what to do about each.

03

In the call centre, the agent groups the month's complaints and reviews into themes by restaurant and quotes the comments behind each. Operations staff check the themes and choose what to change.

04

In head office, the agent sets each delivery note against its purchase order and invoice and drafts the claims. An accounts lead raises them.

05

For franchisees, the agent answers questions from the operations manual with the section cited, and passes anything about the agreement itself to the franchise team.

Every one of these reads documents or records the operator already has, changes nothing by itself, and leaves a record of who asked and what was shown.

PART 07

Seven questions before an agent touches a store

01 **Where does it run?**
On our infrastructure, in our private cloud, or in the supplier's service? Where is the model?

02 **What can it read?**
Does it keep each person's permissions, or read as one all-seeing account?

03 **What can it change?**
Which tools change a price, an order, a refund or a roster, and does each wait for a person?

04 **Who decides?**
Show us the step where a named person approves, and what the agent does when nobody has.

05 **What does it tell the customer?**
Does it say it is an AI, and what happens when it does not know?

06 **What is recorded?**
Who asked, which tool ran, what was decided, and for how long is that kept?

07 **Can we stop it?**
One agent, one tool, or all of them, at once, by us.

PART 08

Where to start

Pick work where reading is slow, the decision stays with a person, and no customer is waiting on the agent.

Questions on your own procedures

Colleagues ask what the procedures say and get the answer with the procedure cited. No customer data is involved, so it can run on the first day.

Complaints and reviews, grouped

The agent groups a month of feedback into themes and quotes the comments behind each. Operations staff check and decide.

Delivery discrepancies

The agent sets the delivery note against the order and the invoice and drafts the claim. A person raises it.

One team, one region, invented data first. Then beside the manual process, measured against it.

About Katonic AI

Katonic AI builds the platform that organisations use to build, run and govern AI agents on their own infrastructure: in their own data centre, in a private or sovereign cloud, or with no connection to the internet at all. The platform is in production with governments, banks, telecom operators and utilities. Katonic is certified to ISO 27001.

KATONIC.AI

SOURCES

Read on 21 November 2026. Figures are the publishers' own and were not audited by us.

01 Ahold Delhaize newsroom, "Transforming retail: 10 questions about how AI is empowering associates", 4 September 2026. newsroom.aholddelhaize.com

02 Target press release, "Target to roll out transformative GenAI technology to its store team members chainwide", 20 June 2024. corporate.target.com

03 Microsoft Source Asia, on Coles and the mycoles Assistant, 24 October 2025. news.microsoft.com

04 CNBC, carried by NBC News, "McDonald's to end AI drive-thru test with IBM", 17 June 2024. nbcnews.com

05 Wendy's, "Drive-thru innovation: Wendy's FreshAi", 11 December 2023. wendys.com

06 RNZ, "Australian supermarket giant reins in AI assistant claiming to be human", 27 February 2026. rnz.co.nz

07 New Zealand Herald, on the Pak'nSave meal assistant, 11 August 2023. nzherald.co.nz

08 Moffatt v. Air Canada, 2024 BCCRT 149, British Columbia Civil Resolution Tribunal, 14 February 2024, paragraph 27.

09 PCI Security Standards Council, PCI DSS version 4.0.1, June 2024, sections 2 and 4.

10 Regulation (EU) 2024/1689 (AI Act), Article 50(1); Annex III point 4(b); Article 5(1)(f); Article 26(7).

11 Directive (EU) 2019/2161, inserting Article 6(1)(ea) of Directive 2011/83/EU and Article 6a of Directive 98/6/EC.

12 Digital Markets, Competition and Consumers Act 2024 (UK), section 230(4) and Schedule 20 paragraph 13.

13 Regulation (EU) 2026/1744, Official Journal 24 July 2026.

14 Regulation (EU) No 1169/2011, Article 44(1)(a) with Article 9(1)(c).

15 Food Standards Agency, labelling guidance for prepacked for direct sale food, 17 June 2021. gov.uk

16 Template libraries of agent and workflow platforms, searched 21 November 2026; and 24 interface descriptions fetched from their publishers the same day (Katonic research).