# Retail insights 2026

## Practical actions retail leaders can take to prepare for the technology landscape of tomorrow

/thoughtworks

Design. Engineering. AI.

# Introduction

There's never an easy time to be a retailer. The challenges remain the same — attract and retain customers, find the best way to balance shopper demands and supply chain pressures — but the environment is constantly changing.

Right now, the conversations we have with Thoughtworks' retail clients are dominated by unreliable supply chains, soaring customer acquisition costs, and a market where customer attitudes are difficult to predict. They're looking for ways to navigate this landscape and boost profitability, all against a backdrop of rapid technological evolution.

Across the retail industry, AI is expanding its footprint. Some areas of retail businesses are already very familiar with AI-supported workflows. Supply chain, for example, is a department that's frequently ahead on machine learning deployment thanks to early advances in predictive analytics and ample opportunities for complex optimization. Others are just beginning to experience its potential, which is constantly evolving as generative and agentic capabilities develop. Gartner predicts that 40% of enterprise apps will include a task-specific agent by the end of 2026, up from just 5% in 2025.

This isn't about making futuristic projections or speculation. We're not here to tell you about the trends that may or may not materialize; we're here to help you focus on what's directly ahead of you. This report outlines 10 realistic, practical moves retail leaders can make now to take advantage of what's already available and create a foundation for what's coming in 2026 and beyond.

Scaling AI in retail ultimately comes down to three imperatives:

- Rebuild the core systems and data your business depends on.
- Rewire operations so intelligence can move safely across workflows. Consolidation of subsidiaries to gain higher efficiencies.
- Reimagine value by embedding AI into products, services and experiences.

When these three shifts are executed in a coordinated way, retail organizations move beyond pilots and deliver AI that works.

# Build a foundation for transformative technology

The first priority for retail leaders is to rebuild, strengthening the core systems and data foundations that daily operations depend on. With AI dominating the technology conversation, many retailers are pushing hard to embed it in their operations. Implemented in the right way, it's transformative, but there are early risks that can derail an AI strategy before it has a chance to deliver any value. When organizations start provisioning AI solutions, there's a fine line to walk between too specific and too general.

On the specific side, some organizations end up building a collection of unintegrated solutions that only work in isolated use cases. Say their team creates an agent that monitors fabric supplier pricing, a generative AI chatbot for handling returns, and a visual search tool for their online store. There's no connection or cohesion; they may work individually, but there's no way to use the underlying models or processes for other use cases or to compose use cases into larger workflows. This leads to uneven progress across the business and yet more systems that can't talk to each other.

On the other side of the line are those that build too generally. In these organizations, data and technology teams create speculatively designed solutions that follow theoretical best practices and look great in the planning stages. But they aren't designed to meet specific needs. So, the organization winds up with an overengineered system that takes a long time to build and can't serve any real-life use cases.

The best way to steer between these two failure modes is through a combination of two techniques: thin slicing and thinking of data as a product.

## Thin slices: rapid value through incremental work

Thin slices are narrow, end-to-end increments that enable retailers to rapidly create and deploy innovative customer experiences and colleague tools. Thin slices start providing value rapidly and, because they're actively in use from very early on, users can provide real feedback that allows teams to iterate and improve quickly.

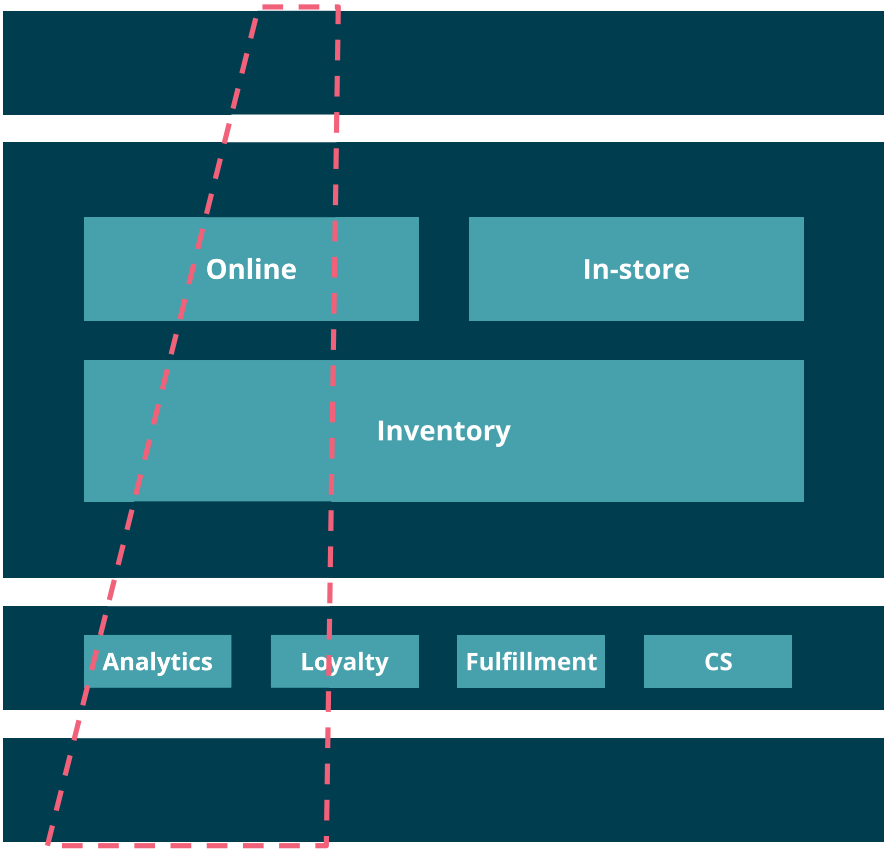


If retailers want to adopt this approach, though, their data and data infrastructure might be a barrier. When systems are disconnected and unable to share data, AI can become just as siloed as any other tool.

However, a unified data platform can form the foundation for hundreds of use cases by integrating data from all the different business channels, including point of sale, warehouse, e-commerce, marketing and supply chain. It's vital for retailers that are exploring the potential of agentic AI.

But retailers shouldn't simply create a pool of unstructured data; it must be curated and responsive. The platform needs to be federated, allowing different departments to contribute value independently rather than relying on a central data team. Data should be organized according to its domain to ensure that the central 'data brain' can make new connections as needed.

Agentic systems need reliable product data, such as titles, attributes, images, pricing, and availability, that they can parse and reason over. The reality is, many retailers have unreliable product data, with duplicate SKUs, inconsistent attributes, missing specs, and mismatched images across ERP, e-commerce, PIM, and supplier feeds, which can lead to incorrect recommendations, stockouts, and mispriced items.

## Data as a product: development driven by the use case

One of the key components of the 'data as a product' approach is to start with a use case and work backward to identify and organize the information needed to serve that use case. This helps ensure each thin slice is supported by the data required to deliver the target outcome, whether that's a customer service feature or an internal tool.

When data is managed like this, retailers can support new use cases by simply assembling capabilities and information in different ways. It's both faster and more efficient; Forrester reports that organizations

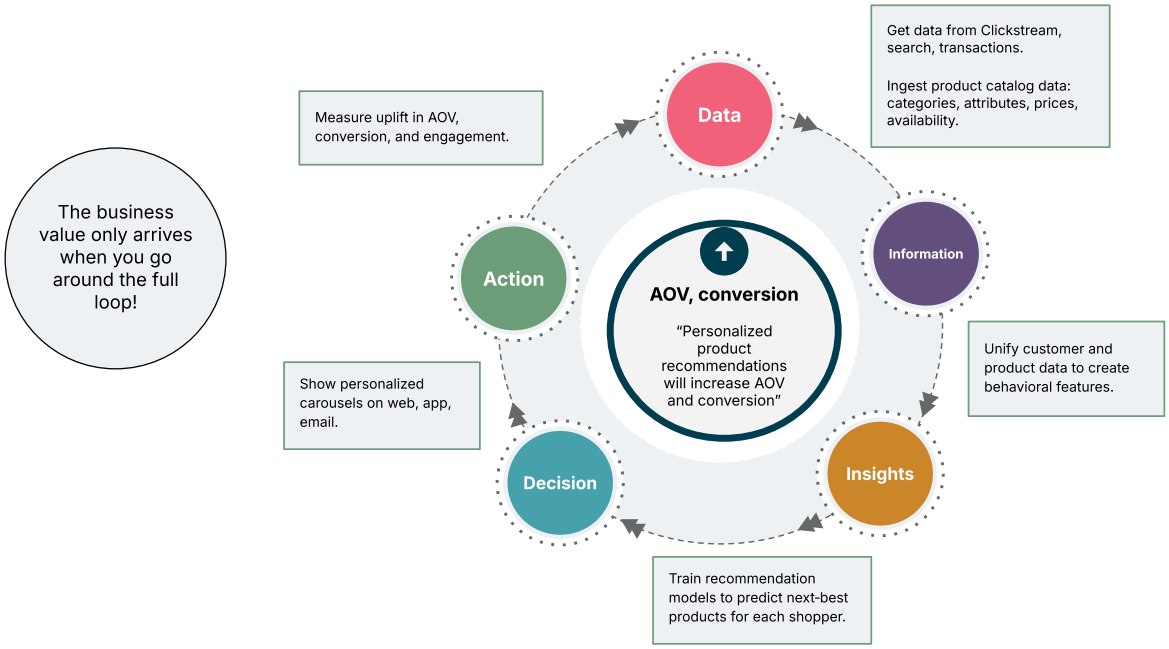


with a product-centric data architecture can deliver analytics as much as 60% faster, for example. Thin slices can be expanded and adapted using this foundation, steadily building an AI-enabled infrastructure across the organization.

Rebuilding the core is not a one-time modernization initiative. It is the ongoing discipline of ensuring systems, data and architecture remain stable, adaptable and ready for intelligent automation at scale.

# Embed AI strategically with an outcomes-driven approach

Once the core is rebuilt, the next imperative is to rewire, making data, processes and systems accessible so intelligence can operate consistently across the business.

Few AI projects successfully make the leap from proof of concept to an effective, embedded component of the business. Researchers report that 95% of generative AI pilots are currently failing to deliver any value, and 42% of organizations scrap most of their AI initiatives before they reach production.

The complexity of implementing AI doesn't come from the AI model itself. For most organizations, model development and iteration are handled by AI research and product companies like OpenAI or Anthropic. The real difficulty comes from finding the right approach and ensuring the retailer can provide the high-quality context and tools — including the data foundation — that the model needs to be effective.

One strategic approach to implementing AI focuses on identifying a small set of foundational capabilities that can support a wide variety of innovative user experiences. A great first step is to look at a broad list of capabilities that the organization is interested in and see what foundations they share, as shown in the figure below:

**Insights: platform readiness falls into buckets**

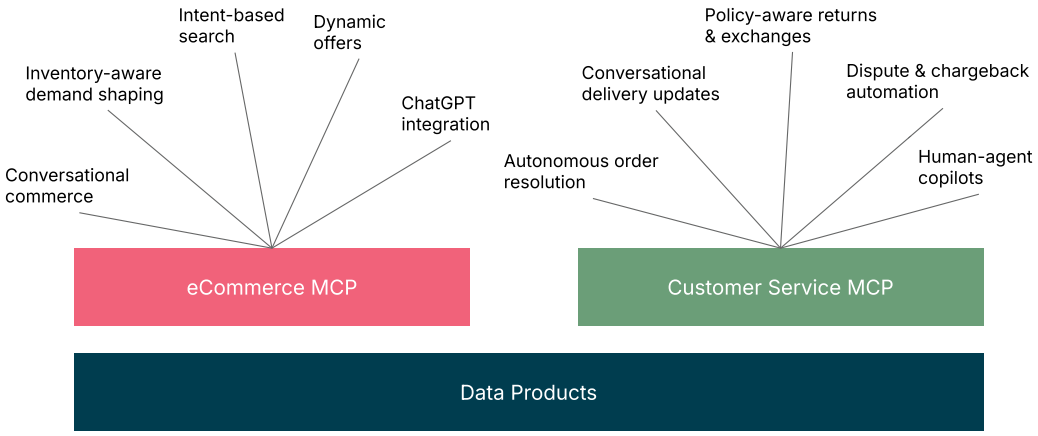

By building a collection of base components that can be mixed and matched in different ways, retailers can accelerate the development and deployment of new AI-driven experiences. This also prevents teams from independently producing the same components and duplicating effort, while ensuring that every new capability integrates seamlessly into the foundation.

## Becoming AI-first is a transition, not a switch

Many organizations — in retail and other industries — are aspiring to become 'AI-first' businesses, and this combination of thin slicing and data products can enable that shift. But it's not a simple switch that can be flipped. Becoming AI-first is most practical and effective when it's approached in two phases: first through a transitional model and then onto a target operating model.

The target operating model might be an environment where tasks and workloads are orchestrated seamlessly between automation and humans, with very little manual intervention required to manage the AI. Human expertise is layered on top of a robust set of proven AI workflows to maximize the value of both. At this level of maturity, it's completely intuitive to rebalance the mix of responsibilities as models improve or business needs change, and easy to introduce new, integrated capabilities.

But in most organizations, today's AI capabilities and foundations aren't yet ready for that.

A transitional AI-first model uses the thin slice approach to build general capabilities, which can then be refactored to answer new demands and serve as many use cases as possible. It starts with small amounts of automation that are closely monitored by human experts for validation and oversight. Over time, as the AI becomes more embedded into the business and increasingly shaped by its processes and knowledge, more tasks can be routed through to create that target orchestration engine.

Rewiring is what transforms isolated experiments into connected, scalable capabilities. It enables intelligence to move across workflows rather than remaining confined to individual tools or departments.

# Create a culture of experimentation and role-driven upskilling

Rebuilding and rewiring create the conditions for change. The final imperative is to reimagine, redefining how value is created when humans and intelligent systems work together.

If retail organizations want to achieve seamless orchestration between AI and humans, it's not just the technology that must be improved. This shift is driving a need for new skills in the retail workforce to ensure colleagues can work effectively with their AI-enabled tools.

It's important for supporting adoption rates — because no one wants to use a tool they don't understand — and to futureproof the workforce as the retail world continues to change and roles and responsibilities evolve.

In light of this, leaders may think that the best action is to set a goal to upskill the whole organization in AI. But that's far too fuzzy to be actionable. 'Working with AI' isn't a single, discrete ability; training must be contextual and flexible to match the broad and varied nature of the technology.

## Empower employees to discover their own AI skills

Of course, just as with any technology, there's a hygiene level of helping everyone understand the basics, such as what an agent is. But with so many different ways to apply AI, from workflow automation to interactive agents, there's no neatly packaged training course that will comprehensively teach people to use AI effectively.

Instead, retailers should create a culture where people actively think about their individual roles and how AI could benefit them. Which processes could be delegated to an agent? Where could faster insights or data-led advice be useful? How could the colleague use their time in more productive ways if they were supported by AI? The value of using AI will be very different for someone in a fashion merchandising role, for example, compared to someone who manages a CPG warehouse.

Retailers should provide both opportunities for individual targeted learning and routes for teams to explore together, especially for departments that work in tandem with each other.

AI literacy is a foundational skill. The mix of competencies will vary from role to role and department to department, but everyone in the business should have a good understanding of:

- The tools they're expected to use and how they function
- Where their expertise or oversight is needed
- The systems and safeguards that are in place to monitor the AI
- When and how to push back on AI-driven decisions or recommendations
- The role they play in the continuous improvement of the retailer's tools

Reimagining value means rethinking roles, workflows and customer experiences, not simply automating existing processes, but designing new ones around AI's capabilities.

# Design for trust and transparency and test rigorously

Trust is the heart of loyalty — and loyalty is everything in retail. To enable innovation, teams should be able to experiment with AI and adapt AI-driven workflows for their use cases. But at the same time, there must be a central structure that ensures everybody is working toward the same approved goals and within accepted parameters.

Without trust, neither rebuilding nor rewiring can scale. Governance and transparency ensure intelligent systems operate within defined boundaries while remaining adaptable as business needs evolve.

## Lax oversight can put retailer reputations at risk

It's easy to test deterministic actions — this button does this — but more complex interactions can be more difficult to assess. AI, especially newer and more experimental models, can behave in unexpected ways. These hard-to-assess cases are far more common when there's an agent working with any degree of autonomy, or there's back and forth between a system and a human.

However, continuous testing, evaluations of model effectiveness and bias and close monitoring are worth the effort to protect people, data and the business.

Bias, unclear data-handling processes, or a lack of guardrails for models or agents can be devastating for a business. Retailers can't risk a dynamic pricing agent discriminating against certain demographics, or an AI assistant giving an offensive answer to an online shopper's query. Losses, fines, lawsuits and reputational damage can put the entire organization in jeopardy; transparency, auditability and traceability are paramount.

End-user trust in agents comes from:

- **Clarity:** Understanding what the agent is and what it can do
- **Management:** Offering easy ways to adjust, override, and configure
- **Transparency:** Providing previews and explanations
- **Collaboration:** Enabling Fast feedback loops that show the agent will listen and adapt

**Enabling safe development for a leading ecommerce site's AI chatbot**

This Thoughtworks client wanted to create an AI-powered chatbot to give shoppers a new way to discover the latest fashion by providing curated outfit suggestions.

An OpenAI LLM responds to shoppers' queries, while another model finds matching clothing items that fit their requested style. To make sure each iteration of the chatbot behaved as expected and in line with the brand, we built LLM observability and evaluation frameworks.

**Evaluations**

Evals are rigorous quality checks that build confidence in AI models while ensuring they're performing reliably and completing tasks correctly.

They can help teams mitigate common AI risks such as bias, or fix areas that are vulnerable to exploitation. Evaluations aren't one and done; models should be assessed regularly, especially if they're learning and changing or working autonomously.

**Guardrails**

Guardrails provide boundaries for AI, ensuring models, tools and agents don't operate beyond their domain or capabilities and preventing them from producing incorrect or biased outputs.

They can exist at all levels, from training to deployment. An early-stage guardrail may filter out harmful data, to prevent it from informing the model's future outputs. Later down the line, human experts can provide feedback to demonstrate the correct responses to different inputs, or define content moderation rules to prevent inappropriate answers or decisions.

Retailers can bring these complex and automated interactions under control with evaluations, or ‘evals’, and guardrails.

These measures enable retailers to assure shoppers (and legislators and other governing bodies) that their systems are carefully designed, tested and monitored, alongside all the usual reassurances about data protection and privacy.

Evals and guardrails can also help make colleagues more comfortable with the tools they’re using and provide transparency for them, too. Better understanding leads to more effective use, helping retailers extract more value from their AI investments and supporting employees to work more efficiently and confidently.

**Evals and guardrails for retail:**

- Task-success and tool-use monitoring in production
- LLM-as-judge evals to assess quality of responses
- Agent governance: What bots are running? Who owns them? What are they allowed to do?
- FinOps to track token spend
- Ethical frameworks to monitor bias e.g. gender
- Benchmarking to evaluate new LLM models as they are released

# 10 key actions for retail leaders in 2026

Retail is ultimately about delivering experiences that keep shoppers coming back. So everything retailers develop must be focused on providing those high-quality, loyalty-building experiences — whether it's directly facilitating an in-store interaction or supporting the supply chain. More and more of these experiences will rely on AI, from basic analytics to advanced decision engines.

Any retailer can deploy an AI solution. Operationalizing AI requires more than isolated initiatives. It requires rebuilding the core, rewiring workflows and reimagining value as an integrated transformation.

If you want to truly operationalize AI in your retail organization, you need to follow these 10 key actions:

- Steer in the middle to avoid development pitfalls — not too specific, not too general
- Start with thin slices to see impact and value quickly
- Create a data foundation that can support a broad range of use cases
- Focus on reusable components that you can mix and match
- Build solutions based on outcomes to ensure they provide needed capabilities
- Follow a staged transition plan to embed AI incrementally
- Encourage teams to explore AI in ways that suit them
- Help employees develop AI literacy for responsible AI use
- Re-think core business processes using AI, instead of simply automating
- Guide model development with evals and guardrails

Retailers don’t lack AI tools. They lack alignment across foundations, workflows and value creation.

When organizations rebuild, rewire and reimagine in concert, AI shifts from experimentation to sustained operational advantage.

That is how retail organizations deliver AI that works.

Thoughtworks is a uniquely capable partner for retailers navigating the world of rapidly expanding AI capabilities. Driven by decades of expertise, leading-edge design thinking, and engineering excellence, we can help you transform the way your organization approaches development, deployment and everything else that makes the modern enterprise run.

We can guide you through these key actions and onward in your journey. From identifying and creating your first thin slice to deploying complex agentic systems at scale, our consultants and engineers have the frameworks, strategies and delivery expertise you need to thrive in a demanding retail environment.

**Explore our retail and e-commerce success stories to discover how these capabilities translate into real-world results.**

We are a global technology consultancy that delivers extraordinary impact by blending design, engineering and AI expertise.

For over 30 years, our culture of innovation and technological excellence has helped clients strengthen their enterprise systems, scale with agility and create seamless digital experiences.

We're dedicated to solving our clients' most critical challenges, combining AI and human ingenuity to turn their ambitious ideas into reality.

**thoughtworks.com**

/thoughtworks

**Design. Engineering. AI.**