



Retail Hive Mind Guide



# Laying the Foundations for Agentic AI

## Retail Hive Mind Guide

Your introduction to preparing for Agentic AI: how to govern and manage product and customer data

Supported by:

**rackspace**  
technology®

Written by Ed Lawson,  
Head of Content,  
The Retail Hive



# The Problem


Retail leaders are gearing up to embrace agentic AI: autonomous systems that can interpret context and act on the consumer's behalf. The possibilities and promise of this new era come with a host of uncertainties and understandable caution. In Hive conversations on the subject, it's very evident that retailers worry about disintermediation with the customers they have fought so hard to engage and convert cement as 'loyal'.

Automating engagement brings concerns around 'authenticity' - retailers want shoppers to feel aligned with their ethics and principles, and to engage as personally and meaningfully as possible. The concept of agent seems to stand against these pillars, creating considerable reticence. Brand value is everything, so retailers are treading carefully.

Operational readiness is another sticking point. Many processes still rely on workarounds and custom fixes, highlighting the importance of data stewardship and governance.

*"It feels like the next big advancement in AI, but how do we take it from hype to safe, useful practice?"*

**Retail Hive member**



And finally, there's the existential question of relevance. As discovery shifts from traditional search to LLMs and agents, retailers risk invisibility if their content isn't geared up for this new reality.

The question isn't *whether* agents will arrive (they're already here) but *how* retailers should be ready to meet them capably and safely.

# Current Trends



- **The rise of personal agents** - many shoppers will soon rely on their own AI agents to transact on their behalf, instructing them to find a product at a certain price, weigh up shipping costs and vendor reliability, and even delay purchase until conditions are right. For retailers, this shifts the customer journey from active search to delegated intent.
- **Experimentation at the edges:** most retailers are running low-risk pilots in customer service, content creation and marketing automation.
- **Shadow adoption** – employees are hacking generative AI into workflows without approval, highlighting both opportunity and risk.
- **Standards race:** Model Context Protocol (MCP) and Agent-to-Agent (A2A) are emerging as the technical 'rails' that will underpin agents interactions.
- **Personalisation:** progressing from segmented messaging to a more genuinely personalised offer, using AI to extrapolate preferences and needs.

# Avoiding the Pitfalls

Retailers are running into a familiar set of issues as they experiment with agentic AI. Weak data foundations remain the most common blocker. Fragmented systems, inconsistent definitions and poor data quality limit what agents can safely do and make explainability difficult when results are challenged. At the same time, hype can drive premature adoption. **Teams are tempted to deploy agents before agreeing what outcomes matter**, how success will be measured or whether the organisation is ready to support them beyond the initial pilot.

Shadow adoption adds another layer of risk. Pilots often move faster than governance, with tools connected to data sources before ownership, scope and controls are clear. Early results can look promising, but performance drifts as prompts evolve, data changes and workflows shift. **As Atif Sheikh, Strategy Consultant at Rackspace puts it, “this is a lifetime cycle... not a fire and forget solution.”** Without guardrails, retailers risk inconsistent outputs, unclear lineage and growing operational and compliance exposure.

Finally, malicious or unintended traffic is becoming a real concern. As agentic browsers and automated agents grow, retailers need to think about how AI interacts with their sites, APIs and offers. **Poor controls can expose promotions, distort demand signals or create new security and fraud risks.**

## How to avoid

- ✓ Create an owner for each priority data object and outcome so there is a single accountable ‘captain’.
- ✓ Set access rules early, defining which datasets AI can use, for what purpose and in which layer.
- ✓ Build a data quality feedback loop so poor inputs are corrected rather than normalised.
- ✓ Put day two ownership in place with monitoring, review cadence and model changes tracked.
- ✓ Treat unapproved tools and connectors as a control issue, not a training gap.



# What the Experts Say

“There is no such thing as an IT strategy, a technology strategy or a data strategy. There is only the outcome you’re trying to achieve. Agentic AI points to a future where payment providers play a far more active role in the shopping journey. If agentic tools start purchasing on behalf of customers, retailer websites may become less central, with engagement shifting towards media platforms and payment providers’ agentic ecosystems..”

**Benjamin Greenaway**  
Head of IT & Development, The Fold




"Agentic AI challenges retailers to protect and express brand values in environments they no longer fully control. As the traditional ecommerce journey gives way to AI mediated decision making, brands must rethink how they build trust, relevance and emotional connection with customers beyond their own channels."

**Tom Hedges**  
Senior Project Manager, AllSaints

# Implementation Tips



- **Clean your data** Tidy up product hierarchies, item reference files and backend records before layering agents on top.
- **Experiment in safe zones** Start with internal uses like content drafting, research, and marketing automation where brand risk is low.
- **Structure content for AEO** Use schema.org tagging, FAQs, and detailed answers to nuanced questions.
- **Leverage existing platforms** Work within Shopify, Salesforce, or other ecosystems with built-in guardrails.
- **Harden governance** Put in place controlled sandboxes, clear SOPs, and strict infosec policies before scaling.
- **Create an LLM.text file** This emerging standard acts like a sitemap for agents, guiding them on how to interact with your site safely and effectively.



## Checklist:

- Are your product hierarchies clean?

---

- Do you have tight SOPs for returns and refunds?

---

- Is your content structured for agentic search?

---

- Have you defined acceptable error rates?

# Structure and Governance

Strong governance starts with clear separation between foundations, preparation and outcomes. For Rackspace, this structure is essential if agentic AI is going to scale safely and deliver real value.

1

## Bronze layer: Foundation and control

### What it is

Raw data landed from source systems with consistent handling, access controls and lineage.

### Why it matters

Without this foundation, compliance and explainability break down.

### Ownership focus

Central data and IT teams own standards, security and access guardrails.



Atif Sheikh,  
Rackspace

*Storing customer data in multiple inconsistent ways quickly becomes a business risk.*

2

## Silver layer: Quality and consistency

### What it is

Data that has been cleaned, deduplicated and aligned to shared definitions.

### Why it matters

This is where trust is built. AI models are only as good as the quality of the data they learn from. Garbage in still means garbage out.

### Ownership focus

Shared responsibility. Central teams enable tooling and rules. Business teams validate meaning and usability.

**rackspace**  
technology®

3

## Gold layer: Workloads and outcomes

### What it is

Curated datasets assembled for specific business goals such as personalisation, churn reduction or service automation.

### Why it matters

Agents should only access data that is relevant to their purpose. Selective access reduces risk and improves results.

### Ownership focus

Business owners take accountability. As Rackspace sees it, every data object should have a 'captain', and ownership should sit where the impact lies.



# Key Challenges

## Supply Chain Team

Supply chain leaders point to fragmented systems and a wide mix of available technologies, which make it hard to keep pace with AI-enabled expectations.

## IT Teams

Many retailers, particularly smaller or mid-sized, or with limited resources, find that staff are already using AI tools independently, creating governance headaches for overstretched IT teams.

## Customer-Facing teams

Customer-facing teams are worried about operational efficiency pressure, fraud, and brand risk if automation gets it wrong. And for those relying on search to drive discovery, the shift to agent-based recommendations is both exciting and threatening.



**The consensus:** these challenges are shared across departments and retail segments, from food to fashion, logistics to luxury

# Next Steps

1

## **Start small and deliberate:**

- Launch controlled pilots in low-risk areas.
- Monitor outcomes closely to learn and iterate.

2

## **Prioritise data hygiene:**

- Fix broken Standard Operating Procedures (SOPs).
- Standardise definitions across teams and systems.
- Clean and reconcile backend records to ensure readiness for agentic integration.

3


## **Prepare for discovery shifts:**

- Optimise content for Answer Engine Optimisation (AEO).
- Adapt content for agent-based search experiences.

4

## **Stay connected to emerging standards:**

- Track developments in Merchant Capability Profiles (MCP) and Agent-to-Agent (A2A) protocols.
- Leverage the Hive community to compare notes and share best practices.



# Cheat Sheet:

## Agentic AI Dos & Don'ts for Retailers



### Dos:

- Do start with low-risk internal pilots before scaling to customer-facing use cases.
- Do clean and standardise data across systems before introducing automation.
- Do optimise content for AEO to stay visible in agent-driven discovery.
- Do set clear governance guardrails and sandbox experimentation.



### Don'ts:

- Don't assume agents will 'work it out', flaky SOPs will break automation.
- Don't rush into vendor partnerships without due diligence; some won't last the hype cycle.
- Don't overlook security: treat malicious bots as the default, not the exception.
- Don't lose the human touch. Design jobs where people and AI complement each other.





# Hive



**rackspace**  
technology®

Rackspace Technology empowers today's technology, data, AI and business leaders with secure, AI-first multicloud strategy and delivery.

As architects of Fanatical Experience™, we integrate application modernization, data governance, security resilience and FAIR™—our Foundry for AI—across hybrid and hyperscale cloud platforms.

We guide organisations through ideation, incubation and industrialization of responsible AI at enterprise scale, accelerating outcomes with governance, automation and cloud-smart workload placement. With proven leadership recognized in Gen-AI services and cloud managed services, plus global reach across 120+ countries, Rackspace delivers agile, future-ready infrastructure that drives measurable business transformation.

