CIKLUM

From Automation to Agentic Decisioning (Autonomy): A Paradigm Shift

The Retail and Consumer Packaged Goods (RCPG) industries stand at the crossroads of a transformative phase, driven by the pervasive and ever-evolving capabilities of Artificial Intelligence. In a landscape characterized by unprecedented market volatility, fragmented and changing consumer demands, and increasing pressure for lean, sustainable operations, AI is no longer a futuristic concept but a present-day imperative for sustained growth and to be ready for what is ahead.

As RCPG businesses enable and adapt to the shift, we see them moving from traditional automation to agentic decisioning, supported by advancements in AI. This evolution marks a transition from systems that merely execute predefined rules to intelligent, autonomous agents capable of making real-time decisions, adapting to dynamic environments, and even initiating actions with limited human intervention. This transformation promises unparalleled efficiency, personalization, and resilience across the business value chain.

AI is Fundamentally Reshaping the Retail Landscape

AI is reshaping the retail landscape -transforming both the customer journey and operational efficiency across the value chain. From the moment a consumer contemplates a purchase to the post-sale experience, AI infuses intelligence at every touchpoint. One of its most visible applications is personalization at scale: AI algorithms analyze vast datasets—including purchase history, browsing behavior, and individual preferences—to deliver hyper-targeted recommendations, dynamic content, and personalized offers that drive engagement and conversion. Beyond the digital realm, AI-powered computer vision and sensor technologies can transform physical stores by providing insights into shopper engagement with displays, product placements, and in-store navigation, allowing for real-time adjustments to optimize layouts and promotions.

Moving to the backend processes, AI agents can optimize sourcing decisions by evaluating supplier performance, pricing, reliability, and risk profiles in real-time. They can even initiate automated procurement decisions, negotiate with suppliers, and suggest alternative suppliers when disruptions or better terms are identified. Instead of simply flagging low stock, agentic systems can autonomously initiate purchase orders, negotiate with suppliers, and even redirect inventory between locations based on real-time demand signals. They dynamically calculate optimal replenishment quantities and trigger orders based on predicted demand and lead times, minimizing stockouts and overstock situations.

For CPG companies, the impact of AI extends across the entire product lifecycle and supply chain. Hyper-accurate demand forecasting can be a game-changer, leveraging deep learning algorithms that analyze millions of historical data points, external factors and social media sentiment, and real-time sales data. This precision can minimize stockouts, reduces excess inventory, and ultimately increases sales by ensuring product availability.

AI can revolutionize supply chain optimization, enabling through predictive modeling and automation. From optimizing logistics and warehouse management to dynamically adjusting reorder points and safety stock levels, AI provides the acuity needed for lean yet resilient operations in the face of geopolitical shifts and unforeseen events. Further, AI is accelerating product innovation and development. AI tools can analyze consumer preferences and market trends to rapidly ideate new product concepts, suggest formulations, and even optimize packaging designs, significantly shortening development cycles. AI-driven market research provides deep insights into emerging trends, allowing CPG brands to stay ahead of the curve and develop products that truly resonate with evolving consumer preferences.

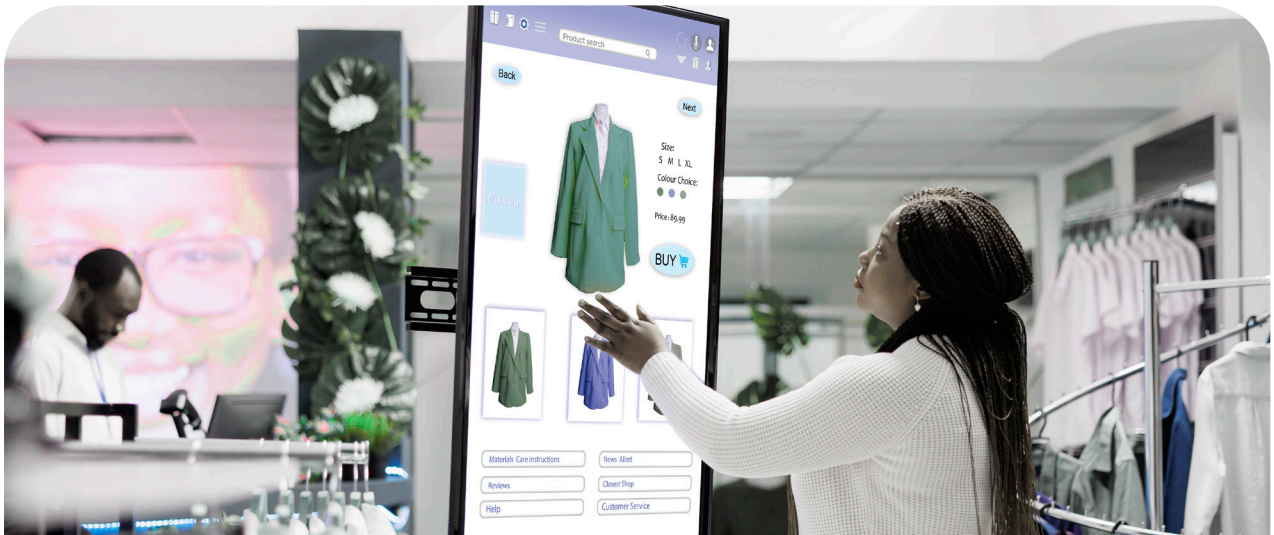

At Ciklum, we believe in taking an AI led business value approach and have partnered with clients across the retail, consumer products, manufacturing, automotive, healthcare, banking and other industries in enabling AI in their day-to-day work processes. Our suite of services

Agentic AI Consulting and Strategy: Providing end-to-end advisory services, from identifying high-impact use cases (using our PRIME framework) and developing an agentic AI roadmap. This includes Proof of Concept (PoC) development and Minimum Viable Product (MVP) creation.

Data Foundation and Integration: Building robust data pipelines and unified data layers that integrate disparate data sources (POS, inventory, CRM, external market data) into a "single source of truth." This is crucial as agentic AI thrives on high-quality, accessible data.

Custom AI / Multi-Agent Development and Orchestration: Designing, developing, and deploying bespoke AI agents tailored to specific business needs. Creating and managing networks of interconnected AI agents that collaborate to achieve complex organizational goals, ensuring seamless communication and coordination across business functions.

Platform and API Integration: Ensuring seamless integration of new AI agent solutions with existing legacy systems, ERPs, warehouse management systems, and other third-party platforms through robust APIs.

Scalability and Governance Solutions: Providing frameworks and tools for scaling AI agent deployments, implementing "human-in-the-loop" workflows for oversight, and establishing robust governance models for ethical AI use, bias mitigation, and data privacy..

Support and Maintenance: Offering ongoing support, monitoring, and maintenance for deployed AI agents, including regular updates, performance optimization, and troubleshooting.

We are pleased to present this Everest Group report examining the industry-wide the industry-wide evolution from rule-based automation to AI-driven reasoning. We explore how Retail and Consumer Packaged Goods (RCPG) businesses can strategically apply AI across their value chains and outline key considerations for scaling these efforts using a structured strategic framework. The shift to AI-led decision-making is not a simple technological upgrade. It represents a fundamental reimagining of how retail and CPG enterprises operate, compete, and deliver value. With deep capabilities in AI, data, systems integration, and digital engineering, Ciklum is well-positioned to serve as a strategic partner on this journey, helping organizations unlock new levels of efficiency, agility, and customer engagement.

CIKLUM

Experiences of tomorrow.
Engineered together.

Contact us:

Everest Group®

September 2025

From Automation to Autonomy: the Agentic AI Shift in Retail and CPG (RCPG)

Contents

Introduction

Retail and Consumer Packaged Goods (RCPG) companies operate amid fierce competition and rapidly evolving consumer demands. Thin profit margins, tight supply-chain constraints, and an ever-growing emphasis on customer experience have made operational efficiency and innovation essential. Against this backdrop, Artificial Intelligence (AI) has emerged as a critical enabler for increasing efficiency, reducing costs, and creating differentiated user experiences, helping RCPG organizations stay ahead in a volatile market.

This Viewpoint presents insights from Everest Group's ongoing research into the evolving role of AI – particularly agentic AI – within the RCPG industry. It is designed for enterprise leaders, strategists, and technology decision-makers seeking to understand how AI is transforming the RCPG value chain.

We discuss:

- A brief overview of AI's evolution – from rule-based systems to predictive Machine Learning (ML) and the emergence of agentic AI
- The landscape of current AI applications across the RCPG value chain
- Use cases that demonstrate active adoption of agentic AI
- Key challenges and risks in adopting AI and agentic AI at scale
- Enterprise implications by AI maturity stage, including build-vs.-buy considerations
- Future outlook on emerging opportunities and shifts in the ecosystem

Enterprise leaders driving AI initiatives can use this Viewpoint to benchmark their organizations' AI journeys, identify high-impact opportunities, and shape forward-looking strategies aligned with their AI maturity and business priorities.

From rules to reasoning: the rise of agentic AI and Systems of Execution (SoE)

AI has transitioned from rules-based systems and machine learning to agentic models capable of autonomous action. Agentic AI differs fundamentally from previous AI waves as it embeds reasoning frameworks like ReAct (Reason+Act) and chain of thought into technology orchestration layers. These components allow agents to not just generate outputs but plan, execute, and course-correct autonomously across multiple decision loops.

This section explores the journey of AI through key milestones, the emergence of agentic AI, and the growing importance of SoE in orchestrating dynamic, self-driven enterprise environments. Exhibit 1 illustrates this evolution, highlighting the shift from early AI systems to today's agentic AI and SoE-led enterprise models.

Exhibit 1: AI's evolution over the years

Source: Everest Group (2025)

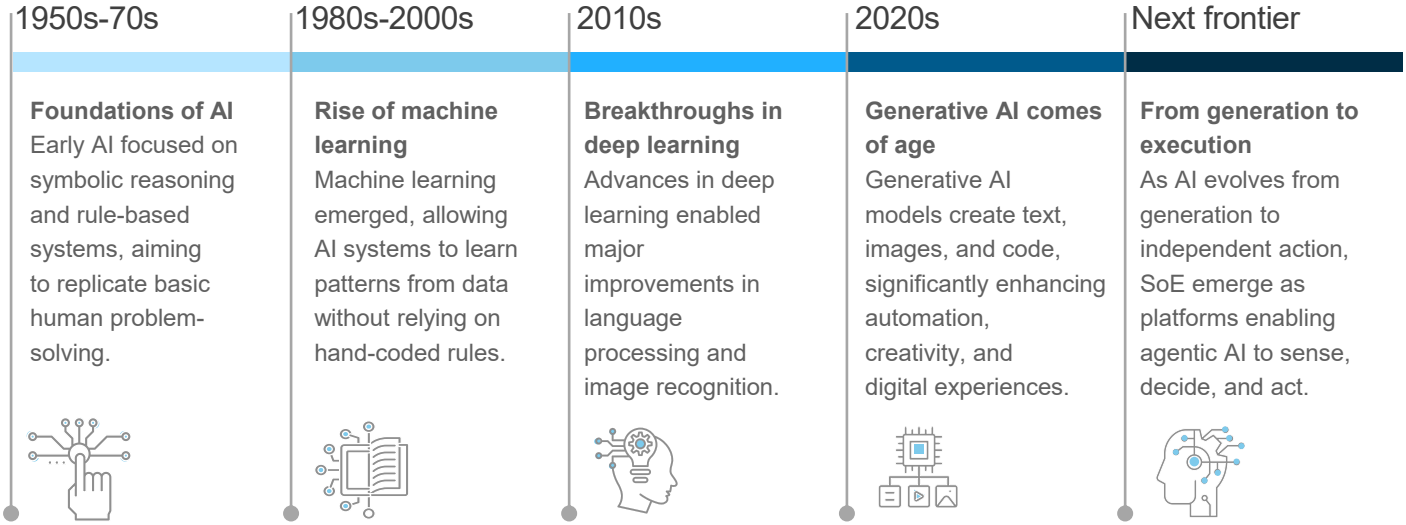

As the timeline shows, AI's evolution is marked by major shifts in how intelligence is applied. Today, we are entering the era of agentic AI – systems capable of autonomous decision-making and self-directed action. Unlike earlier models that simply responded to prompts, agentic AI acts independently, pursuing goals and adapting to changing contexts. This marks a critical shift from systems that assist to systems that act, laying the groundwork for the next wave of enterprise transformation.

Agentic AI

Agentic AI represents the next major frontier beyond generative AI. While traditional chatbots respond reactively to prompts, agentic systems operate with autonomy – pursuing goals, adapting to context, and executing multi-step tasks without constant human oversight. These agents achieve autonomous, goal-driven behavior by pairing reasoning-enabled language models with structured knowledge bases for context, reinforcement learning loops for continual improvement, and Retrieval-augmented Generation (RAG) for on-demand data access. Over time, they create a compounding "data flywheel" effect – learning from every interaction and continuously improving. Unlike generative AI, which creates content based on instructions, agentic AI drives outcomes, optimizing supply chains, enhancing customer experiences, and triggering decisions across complex systems, all without waiting for a prompt.

SoE

To truly unlock the value of agentic AI, enterprises must rethink how actions are initiated and executed across their digital landscapes, shifting from passive data architectures to dynamic SoE, which provide the architectural layer that activates agentic AI. Unlike systems of record and systems of engagement, which passively store and display data, SoE connect data to action using autonomous agents. For example, a lead-to-cash agent may detect buyers' intent, apply pricing logic, and generate contracts without human prompts. Through structured planning loops, SoE allow enterprises to shift from insight latency to execution velocity.

Exhibit 2 presents an overview of systems of record, engagement, and execution.

Exhibit 2: Systems of record, engagement, and execution

Source: Everest Group (2025)

SYSTEMS OF **Execution**

Agentic AI platforms | Pricing/Replenishment bots

Orchestration engines | Decision automation hubs

SYSTEMS OF **Engagement**

E-commerce platforms | Customer portals

PoS systems | Chatbots

SYSTEMS OF **Record**

ERP suites | SCM platforms

CRM systems | PLM tools

AI across the RCPG value chain

As AI evolves from generating insights to driving autonomous, real-time actions, its impact is becoming increasingly tangible across the RCPG value chain. Agentic AI is being applied across procurement, supply chain, sales, marketing, and customer operations. While agentic AI autonomously drives many retail functions, enterprise control remains layered. In returns processing, for instance, humans remain in the loop, approving edge-case exceptions, while shelf pricing agents operate on the loop, alerting supervisors only when anomalies breach thresholds.

Beyond replacing repetitive tasks, AI augments human roles by dividing execution responsibilities:

- **In the loop:** humans approve exceptions (for example, disputed returns)
- **On the loop:** agents autonomously act but escalate boundary conditions (for example, digital shelf pricing errors)
- **Outside the loop:** full autonomy for routine processes (for example, reorder triggers for high-volume SKUs)

This orchestration enables not just efficiency but resilience – humans govern risk while AI maximizes responsiveness. Exhibit 3 illustrates the evolving dynamics of human-AI interaction in agentic AI processes, mapping the shift from manual oversight to autonomous execution across key decisioning stages.

Exhibit 3: Human and AI interaction in agentic AI processes

Source: Everest Group (2025)

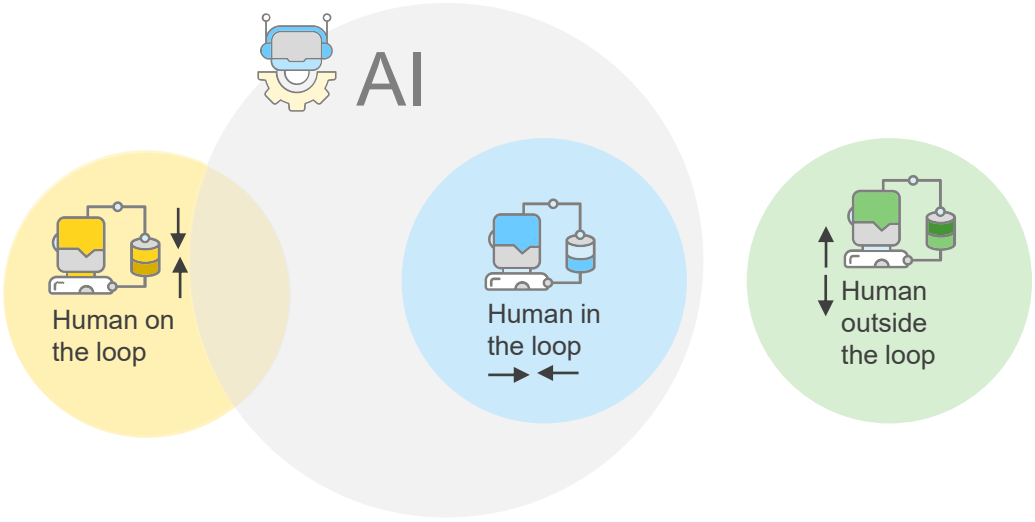

Exhibit 4 underscores AI's transformative potential in RCPG to assist leaders in identifying high-value use cases, demonstrating key current deployments alongside tomorrow's advances.

Exhibit 4: AI across RCPG use cases

Source: Everest Group (2025)

	Use case	Current AI capabilities	Future outlook
Procurement 	AI-based payment terms optimization	Calculates the total cost impact of different payment terms, for instance net 30 vs. net 60	Near real-time renegotiations if interest rates/credit conditions shift
	Lead-to-cash optimization	Automates quote generation, validates order terms, and streamlines contract approvals across customer segments	Autonomous agents will negotiate, execute, and reconcile end-to-end deals in real time based on buyer behavior and risk thresholds
Supply chain management 	Dynamic assortment optimization	Predicts consumer preferences across segments, suggesting which SKUs to expand or reduce in certain stores/regions, improving shelf space usage	Store-specific (or e-commerce site-specific) AI-driven assortments that continuously adapt to local trends and demand patterns
	AI-driven early warning on supply disruptions	Aggregates supplier performance data to anticipate inbound delivery disruptions and trigger contingency plans	Standardized supplier risk analytics widely used across industries, with automatic re-sourcing or buffer stock strategies triggered by red flags
Merchandising 	Dynamic pricing	Scrapes competitor websites or public data to gauge market price levels, auto-recommending price updates (match, undercut, or premium) per SKU	Autonomous, real-time pricing agents will self-optimize using live data across inventory, competitor, and macro-trend feeds (weather, social sentiment) and adjust prices continuously
	Product Lifecycle Management (PLM)	Tracks SKUs across launch, growth, maturity, decline, prompting end of life markdowns	Agentic lifecycle orchestration that triggers price adjustments, re-promotions, or kill decisions for underperforming SKUs
Development and manufacturing 	New concept design	Generative AI for idea generation (formulations, packaging concepts)	Agentic AI co-designing new products with real-time feedback loops from pilot users, social media
	Production planning and scheduling	Demand-driven scheduling to align daily or weekly production volume with forecasts	Agentic, near real-time scheduling that continuously juggles production lines, especially under sudden disruptions like raw material delay, machine breakdown, etc.
	Quality control	Correlates sensor data (temp, humidity) with product quality to warn of out-of-specification conditions	Agentic process control that self-adjusts mixing or cooking parameters mid-run to keep quality consistent
Customer engagement 	Voice-of-the-customer war room	An AI-driven dashboard integrates real-time feedback across channels, giving cross-functional teams simultaneous visibility	Unified brand “nerve center” for quick decisions, turning continuous feedback into a competitive advantage

Exhibit 4: AI across RCPG use cases (continued)

Source: Everest Group (2025)

	Use case	Current AI capabilities	Future outlook
Customer engagement 	Social commerce optimization	Integrates “shoppable posts” with AI suggesting best times, product combos, or price points for maximum conversion on social platforms	Social channels evolve into major e-commerce drivers with advanced AI bridging social listening, dynamic pricing, and real-time inventory
Store operations 	AI-driven workforce management	Demand-driven shift scheduling using historical foot traffic and promotions data	Agentic scheduling adjusting staff hours or reassignments in real time, e.g., if foot traffic spikes suddenly
	POS innovations	Computer vision self-checkout in select stores with limited scale	Widespread “just-walk-out” or frictionless checkout tech in mainstream grocery or C-stores
	AI-driven loss prevention	Theft detection using camera feeds for suspicious behavior or self-checkout mis-scans	Fully autonomous suspicious-activity detection with near-zero false positives, integrating multi-sensor data such as RFID and weight sensors
	Shopper movement heatmaps and shelf layout	Tracks how shoppers navigate aisles, revealing traffic flow and dwell-time to guide item placement for max visibility	Predictive, privacy-safe layout shifts, personalized by shopper segments and integrated with marketing and inventory systems for hyper-optimized in-store journeys
Sales and trade marketing 	Omnichannel demand sensing	Consolidates e-commerce, in-store PoS, and third-party marketplace data to predict total brand demand across all channels	Real-time models that adapt to new channels, products, and market shocks, personalizing stock, pricing, and promotions without manual retraining
	Dynamic promotions	Suggests optimal discount timing and levels based on historical sales data and competitor benchmarks	Agents autonomously tailor promotions in real time using product attributes, shopper behavior, and inventory context deploying differentiated offers across regions, channels, and segments
Digital commerce 	Fraud and abuse detection	Flags suspicious transactions, abnormal return patterns, or fraudulent promotions, protecting platform integrity	Autonomous agents will both flag anomalies and draft audit reports, propose adaptive rules, and orchestrate cross-system investigations
	Autonomous shopping	Chatbots and voice assistants automate Q&A, product recommendations, and checkout, streamlining user journeys	Highly conversational bots using large language models with advanced contextual understanding and multilingual voice support
	Dynamic product content and layout	Determines optimal digital “storefront” placements, product groupings, or promotional banners to increase conversions	Self-adapting website product displays that rearrange based on real-time shopper interactions and brand priorities

While Exhibit 4 covers use cases from an RCPG value chain perspective, enterprises are also using agentic AI to transform internal IT operations. AI agents are helping to triage support tickets, automate root-cause analysis, and prioritize incidents based on business impact. Additionally, they assist in provisioning software, managing access rights, surfacing relevant training content, and guiding employees through common troubleshooting workflows. These use cases reduce IT backlog, accelerate resolution times, and improve the employee experience, particularly in hybrid work environments. As IT functions mature, SoE-based agents will act as autonomous service desk coordinators across enterprise support ecosystems.

A lead-to-cash agent may detect buyers' intent, apply pricing logic, and generate contracts without human prompts. Through structured planning loops, SoE allow enterprises to shift from insight latency to execution velocity.

Agentic AI in action: early deployments reshaping RCPG

As the RCPG industry moves beyond descriptive analytics and prescriptive recommendations, a new paradigm is emerging – agentic AI. These systems go beyond assisting decisions – they act on them, adapt through feedback, and operate with increasing autonomy. Leading RCPG players are beginning to operationalize these self-directed systems, signaling a shift toward a more autonomous, intelligent enterprise. Exhibit 5 outlines key agentic AI use cases.

Exhibit 5: Agentic AI use cases

Source: Everest Group (2025)

	Agent in action	Future outlook
Autonomous shopping 	**Case in point:** Amazon Buy for Me is an autonomous shopping AI agent that sources non-Amazon listings, executes end-to-end purchases with secure credential handling, and integrates order tracking into the Amazon interface. **Deployment model:** It runs on AWS Lambda, proprietary LLM APIs and real-time feedback loops.	Agents will auto-replenish essentials and place orders based on customer usage patterns. They will autonomously suggest bulk-buy deals, manage subscriptions, and apply personalized discounts.
Lead-to-cash optimization 	**Case in point:** PepsiCo has deployed agentic AI to automate parts of its lead-to-cash cycle, including purchase order handling and data reconciliation, enabling faster order processing, reduced manual intervention, and more responsive sales operations. **Deployment model:** It uses UiPath's automation platform integrated with ERP systems, featuring agentic workflows and real-time exception handling.	Agents will autonomously generate dynamic quotes, negotiate pricing within approved thresholds, and initiate contract workflows in real time. They will orchestrate end-to-end invoicing, apply contextual payment terms, and trigger escalation workflows for at-risk receivables.
Dynamic pricing 	**Case in point:** Fashion e-retailer ASOS has deployed an AI-driven multi-agent pricing system that offers real-time, personalized price adjustments. When shoppers linger or hesitate at checkout, a negotiation-style chatbot is triggered, offering time-bound discounts or tailored incentives. **Deployment model:** The system operates via a multi-agent architecture combining behavioral analytics, real-time decisioning, and generative AI interfaces.	Dynamic price adjustments tailored to customer segments, loyalty tiers, or even individual purchasing behavior, powered by predictive propensity models. Standardized agent APIs will enable seamless integration across ERPs, CRMs, and e-commerce systems for always-on pricing.

Beyond the use cases described above, agentic and semi-agentic AI solutions are quietly transforming retail operations. No longer confined to pilots, these semi-autonomous agents now underpin workflows, boosting efficiency, enhancing agility, and enabling smarter, faster decisions at every value chain stage.

Walmart's My Assistant, for instance, coordinates around 45 AI agents – from shelf-restock alerts to demand forecasting – freeing associates to focus on customer engagement. Amazon's Interests feature monitors the catalog for user prompts, alerting shoppers to restocks or deals, and future versions may auto-purchase items based on past customer behavior. Meanwhile, Nova Act – Amazon's internal assistant capable of booking travel, making purchases, and managing calendars – could be adapted for RCPG workflows, such as reordering raw materials, adjusting pack sizes, or rescheduling supply chains in response to seasonal or demographic shifts in the future.

Yet with greater autonomy comes greater risk. Persistent monitoring features such as "Interests" may feel invasive without transparent controls. In-store robots like Kroger's Tally and Barney could soon update digital shelf labels without human oversight. But what happens when a pricing error goes unnoticed during a Saturday rush? Similarly, fully autonomous production lines offer the promise of "lights-out" efficiency, but a single misstep – whether a faulty sensor or a flawed recipe update – can lead to significant waste or downtime. As RCPG leaders scale these systems, building safeguards will be essential to maintaining accuracy, fairness, and consumer trust.

Speed, scale, and the need for control: navigating the challenges of scaling AI

While agentic AI holds transformative potential, it also introduces complex challenges. On the one hand, organizations stand to gain significantly from AI initiatives through greater efficiency, speed, and autonomy. On the other, these advances come with risks that cannot be ignored, including decisions made without sufficient oversight or data quality issues that trigger flawed actions at scale. It becomes imperative for enterprises to recognize these challenges early, build awareness across teams, and take precautionary steps to mitigate them.

Exhibit 6 outlines key hurdles facing enterprise AI adoption and the considerations needed to address them effectively.

Exhibit 6: Challenges of AI adoption

Source: Everest Group (2025)

	How it plays out	Recommendations
Data accuracy- and bias-related concerns 	AI-driven forecasts misprice promotions or reorder volumes AI proves ineffective, as constant manual overrides are required	Implement end-to-end data-quality pipelines and validation checks Deploy bias-detection tools and fairness audits with alerts for AI actions that go beyond defined thresholds Establish AI ethics and governance forums
Insufficient proprietary data for custom models 	Generic, “one-size-fits-all” bundles and pricing Poor personalization of agentic recommendations	Fine-tune SLMs (e.g. Blue Yonder Luminate’s Retail AI) on your business data using training platforms like Amazon Bedrock or H2O.ai LLM Studio Augment with synthetic data and transfer-learning Form industry consortia or partnerships for data-sharing
Inadequate generative and agentic AI expertise 	Agentic pilots stall in the PoC stage Heavy reliance on expensive external consultants High time to value	Launch an internal AI CoE with “agent architect” roles Upskill teams in prompt engineering, MLOps, and reinforcement learning Adopt low-code/no-code agentic platforms
Weak financial justification and RoI visibility 	Budget cuts on AI/agentic initiatives Difficulty quantifying incremental margin uplift from autonomous agents	Begin with narrow, high-impact agentic use cases. Use hyperscaler credits to run pilots at zero upfront cost Track pilot metrics such as cycle time and margin lift to build a tangible business case Model cost-benefit scenarios early along with partners
Intellectual property leakage and unauthorized disclosure 	Proprietary algorithms or training data inadvertently exposed to third parties Competitors reverse-engineer IP from shared outputs	Apply anonymization, encryption, and safeguards to prevent identifying individuals (at least for use cases that don’t demand individual-level predictions) Govern via strict access controls and audit logs Leverage federated learning to preserve data sovereignty
Change management and organizational readiness 	Teams resist AI integration due to unclear role impacts Execution bottlenecks caused by the lack of AI process ownership Misalignment between business and tech functions	Define new roles (e.g., agent supervisors, orchestration owners) early in the transformation Invest in training on AI collaboration models (in the loop, on the loop) Communicate a clear “human + AI” narrative to foster trust and alignment

Scaling agentic AI requires more than advanced technology. It demands readiness across data, talent, and governance. As organizations move from pilots to enterprise-wide adoption, they will need to proactively address these challenges. Evidence shows that RoI accelerates when AI augments human expertise rather than replaces it, enabling hybrid models that balance autonomy with oversight. Organizations that lay the right foundations now will be best positioned to lead in an autonomous, AI-driven future.

Walmart's My Assistant coordinates around 45 AI agents – from shelf-restock alerts to demand forecasting – freeing associates to focus on customer engagement.

Charting the AI maturity journey in RCPG

As RCPG companies embed AI deeper into their operations, they progress through four maturity stages: Basic, Typical, Advanced, and AI Leaders. This section lays out a clear, four-stage AI maturity framework tailored for RCPG organizations.

Exhibit 7 defines each stage, with targeted short- and long-term priorities alongside build-vs.-buy guidance to help enterprises assess where they are today and plan their next steps toward enterprise-wide, advanced AI capabilities.

Exhibit 7: Enterprise AI maturity framework

Source: Everest Group (2025)

		Basic	Typical	Advanced	AI Leaders
		Early-stage with minimal or ad hoc AI use Focused on isolated use cases such as marketing, customer service, and automation Teams depend on spreadsheets and siloed data, lacking integrated systems	Multiple PoCs across functions without a unified AI strategy Early investment in central data platforms, but silos persist Expanding AI team and metrics, but lacking standardization and full integration	Formalized AI strategy with budget, cross-functional teams, and RoI goals AI integrated across major business units with more centralized data and systematic RoI tracking Early experiments with agentic AI are in progress but not yet scaled	AI embedded across all enterprise functions in strategic decision-making Centralized and federated teams manage advanced infrastructure, including cloud, edge, IoT, and MLOps Agentic AI systems are in pilot and support use cases ranging from R&D to customer experience
Short-term focus	**Use cases**	Pilot low-cost, low-risk AI tools, such as chatbots and simple fore-casting to show quick wins	Scale successful pilots in areas such as inventory management and marketing across regions or product lines	Enable advanced personalization using generative AI for content and recommendations	Lead with agentic AI to develop new models such as autonomous supply chains and checkout-free stores
	Data	Start cleaning, centralizing, and labeling key data, such as product, customer, and store	Deploy a data lake or warehouse to unify and standardize data	Scale AI from siloed tools to integrated, cross-functional ecosystems	Co-create AI standards with partners, start-ups, and industry groups
	Investments	Hire or upskill data talent	Establish an AI CoE to define best practices and governance	Build strong MLOps pipelines for monitoring, retraining, and compliance	Drive innovation through M&As and strategic alliances
Long-term focus	**AI strategy**	Build a structured AI roadmap aligned with key business goals	Expand pilots to multi-channel and multi-brand programs and start testing agentic AI across themes such as autonomous robots and supply chain automation	Align AI with enterprise-wide strategies such as robotics, D2C, and virtual advisors; scale adaptive and agentic AI for real-time operational decisions	Experiment with self-evolving agents and autonomously generate and retire sub-agents to create new strategies
	Governance	Establish basic governance frameworks for privacy, ethics, and security	Define standards for MLOps, compliance, and ethical AI	Enhance governance to manage bias, privacy, and ethical risks	Adhere to global AI ethics while exploring models that enable autonomous compliance and transparent audit trails
Build vs. buy considerations		Depend on off-the-shelf tools such as CRM and vision or speech APIs Use basic connectors or rely on third-party consultants for deployment	Use a hybrid model combining off-the-shelf tools with custom integration layers Collaborate with vendors focused on RCPG-specific capabilities Prioritize data integration through both internal builds and external solutions	Build core IP in-house in areas such as forecasting and manufacturing control Acquire or partner for specialized modules such as personalization and automation Use cloud platforms to scale efficiently while retaining control	Develop proprietary AI solutions in-house through focused R&D teams Co-develop breakthroughs with select tech partners Use scalable, cloud-agnostic architecture for global agility

With the above framework, RCPG leaders can assess their current AI maturity, prioritize the most impactful initiatives, and make informed build-vs.-buy decisions.

Conclusion

AI stands at an inflection point in the RCPG sector. With the emergence of agentic capabilities and SoE integrations, enterprises have an opportunity to move beyond insight generation toward autonomous action. These systems promise to drive responsiveness, optimize operations, and unlock new efficiencies across the value chain. Yet, as the use cases and challenges show, the road to adoption is complex and requires deliberate planning.

To capitalize on this shift, organizations must act with focus and foresight. This means assessing AI readiness, prioritizing practical, high-impact use cases, and investing in foundational capabilities, particularly around data quality, system integration, and ethical governance. Equally important is building the right partnerships to accelerate learning and reduce execution risk. As systems of execution move to the forefront, the most successful enterprises will be those that treat autonomy not as a distant goal but as a strategic necessity – integrating it into both frontline operations and enterprise-wide decision-making.

Everest Group®
With you on the journey

Everest Group is a leading research firm helping business leaders make confident decisions. We guide clients through today's market challenges and strengthen their strategies by applying contextualized problem-solving to their unique situations. This drives maximized operational and financial performance and transformative experiences. Our deep expertise and tenacious research focused on technology, business processes, and engineering through the lenses of talent, sustainability, and sourcing delivers precise and action-oriented guidance. Find further details and in-depth content at www.everestgrp.com.

This study was funded, in part, by Ciklum

For more information about Everest Group, please contact:

+1-214-451-3000

info@everestgrp.com

For more information about this topic please contact the author(s):

Manu Aggarwal, Partner
manu.aggarwal@everestgrp.com

Abhishek AK, Vice President
abhishek.ak@everestgrp.com

Abhilasha Sharma, Practice Director
abhilasha.sharma@everestgrp.com

Aakash Verma, Senior Analyst
a.verma@everestgrp.com

Notice and Disclaimers