KPMG

Competing smarter:

Modernizing retail operations in the age of AI

Accelerating progress with targeted advancements in data, platforms, and AI.

Introduction

Retail has become bifurcated. Big-box leaders and digital-first players deploy technology and capital at scale, while most retailers must protect loyal customers and brand equity under tight budget, talent, and time constraints.

The lesson from the last two years is that meaningful modernization—upgrading legacy operations to deliver digital-first capabilities—no longer requires a multiyear replatforming effort. When sequenced well, targeted shifts across data, core retail platforms, and AI can deliver measurable outcomes in months and help fund what comes next.

This playbook outlines practical ways to sharpen decisions and improve speed without large-scale transformation. The focus is on outcomes, with examples throughout to help teams move quickly from opportunity to execution.

Targeted pilots are critical for guiding modernization efforts that deliver results quickly. Data upgrades and platform overlays typically show value within months, while AI can generate a cascade of steady improvements. For example, we've seen this targeted approach help retailers improve inventory efficiencies by 20 percent, unlock 10 to 15 percent more capacity, and improve margins by 7 percent through AI-driven pricing precision.

For these initiatives, starting small is a feature, not a limitation. Pilots are narrowly scoped so that teams can test assumptions and validate impact before committing additional time or capital. This phased approach reduces risk and builds momentum, enabling organizations to learn, adapt, and fund improvements from early wins.

Targeted pilots can help retailers:

Improve **inventory efficiencies** by
20 percent

Unlock
10 to 15 percent
more **capacity**

Improve **margins** by
7 percent
through AI-driven pricing precision

The case for modernization

Retailers today face a convergence of pressures. Shoppers demand more responsive service at every touchpoint, while margin compression, unpredictable labor markets, and ongoing supply chain uncertainty require retailers to make faster decisions with less room for error. But meeting these challenges is difficult amid legacy data, systems, and process roadblocks. The result is a widening gap between what customers expect and what many organizations can reliably deliver.

But for leaders who navigate this complexity, the prize is significant. Modernization offers a strategic opportunity for agile mid-market players to leapfrog larger, slower-moving competitors by building a data-savvy organization that can match the capabilities of the industry giants.

Here's a closer look at how to make it happen with three shifts retailers can make today to accelerate outcomes:

1. Activate the data you already have

Most retailers sit on rich operational and customer data—point of sale (POS), e-commerce, loyalty, inventory, and supply systems—but much of it remains siloed. This leads to blind spots that degrade customer promises and force teams to rely on instinct rather than insight. Connecting a few critical flows (such as linking digital inventory to store availability or merging loyalty IDs with POS transaction data) can improve visibility and trust quickly, without rebuilding the data platform.

What to do now

To bypass multiyear timelines and cost commitments, focus on these high-velocity data actions:

Begin by establishing a single source of inventory truth. When systems disagree about what is actually available to sell, cancellations rise, pickup promises wobble, and confidence erodes. Correcting this is critical for stabilizing operations today while also establishing the groundwork required to earn the trust of the AI shopping agents of tomorrow. This data discipline creates myriad benefits. As just one example, we've seen this help retailers quickly identify and reduce low-value items by 20 percent, freeing up working capital for higher-performing products.

In parallel, unify customer profiles. Link loyalty and customer relationship management (CRM) records so that even simple journeys—welcome, lapsing, high-value thank-yous—are based on integrated purchase and browsing history. Personalization has become a primary driver of top-line growth: Retailers that deliver personalization at scale typically achieve 5 percent to 15 percent higher revenue growth than peers, while leaders in AI-driven personalization have seen revenue uplifts of up to 40 percent from improved conversion rates and customer lifetime value.

And don't overlook the value of unstructured data. AI can mine disparate sources such as product reviews and chat logs to rapidly pinpoint quality issues and service gaps—turning raw feedback into high-value wins.

Finance can improve speed and visibility as well. A one-day view into margin and shrink across products, stores, and channels turns backward-looking forensics into real-time levers. Combine core elements—price, discounts, costs, freight, returns, adjustments—into a business intelligence model with exception views that highlight unusual erosion or loss patterns.

Finally, extend visibility to strategic suppliers. Giving partners a live window into sell-through, inventory, and forecast views can improve in-stock rates, reduce rush shipments, and align production more closely to demand.

To illustrate how these principles work in a real-world setting, consider the following example. A national beauty specialty retailer sought to lift repeat purchases without building a full personalization engine. By linking loyalty IDs and CRM to POS receipts and browsing history, the team created unified profiles and launched simple journeys targeted to lapsing and high-value segments. Performance signals fed back into merchandising and pricing. The outcome: more relevant offers, higher engagement, and a measurable lift in repeat purchase—achieved within the target 16-week pilot window.

2. Modernize retail platforms without replacing them

Core retail systems—such as POS, order management, inventory, and fulfillment—shape nearly every customer experience. For most organizations, replacing these platforms is costly and risky. Instead, targeted upgrades and modular overlays can stabilize execution, improve reliability, and enhance customer experience while protecting what works.

What to do now

To stabilize execution without the high cost and operational risk of replacing core systems, use targeted overlays to bridge the gap between digital intent and store reality.

Prioritize high-impact journeys first. Improving buy-online-pickup-in-store performance is one common place to start, with the broader goal of steadily reducing friction between digital shopping intent and store-level execution. Piloting new service routines and lightweight digital tools in a handful of stores helps teams refine operating steps, customer communications, and exception handling before scaling. The goal is to deliver measurable improvements quickly, without disrupting core operations.

In practice, these upgrades typically combine several capabilities that reinforce one another. Mobile checkout can relieve congestion during peak periods, while automated task cues help associates prioritize picking, staging, and

replenishment activities. Crucially, these changes can transform the employee experience. Eliminating manual workarounds empowers associates to focus on the customer, turning them from administrative troubleshooters into effective brand ambassadors.

With lightweight order management system overlays that reveal real-time inventory, order status, and customer context, associates gain the kind of immediate insight that supports stronger judgment on the floor. The workflow becomes less about reacting to issues and more about advising, anticipating, and shaping the interaction—a shift that elevates the role and improves consistency across stores.

Tightening handoffs between e-commerce, store operations, and fulfillment systems also reduces delays and missed pickups, improving reliability and accountability across teams. Together, these modular changes produce measurable gains in customer satisfaction and labor efficiency within a short pilot window.

Using a "digital twin" to unlock capacity and reduce execution risk

The challenge

Retailers and their suppliers are operating with increasingly tight capacity buffers while facing more frequent demand swings, supply disruptions, and new product launches. Planning decisions—where to run production, how to sequence changeovers, and when to shift capacity—are often made with limited ability to test downstream impacts, leading to overtime spikes, stock-outs, and elevated launch risk.

The approach

A digital twin creates a virtual model of the factory and fulfillment network, encoding real-world constraints such as capacities, routings, changeover times, minimum order quantities, and quality risk. Using AI-enabled simulation, teams can model demand spikes, supply interruptions, and new-product ramp-ups across internal and contract manufacturing sites. This allows planners to test scenarios before acting—optimizing where to run what, when, and on which lines—generating contingency plans if conditions change. The digital twin feeds improved capacity and allocations scenarios back into existing planning and scheduling systems.

The outcomes

Based on our client work and observed implementations, digital twins have helped organizations unlock 10 percent to 15 percent effective capacity, improve utilization, and reduce overtime and stock-outs. They also derisk major launches by enabling teams to pressure-test plans in advance—turning high-stakes decisions into modeled scenarios rather than costly real-world experiments.

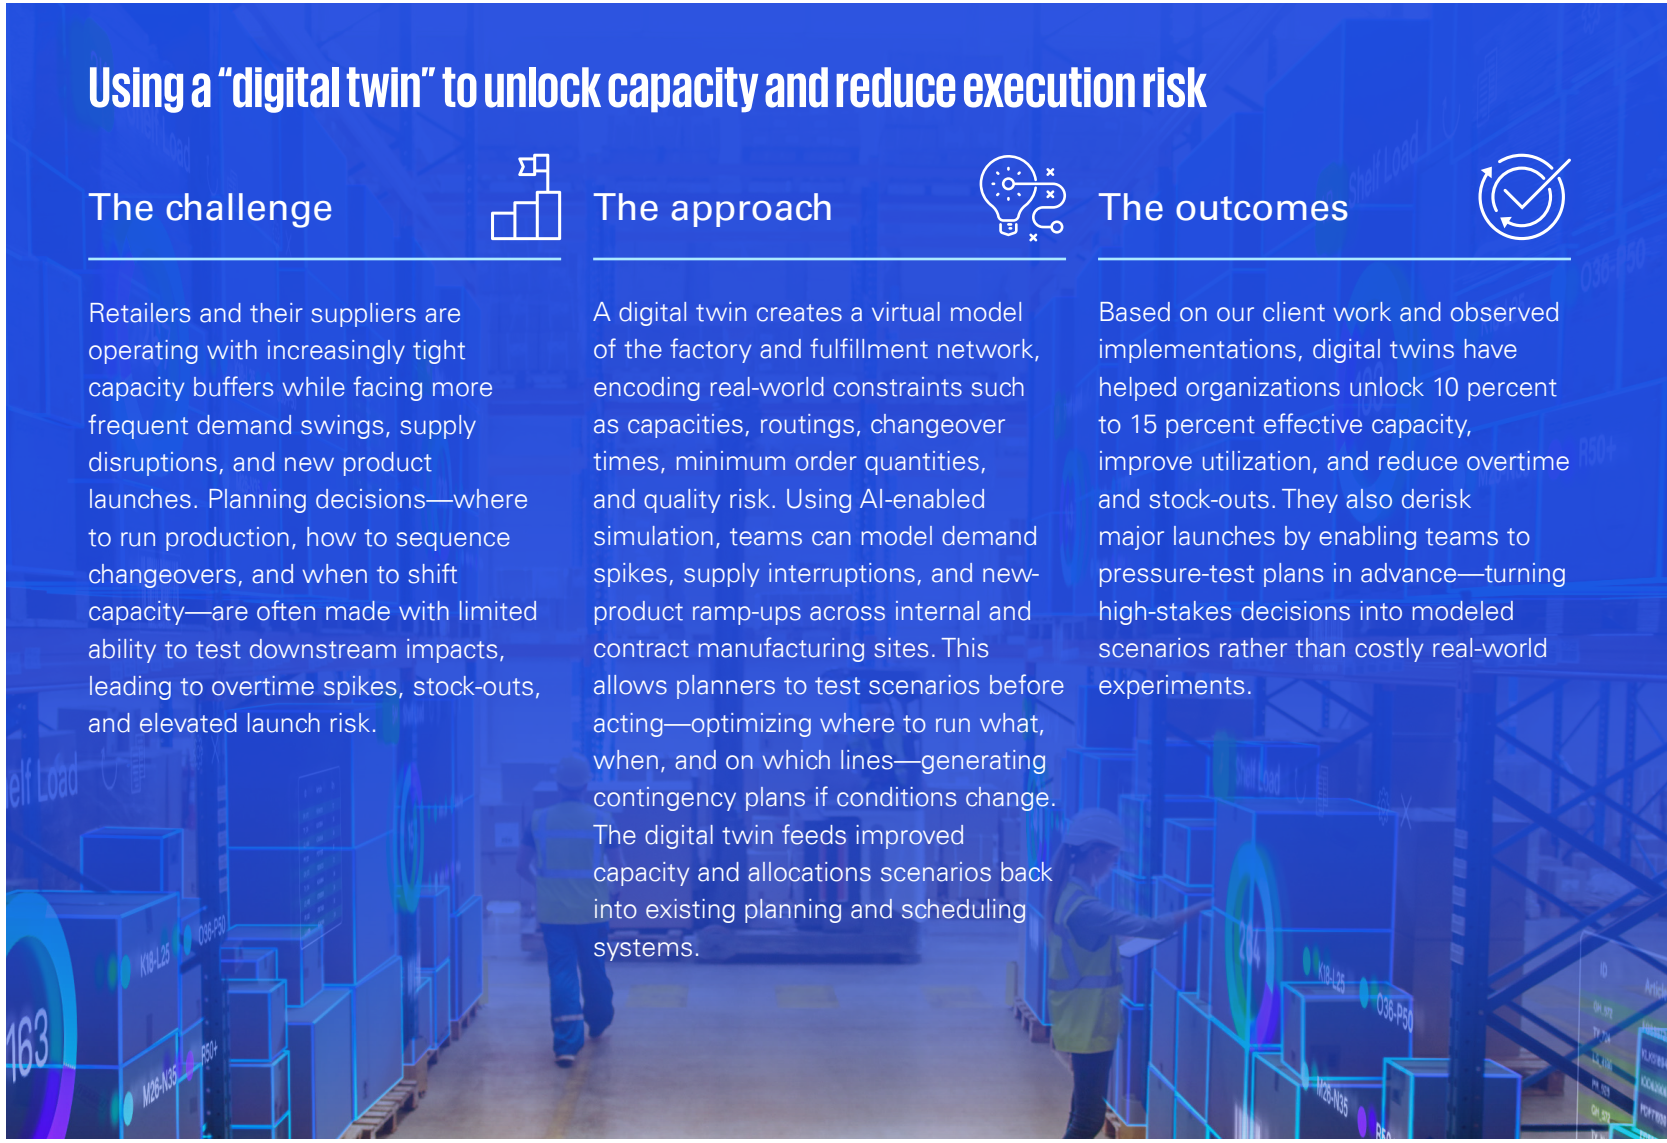

3. Apply AI precisely to enhance the speed, accuracy, and value of core decisions

AI's advantage in retail is precision. Applied to targeted decisions, it improves accuracy and speed in ways teams can act on immediately—without forcing new systems or reengineered processes.

What to do now

To drive rapid value without a massive engineering lift, apply AI surgically to high-stakes decisions where improved accuracy and speed directly boost margins and growth.

Forecasting is a natural starting point. Focus on a single, high-impact category; combine internal signals with external indicators; and use self-healing techniques so the model adjusts as actuals diverge from plan. Surface exception alerts—fast-moving SKUs, stock-out risks, overstated demand curves—so planners intervene earlier, then feed improved forecasts into existing workflows for buys, allocations, and replenishment.

With loyalty and CRM data connected to enhance personalization, predictive scoring can determine which customers are most likely to convert—and which actions matter most. Trigger tailored bundles, lapsing nudges, and individualized promo depth in priority categories. When applied to pricing and promotions, we have seen AI-driven optimization help retailers improve margins by as much as 7 percent by reducing overdiscounting and targeting markdowns more precisely.

AI also serves as a powerful new creative resource. Generative models can help merchandising and marketing teams draft initial product descriptions and promotional copy, accelerating speed-to-market for new items.

In customer experience, GenAI can turn unstructured feedback into structured intelligence that guides product content, service scripts, and operational fixes. And in commercial decisions, pricing and promotion optimization help planners move from blanket markdowns to precise offers that protect margin while shaping demand.

For network design and fulfillment, a "digital twin"—a virtual, AI-powered model of the retailer's fulfillment network—can enable teams to test scenarios for staging inventory, adjusting pickup windows, shifting routing rules, and more. By using AI to simulate how demand, capacity, and operational constraints interact, planners can pressure-test high-effort, high-impact decisions in a relatively low-cost environment before committing to changes in the real world. We've seen this cost-efficient modeling generate 15 percent capacity improvements for retailers.

This digital twin concept extends beyond logistics. Leading brands now use it in the design studio to create virtual product models, and then leverage synthetic customer data to simulate market responses. This allows them to iterate on product development and test for viability faster and more cost efficiently—and long before committing to physical production.

These precise AI applications create cleaner, more reliable data flows. That consistency becomes the foundation for serving future machine customers, which will depend on accurate forecasts and dependable, API-ready inventory signals.

To see this precision in action within core operations, consider a regional grocer that targeted a volatile category to improve allocation. By combining internal signals with relevant external indicators and applying self-healing models, the team reduced both out-of-stocks and excess while improving allocation for priority SKUs. The pilot was completed within the targeted 16 weeks and expanded based on measurable forecast accuracy gains.

Scaling generative AI adoption with a secure foundation

The challenge

A leading global footwear retailer recognized the potential of generative AI (GenAI), but it lacked a clear strategy, secure framework, and internal capabilities to move beyond experimentation. Leadership wanted to enable GenAI adoption across the business while maintaining strong governance, security, and workforce adoption.

The approach

The retailer worked with KPMG to establish a secure, enterprise-grade GenAI framework focused on practical use cases and adoption. This included defining an AI strategy and operating model, selecting and configuring multiple large language models, launching targeted proof-of-concept applications, and pairing deployment with structured learning, change management, and governance. The approach emphasized scaling responsibly rather than pursuing isolated pilots.

The outcomes

In under a year, the retailer deployed a GenAI platform supporting more than 4,500 employees and launched eight GenAI-enabled solutions across customer service, marketing, legal, retail, and other functions. The company also established guidelines for ongoing governance and value realization, enabling GenAI to be embedded into day-to-day work across the organization.

The first 100 days: organizing for speed

The biggest barriers to modernization are organizational silos, not technology hurdles. Success requires breaking down functional walls immediately to create a single, integrated engine for execution. Here are some ways to make that happen.

Days 0–30
Charter the cross-functional pod

Don't treat this as a standard IT project. Stand up a dedicated, integrated team comprising members from data, store operations, merchandising, and finance.

Establish decision rights: Appoint a single owner with the authority to break deadlocks across functions.

Align the resources: Whether physical or virtual, the team must operate outside of standard reporting lines to focus solely on pilot goals.

Define shared metrics: Ensure Finance, Ops, and Merchandising agree on exactly how "success" is measured (e.g., margin versus revenue versus throughput) to prevent disputes later.

Days 31–60
Execute in sprints

Move from monthly reviews to daily execution. The integrated team should operate in one- or two-week sprints to test assumptions rapidly.

Establish a War Room rhythm: Replace weekly status reports with daily stand-ups to triage blockers in real time.

Test cross-cutting workflows: Run pilots (such as AI forecasting or mobile checkout) in the live environment, forcing data, tech, and store teams to solve problems together on the fly.

Days 61–100
Operationalize and self-fund

Once pilots generate data, shift the focus from experimentation to scaling and funding.

Codify the playbook: Translate the "hacks" and workarounds used during the pilot into standard operating procedures for the broader organization.

Pivot to value capture: Use verified financial wins to unlock the budget for the next wave of improvement, transitioning the pilot team into a permanent modernization office.

Risk and governance guardrails

As you run pilots and test new solutions, keep the scope focused and avoid making major changes to your core systems. Use data that's "good enough" for customer-facing promises, but strictly enforce accuracy for key identifiers—such as product or customer IDs—across systems. Set clear checkpoints in advance, with criteria for when to continue or stop a pilot, and review progress every week. From the start, prioritize privacy, security, and responsible use of AI. This reframes governance from a defensive check into a strategic asset. Ultimately, the most resilient retailers will compete and win on trust, making transparent and ethical AI a core brand pillar.

These moves are the first steps toward a new retail model built on disciplined data, an insight-enabled workforce, and AI that strengthens decisions at every level. The retailers that move now will shape the next competitive curve, building the capabilities that intelligent ecosystems and machine customers will expect. The window is open, but it won't stay open for long.

The functional playbook: leading the modernization sprint

Moving quickly from opportunity to execution requires clear ownership. Here's how functional leaders can translate these strategic actions into daily operations.

For the CIO, CTO, and CDO:

Prioritize the layer, not the core: Stand up an API-first orchestration layer that other systems can trust, rather than replatforming the ERP.

Unify signals: Use microservices to centralize inventory visibility and customer identifiers. Prioritize data hygiene where it affects promises (latency, cancellations, phantom stock).

Instrument for speed: Capture specific error metrics, such as pickup accuracy and cancellation rates, and publish weekly exception reports for triage.

For the CFO and FP&A:

Enable the "self-funding" model: Use pilot dashboards to track the specific economics of offer testing and inventory accuracy improvements.

Establish real-time baselines: Define alert thresholds for margin and shrink, and agree on 24- to 48-hour response routines rather than waiting for end-of-month reconciliation.

For merchandising, marketing, and pricing leads:

Start small: Choose one category for AI forecasting and one segment for personalization.

Test with precision: Use elasticity insights to set promo depth and simulate scenarios before launch.

Measure outcomes: Track forecast error, sell-through, and promo ROI explicitly to validate the pilot before expanding to new categories.

For store operations and customer experience leads:

Pilot in the field. Focus mobile checkout and task automation in a select group of stores (10–20) to define service levels and routines before scaling.

Close the loop: Conduct micro-audits to validate phantom stock, and use GenAI to mine unstructured data (reviews/chats) for immediate service fixes.

Track throughput: Measure task completion time, customer wait times, and repeat contact rates to prove efficiency gains.

How KPMG can help

KPMG teams work with retailers to link core data, stabilize platform execution through modular overlays, and apply AI to decisions that move the needle. We tailor the sequencing to the organization's budget and pace, instrument value so pilots fund the next step, and help keep disruption low. Connected capabilities across strategy, operations, data, and technology accelerate execution and scaling. We use these foundational wins to build the capabilities, confidence, and capital required to tackle the truly transformative opportunities of the future: reimagining your business (and workforce) with AI and preparing for an economy of machine customers.

Authors

Brett Benner

US Consulting Leader, Consumer, Retail & Hospitality

bbenner@kpmg.com

Brett is the leader of our Consumer, Retail & Hospitality's Consulting business and brings over 20 years of experience serving our clients. He has deep industry knowledge and has served some of our largest clients as they navigate their complex transformations. Brett helps clients understand how the latest technical solutions and advancements enable organizations across the front, middle, and back office functions. He focuses on helping clients harness these advancements to drive growth, enhance profitability, and create more agile, future-ready organizations.

Sunder Ramakrishnan

Principal, Advisory, Performance Transformation

hsramakrishnan@kpmg.com

Sunder is a principal in the KPMG Consumer and Retail Strategy practice and has 20 years of experience as a consultant and industry executive. He has worked as a strategy consultant and held executive roles at a Fortune 150 retailer. Sunder has deep knowledge in both developing and implementing strategic initiatives across channels and functions in retail, including labor models, e-commerce and omnichannel merchandising, marketing, and operations.

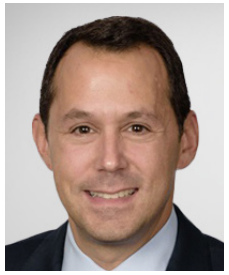

Scott Rankin

Principal, Advisory Practice Leader

scottrankin@kpmg.com

Scott Rankin is the KPMG Advisory practice leader for our Products sectors, which encompasses the Consumer and Retail, Life Sciences, Industrial Manufacturing, and Energy, Natural Resources, and Chemicals sectors. Based in Boston, Scott has a wealth of experience as a strategy and M&A consultant, in addition to senior executive roles at a Fortune 150 retailer and a mobile commerce start-up.

In this role, Scott collaborates with C-level executives to guide strategic transactions and transformations to create sustainable value for clients. His areas of expertise include organic and inorganic growth strategy, operational diligence, operating model transformation, cost optimization, and profit improvement.

Duleep Rodrigo

US Sector Leader, Consumer & Retail

drodrigo@kpmg.com

Duleep serves as the US Sector leader for Consumer and Retail. He has led transformation projects for clients across all phases of the business cycle—from emerging, pre-IPO, platform-based companies to some of the world's largest multinational organizations—by bringing together the best of KPMG's audit, tax, and advisory teams. He frequently engages with C-suite executives and board members to share emerging sector trends, focus areas for investments, transformation, and critical risks and challenges.

Duleep contributes to a variety of KPMG thought leadership, external industry publications and hosts a variety of share forums and webcasts with consumer and retail executives.

Martin Sokalski

Principal, Digital Lighthouse

msokalski@kpmg.com

Martin is the U.S. Advisory AI Leader for the Consumer and Retail sector with over 25 years of experience helping clients capture sustainable value with the help from modern technologies. In today's world of evolving and proliferating artificial intelligence, he helps organizations embrace new opportunities, drive business model transformation, and unlock AI-powered innovation, growth, and operational excellence. Martin guides organizations in adopting and scaling AI, advising on emerging capabilities like agentic commerce and build future-ready organizations.

The authors would like to acknowledge the contributions of Vinay Banari, Karen Henrie, Leah Lockwood, Tara Nelson, and Barry Raghunathan.

Related thought leadership

AI in retail: Global lessons from strategy to storefront

From data overload to data-driven decisions in retail

Intelligent retail

Some or all of the services described herein may not be permissible for KPMG audit clients and their affiliates or related entities.

Please visit us: **kpmg.com** **Subscribe**

DASD-2026-19705