RESEARCH PAPER

AI Reality Check

What Retailers Need to Get Right Before Scaling AI

Connection®
we solve IT®

DELL Technologies
TITANIUM BLACK PARTNER

INDEPENDENT RESEARCH PARTNER

IHL GROUP

CONTENTS

Executive Summary

1: The Reality Gap

2: The Network and Edge Foundation

3: Computer Vision

4: Associate Productivity and Augmentation

5: Inventory Management and RFID

6: The Scaling Roadmap

Executive Summary

Retail has settled the AI question. While generative AI is just getting its sea legs, machine learning has fully reached general acceptance; and the impact potential is massive. IHL's AI Forecast Model projects up to $9.2 trillion in financial impact to the retail industry between 2022 and 2029 if retailers maximize potential, and AI-ready retailers can expect a 5% to 15% cumulative improvement to the P&L. Every retailer of consequence now has several budget lines, a mandate, and has moved well beyond pilots to complete installs in marketing, supply chain, and the store level.

Retail has not settled the sequence question. That is what this research paper is about.

Start with the number that frames everything else. Global inventory distortion runs $1.7 trillion a year in 2026, equal to 6% of global retail sales. That is $4.7 billion every day in product customers wanted to buy and could not find, or product nobody wanted sitting in a back room. While that figure has fallen from 10% of sales in 2021—five straight years of improvement, so the industry is getting better at this—it is still the largest unmanaged cost pool in retail due to ongoing continued macroeconomic forces (like port shutdowns, wars, and tariffs) that expose more holes in execution.

AI is already closing that gap for the retailers who have deployed it. Retailers running AI or machine learning in the functions most affected by distortion post 2.3x the sales growth and 2.5x the profit growth of retailers who have not. Yet only 25% overall have moved from pilots to full rollouts at scale.

BY THE NUMBERS

40%
of retailers have fixed cameras installed.

9%
run a current computer vision deployment.

Source: 2026 IHL CSA Retail Transformation Study

The other 75% have the ambition. What they lack is the foundation underneath it. Only 14% of retailers use edge computing at the store. Only 9% use computer vision, despite 40% having IP enabled fixed cameras mounted in their ceilings.

That distance between installed and intelligent is the honest state of retail AI in 2026, and the store is where it shows up first. Eighty-nine percent of commerce touches a store somewhere in the fulfillment path, which makes the store the endpoint where an AI ambition either becomes an operating result or stays a slide.

Three Questions This Paper Answers

What is actually standing between retailers and AI at scale? Spending is spread almost evenly across six functional areas that mostly do not talk to each other, sitting on networks and edge infrastructure sized for transactions rather than inference.

Which store-level foundations have to be right first? Network and edge, then use of RFID and computer vision on the cameras already installed, then associate-facing AI, with inventory accuracy running underneath all four.

What are early movers seeing in return? Retailers using 5G post 9% sales growth against 7% for non-5G peers. Retailers already deploying edge computing are 687% more likely to have a computer vision deployment running. And the results are showing. Retailers leveraging computer vision expect 26% higher profit growth than non-deployers for 2026.

None of that requires a moonshot. It requires doing four things in the right order.

"The winning formula is not a secret. It is a sequence. And most retailers are attempting to execute it out of order."

Source: The 2026 IHL Inventory Distortion Study

1: The Reality Gap

Picture the Monday steering committee. The AI pilot your team funded in March is on the agenda and the deck says the model is performing to spec. Accuracy is up and to the right. The data science team is pleased. Then a regional VP asks the only question that matters: what changed in my stores?

Nothing changed in the stores. The model was never the constraint.

The pilot ran in a clean cloud environment against historical data. The store runs on a network sized for card authorizations, a backroom with no local compute, and an associate holding a device that cannot reach the model without a strong signal and three seconds of patience. Every month that gap stays open, budget burns, the committee loses confidence, and the competitor who fixed the foundation first widens the lead. This is the reality gap: a foundation problem wearing an AI budget.

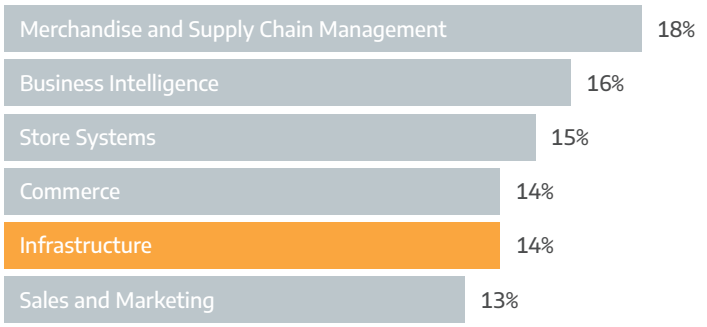
The money is spread evenly across six functions that mostly do not talk to each other.

Retail AI spending in 2026 divides almost evenly. Merchandise and supply chain management takes 18%, business intelligence 16%, store systems 15%, commerce 14%, infrastructure 14%, and sales and marketing 13%.

No single function receives enough concentrated investment to reach production scale, and each of the six sits on its own data. The infrastructure line is the problem. The layer every other line depends on receives roughly one dollar in seven, the same share as commerce, which is why pilots that perform in the lab stall on the way to the shelf.

AI Spend Allocation

Infrastructure receives the same share as commerce, and everything else depends on it.



Source: IHL AI Readiness Index

14%
of AI spend goes to infrastructure, the foundation every other category runs on

Categories shown sum to 90%; the remaining 10% sits in other and unallocated categories.

Readiness is measurable. Most retailers have never measured theirs.

Four questions get an executive into a defensible self-assessment before any consultant walks in the door:

- What percentage of the software stack is cloud- or browser-based today?

- What percentage of the BI and data infrastructure has been modernized since 2020?
- Is edge computing at the store current, or is the last meaningful refresh more than five years old?
- Can the store network carry inference traffic alongside transactions, IoT traffic, video and guest Wi-FI?

A retailer who cannot answer those four questions cannot forecast an AI outcome.

Any projected return is a guess dressed as a business case. Where AI has been deployed in the right order, it works.

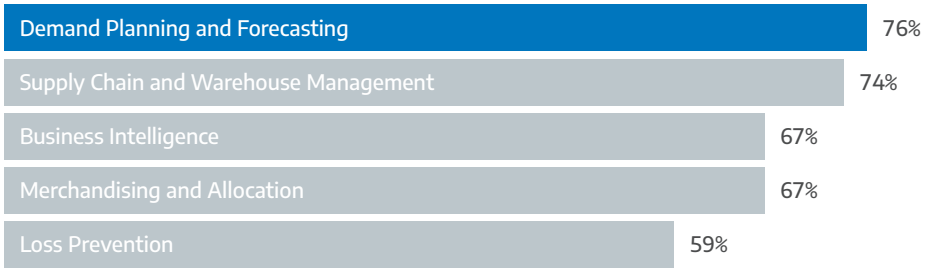
Among retailers who have deployed AI or machine learning in the functions most directly affected by inventory distortion, positive-result rates run high: demand planning and forecasting 76%, supply chain and warehouse management 74%, business intelligence 67%, merchandising and allocation 67%, and loss prevention 59%.

Three out of four demand planning deployments return a positive result. Yet fewer than 25% of retailers have made the deployment at all.

Those results also compound. IHL research has shown retailers deploying AI in distortion-affected functions post 2.3x the sales growth and 2.5x the profit growth of non-deployers.

<25%:
the share of retailers who have deployed AI or ML in distortion-affected functions at all

Positive-result Rates by Function



Source: 2026 IHL Inventory Distortion Study

Three companies show what sequence discipline returns and what skipping it costs.

Walmart is training 2.1 million employees on AI, the largest workforce AI readiness investment in retail history. The company treated human readiness as infrastructure and funded it on that scale.

Tesco tested its Tomoro AI associate tool with 280,000 colleagues before any customer-facing release. The order matters. Internal proof first, customer exposure second.

Lands' End shows the other side of the ledger. A warehouse management system implementation created a one-week backlog that cut first-quarter revenue 9% year over year. The technology eventually worked. The readiness gap still took a quarter of revenue with it on the way through, but it is now seeing multiples of return after the initial misstep.

Each outcome traces to the same layer. Walmart funded people. Tesco funded proof. Lands' End discovered what an unready operation does to a working system. In all three cases the deciding variable sat below the application, in the network, the edge, and the store operation that has to carry the load every day of the year.

2: The Network and Edge Foundation

The Infrastructure Layer AI Cannot Scale Without

Every technology in this paper ends up in the same place. Computer vision, RFID, associate AI, and inventory sensing all terminate at a store—and 89% of commerce touches a store somewhere in the fulfillment path. Whether it is click and collect, local delivery, or walk in shoppers, the store is running on infrastructure only 14% of retailers have brought current.

The sharpest dividing line in retail technology is one most boards have never discussed.

Edge computing currency is the single strongest technology differentiator in the 2026 IHL CSA Retail Transformation Study. Only 14% of retailers use edge computing today.

Currently this explains more variance in AI outcomes than any application layer decision a retailer will make this year. A demand forecasting model and a shelf-monitoring model fail the same way when the store underneath them cannot carry the load.

Four tiers, and most retailers budget for one.

IHL's retail AI processing framework separates where inference actually runs. Each tier does work the others cannot.

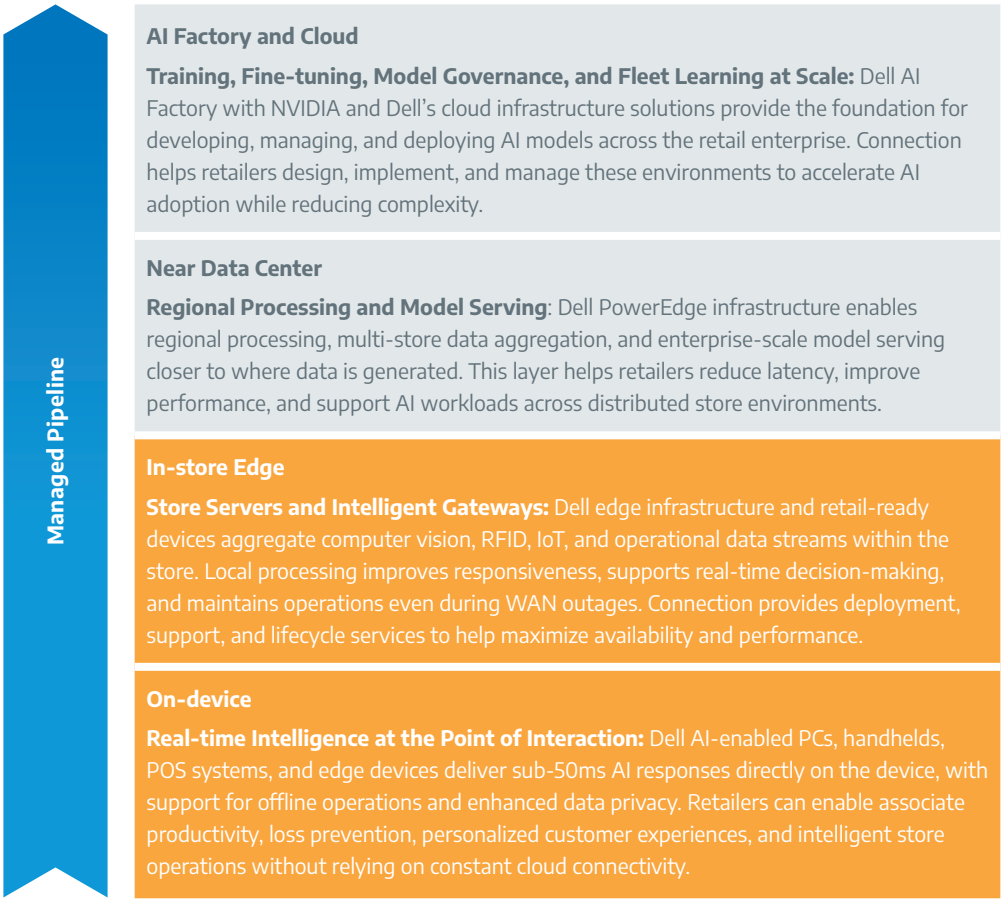




Tier	What Runs There and Why It Matters
On-device	Handhelds, POS terminals, and cameras with NPUs. Sub-50ms response, offline-capable, zero marginal cost per inference, and PII never leaves the device.
In-store Edge	Store servers and gateways aggregating computer vision, RFID, and IoT streams. Survives WAN outages, which is exactly when the store needs it most.
Near Data Center	Regional processing, multi-store aggregation, and fleet model serving.
AI Factory / Cloud	Training, fine-tuning, and fleet learning. Most retailers should consume this capacity through partners rather than build it.

Only 14%
of retailers are current on edge computing.

The Four-tier Retail AI Processing Framework
Powered by Dell Technologies and Connection Services



Source: IHL retail AI processing framework; Source: 2026 IHL How Leading Retailers Outperform

Most AI budgets today fund the fourth tier and quietly assume the first three will be there, but that is simply not the case in most retailers.

The economics decide where the inference lives.

Frontier cloud models bill per query, per associate, per shift, forever. That cost curve rises with adoption, so the more successful a deployment becomes, the more expensive it becomes to keep.

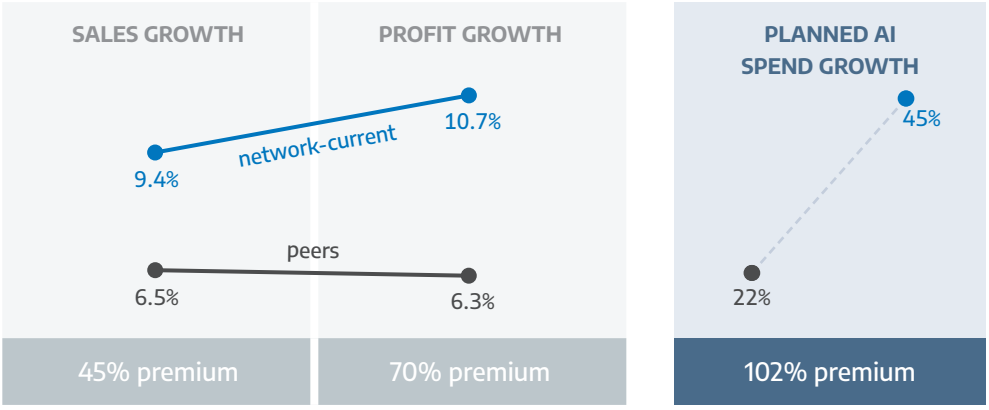
Custom-trained small models running on-device or at the in-store edge ride the hardware refresh cycle retailers already budget for, at zero marginal cost per query. Cost per use falls as adoption climbs. That single difference often decides whether an associate AI tool survives its second budget year.

The network dividend shows up in the P&L, not the IT budget.

Retailers with updated networks post 9% sales growth against 7% those whose networks are holding them back, a 45% premium. Profit growth runs 11% against 6%, a 70% premium.

The same group is pulling away on intent. Retailers with updated networks plan 45% growth in AI spending against 22% for their peers, a 102% premium in investment velocity. The retailers best equipped to run AI are also the ones accelerating hardest into it.

Network-ready Retailers Grow Faster and Accelerate Harder



Source: 2026 IHL How Leading Retailers Outperform

Retailers using edge compute are winning categories they did not buy edge to win.

Retailers currently using edge computing show 615% higher 5G deployment and are 845% more likely to run a microservices architecture. They are 62% more likely to have local delivery fully optimized, 63% against 29%.

Retailers who prioritize WAN investment plan self-checkout at 5x the rate of non-WAN peers and are 250% more likely to plan edge computing investment in the next 24 months.

In other words, the network investment came first and the edge and self-checkout plans followed it. Network capability is the leading indicator of self-service and AI capacity, not a purchase order that trails the AI decision.

Endpoint Density



Source: 2026 IHL How Leading Retailers Outperform
1. Gratis deployed roughly two million Vusion electronic shelf labels across 300 stores (~6,667 total per store)
2. Superdrug deployed 1.7 million Hanshow units across the UK and Ireland.

What the foundation looks like in the market today.

Leaders with updated networks are also 615% more likely to have deployed electronic shelf labels and 102% more aggressive on AI spending.

Electronic shelf labels make the density concrete. Gratis deployed roughly two million Vusion labels across 300 stores, about 6,700 endpoints inside a single store. Superdrug deployed 1.7 million Hanshow units across the UK and Ireland. Every one of those units has to be provisioned, powered, monitored, and kept online. A shelf label is a small thing. Two million of them is an infrastructure program that relies on up-to-date edge computing.

The winning architecture is a managed pipeline running from device to in-store edge server to network to AI factory. No single purchase delivers it, and no single vendor category covers it. The retailers pulling ahead treat those four tiers as one system with one owner, and they buy it that way.

3: Computer Vision

From Loss Prevention to Store Intelligence, Retail's Fastest-growing AI Category

The largest AI deployment opportunity in retail is already bolted to the ceiling. 40% of retailers have fixed cameras installed. Only 9% run a current computer vision deployment in 2026 on them. Most of the installed base is recording footage that nothing is reading outside of loss prevention recordings. Retailers expect an 8,143% increase in computer vision and image recognition adoption from the start of 2025 through 2027.

Retailers current on edge computing are 687% more likely to have a current computer vision deployment. Our earlier discussion of the infrastructure layer is the prerequisite for this section. Video is the heaviest workload a store network carries, and a retailer who has not solved the edge cannot run vision at store scale no matter how good the model is.

The returns here are the least ambiguous in this paper.

Sales winners (those who grew 2025 sales 10% or more) deploy computer vision at 33x the rate of their peers. Retailers with up-to-date deployments expect 26% higher profit growth and IHL research shows first-year ROI often exceeds 300%, with payback under six months through facial-recognition-based loss prevention applications.



Shelf Availability Outcomes

North American Drugstore Chain	4.5 points of on-shelf availability gained	more than $55M annual profit, roughly $12.2M per point
LATAM Grocery Chain	3.5 points gained	more than $28M annual profit
European Grocery Chain	5.0 points gained	more than $4M annual profit

Source: Shelf-Intelligence – The Tech Imperative for Retail

The use case has outgrown loss prevention.

Computer vision now runs shelf intelligence and store operations. Retailers report 15% to 30% productivity gains from vision deployed in distribution centers and store backrooms, an expanding application area beyond the original security use case.

Shelf availability is where it shows up in revenue. A North American drugstore chain improved on-shelf availability 4.5 points and captured more than $55M in annual profit, roughly $12M per point of availability. A LATAM grocery chain improved 3.5 points for more than $28M. A European grocery chain improved 5 points for more than $4M. Same technology, three different scales, all positive.

Where the inference runs decides what it costs.

New solutions now combine computer vision, RFID, and barcode scanning locally at self-checkout with no cloud dependency. They target a specific number: self-checkout shrink runs 3.5% to 4% against 1% at cashier lanes. That is the on-device tier doing production work.

Two more signals point the same direction. Instacart acquired computer vision firm Arpalus and is feeding that intelligence into its Caper Carts and Store View products across a 600,000-shopper network. Costco is combining RFID and computer vision in a warehouse club format where a tagged pallet sitting on the steel above an empty floor position is invisible revenue only a camera catches.

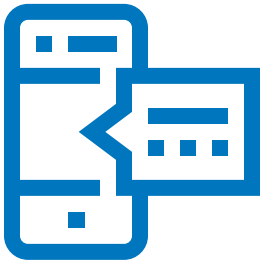
The sequencing implication means a retailer with cameras already mounted and a current edge deployment can put computer vision into production without opening a new capital cycle for either one. A retailer without the edge is buying two projects and calling it one, which is how a vision pilot ends up eighteen months behind its business case.

The cameras are installed. And since the economics are proven, the gate is the edge.

4: Associate Productivity and Augmentation

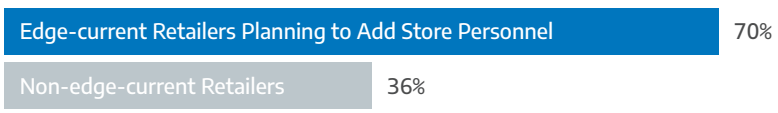
AI Is Best as a Workforce Multiplier, not a Workforce Replacement

The most persistent assumption about retail AI is that it arrives to reduce headcount. The retail deployment data points the other way, and it points from an unexpected direction. Sales winners are 115% more likely to be adding people to the process rather than removing them.



The Hiring Split

The retailers building the infrastructure are the retailers hiring:



Source: 2026 IHL How Leading Retailers Outperform

The retailers building the infrastructure are the retailers also hiring.

70% of edge-current retailers plan to add store personnel, against 36% of retailers who are not using edge computing. The infrastructure story and the workforce story are the same story told from two sides. Retailers who commit to the store commit to improving the total experience for the customer.

The pattern holds for a specific practical reason. A shelf-monitoring model that finds a gap produces a task list that keeps the retailer more in-stock and less likely to lose sales because customers cannot find the merchandise they came into purchase.

The largest recoverable block of time in retail is half an hour a shift.

Three-quarters of frontline workers spend 30 minutes per shift searching for information they do not have. That is the single largest recoverable block of associate time in the research.

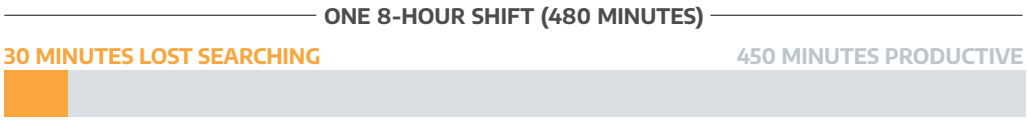
At five shifts a week, 30 minutes a shift is roughly 125 hours a year, about three full working weeks per full-time associate spent hunting for an answer that already exists somewhere inside the company. Multiply that by a store roster, then by a chain.

Unfortunately, too often the cause is structural. The answer exists but it sits in a system the associate on the floor cannot reach from where they are standing.



The 30-minute Gap

30 minutes per shift are lost searching for information, reported by three-quarters of frontline workers.



30 minutes × 5 shifts × 50 weeks = 7,500 minutes (or 125 hours)
≈ 3.1 working weeks per full-time associate.[1]

Source: IHL associate productivity research: Annual conversion is IHL-derived arithmetic
1. Five shifts per week, 50 weeks per year, 40-hour week.

What closes the gap runs in the store, not in the cloud.

On-device associate AI covers task management, product lookup, translation across 18-plus languages, and voice picking. Every one of those functions attacks the same 30 minutes.

The tier matters as much as the application. These functions run at the on-device layer or on the LAN edge discussed earlier, where response is sub-50ms, the tool still works when the WAN does not, and the marginal cost per query is zero. An associate assistant billed per query, per associate, per shift gets rationed by finance in year two. One running on device or on the edge does not.

The sequencing point for a retail leader is narrower than the headline. Start with the function that recovers time an associate already loses, because that return needs no behavior change and no customer exposure. Task management and product lookup clear that bar on day one.

5: Inventory Management and RFID

The Foundation Layer, Refreshed with the 2026 IHL Inventory Distortion Study

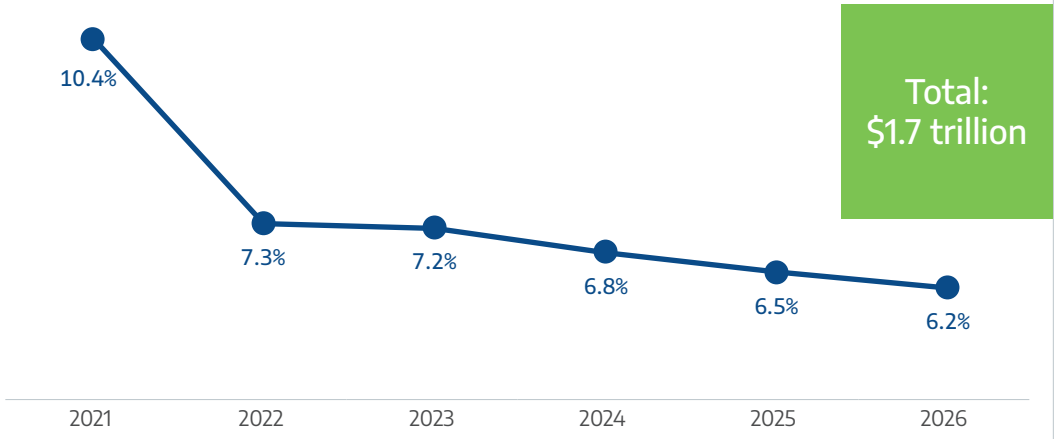
Retail is winning this one, but the battle is not over. Inventory distortion has fallen from 10% of global retail sales in 2021 to 6% in 2026—five consecutive years of improvement. That is the clearest sustained operational gain in the industry, and almost nobody is talking about it because supply disruptions from wars to recalls to tariffs dominate the narrative and prove to be a whack-a-mole for retailers of all size.

Global inventory distortion totals $1.7 trillion a year. For every $100 a retailer rings, $6.20 of demand sits in the wrong place. Out-of-stocks account for 66% of that total and overstock for 34%, which puts roughly two-thirds of the problem in demand the retailer could have served and did not.

The improvement was earned against a headwind, which is what makes it credible. The 2026 study identifies $298 billion in annual losses from supply chain coordination failures, now treated as a permanent operating baseline rather than an exceptional disruption. Retail cut distortion by more than four points of sales while absorbing that as a standing cost.



Five Years of Improvement



Source: 2026 IHL Inventory Distortion Study; IHL065 trendline exhibit

Grocery posted the strongest improvement of any segment IHL measures.

FDCM (food, drug, convenience, and mass merchants) is the best-performing super segment at 5% of segment sales, improved 34% since 2020. Inside it, grocery improved 46%, the strongest result of any segment IHL tracks. Specialty soft goods sit at the other end at 12% with low improvement.

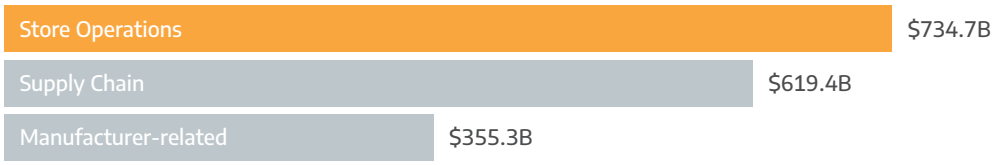
Grocery replenishes daily against short shelf life, which forces the accuracy discipline the improvement data reflects. Segments with longer product cycles have carried less pressure to build it over the last several years.

Store operations carries $734.7B of the total. Supply chain carries $619.4B. Manufacturer-related causes account for $355.3B. Personnel-related distortion alone totals $248B globally, roughly the annual revenue of Costco.

Of the full $1.7 trillion, $768B is directly attributable to empty-shelf moments customers experience in real time. That is the exposure most vulnerable to a shopper switching to same-day delivery rather than waiting for an associate to find merchandise that is "supposed to be in-stock."

Personnel-related distortion totals **$248B globally,** roughly the annual revenue of Costco.

The Cause of Distortion



Source: 2026 IHL Inventory Distortion Study

The benefits land in the channels retailers care most about.

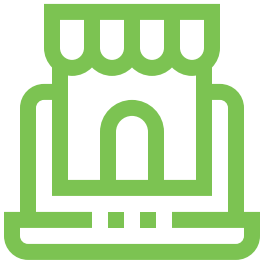
It's not just the impact on the walk-in customers. Retailers who have systemized and prioritized inventory accuracy generate 130% more BOPIS revenue and 43% fewer lost sales from substitutions. Scenario-based and digital-twin planning is associated with a reduction in inventory errors of 25% or more.

In other words, inventory accuracy is the gate on omnichannel revenue. A retailer cannot promise a pickup order it cannot verify, so every point of accuracy converts directly into inventory the store is willing to sell through a digital channel.

RFID is a proven accuracy layer inside a bigger stack.

Sales winners adopt RFID at 16x the rate of Laggards (flat or negative sales growth in 2025), the largest technology adoption-gap differential IHL has ever recorded (except for mobile devices for managers), and retailers project 291% growth in RFID deployment over the next two years.

Where it is deployed, RFID delivers up to 99% inventory accuracy, which is what makes ship-from-store work without phantom-stock errors. RFID earns its place as an accuracy layer inside a stack built on edge infrastructure and computer vision. As a retailer, you only care about which items moved since the last read. All of that processing must be done on the edge or it would overwhelm a WAN network.



The Accuracy Dividend

130%	43%	25%
more BOPIS revenue	fewer lost sales	or greater reduction in inventory errors from scenario-based and digital twin planning

These returns belong to the 22% of retailers who have systemized inventory accuracy, not to all retailers.

Source: 2026 IHL Inventory Distortion Study

The proof is already in production.

Boggi Milano runs RFID across all 105 stores at 99% inventory accuracy, against a 65% to 75% baseline typical of retailers without it, a swing of 24 to 34 points.

Beyond just RFID, Durham Brands moved throughput from 10,000 to 19,000 cases per FTE with an Infios WMS while holding 99% inventory accuracy, a 90% productivity gain on the same headcount leverage.

Daiso Industries replaced rule-based ordering with AI demand forecasting across all 4,625 domestic stores. VF Corporation is deploying an AI inventory engine across more than 1,500 stores for The North Face, Vans, and Timberland beginning in Q2 2026. Kroger is running autonomous inventory drones in live freezer zones down to -20 degrees Fahrenheit, the first grocery chain to do it.

Data accuracy is the layer everything else in this paper stands on. Computer vision needs it to know what a shelf gap means. Associate AI needs it to answer a customer standing in the aisle.

6: The Scaling Roadmap

An Action Framework for Getting the Sequence Right

Everything in this paper reduces to four moves in one order. The order is the strategy. Retailers who run these four in sequence compound each phase into the next. Retailers who start at the application layer pay for the same capability twice, once for a pilot and again when the pilot cannot scale.

The Best Practice Sequence

Phase	The Move	What Done Looks Like
Phase 1	Data Accuracy Audit	You can state inventory accuracy store by store, the share of the software stack that is cloud or browser-based, and the share of BI and data infrastructure modernized since 2020.
Phase 2	Network and Edge Foundation	The store network carries inference alongside transactions, video, and guest traffic. In-store edge compute is current and keeps running through a WAN outage.
Phase 3	Vision and Associate AI Pilots	Computer vision running on cameras is already installed. On-device associate AI is attacking the information gap in a pilot district, measured against recovered minutes per shift.
Phase 4	Scaled Inference on Custom Models	Small models trained for your fleet are running at the device and edge tiers, with AI factory capacity consumed through a partner rather than built.

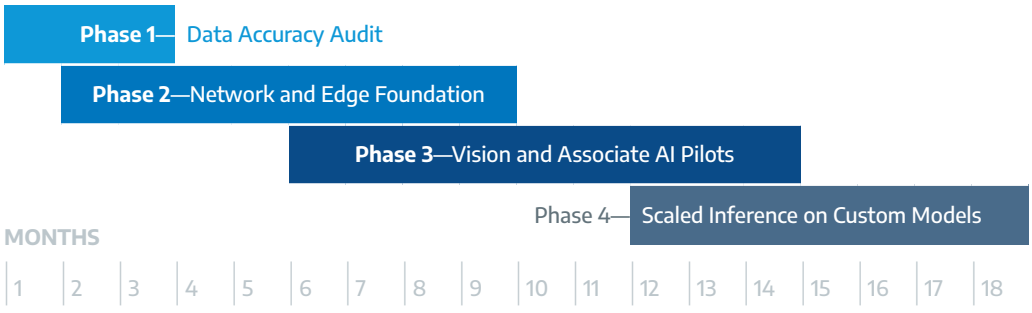
RFID and other point technologies slot into this sequence as accuracy layers inside phases one and three. None of them is a separate track, and none of them fixes a foundation problem underneath it.

Start with accuracy, because every later phase inherits it.

A vision model that finds a shelf gap is only as useful as the inventory record it reconciles against. An associate assistant answering whether a product is in the back room is only as good as the count behind the answer. Accuracy is the cheapest phase in this roadmap and the one most often skipped, because it produces no demo and no press release.

The hardest part of this sequence is political rather than technical. Phase one buys credibility for phases two through four, and it does it without a single thing to show a board.

Sub: The 18-month Sequence



Source: phase definitions from Section 6 of this paper timeline positions illustrative

Build the foundation before the application, because the data runs that direction.

Edge-current retailers are 687% more likely to have computer vision running. Retailers who prioritize WAN investment are 250% more likely to plan edge investment inside 24 months. The order is visible in the data itself: network, then edge, then application.

The foundation phase (data accuracy) carries its own return, which is what makes it fundable on its own terms. Then Phase 2 (network and edge). Network-ready retailers post a 45% premium on sales growth and a 70% premium on profit growth. Phase two does not need phase three to justify itself, and a CFO who sees that will fund the sequence instead of funding one more pilot.

Pilot where the hardware is already paid for.

Phase three (vision and associated AI pilots) has the shortest path to production of anything in this paper, because 40% of retailers are already carrying the capital cost of the cameras. A pilot that needs no new hardware clears procurement faster, proves the edge investment from phase two, and produces the operating result that funds phase four.

Build versus consume at the AI factory tier.

Phase four (scaled inference on custom models) is where retailers most often overreach. The AI factory tier covers training, fine-tuning, and fleet learning, and it is the tier with the highest capital cost and the shallowest retail-specific advantage. Three questions settle the decision:

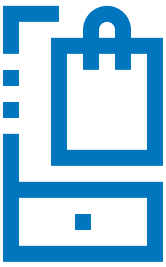
- Does the retailer hold proprietary data at fleet scale that no vendor model can replicate?
- Will a permanent machine learning engineering function be funded through a downturn, not just through a good year?
- Is model differentiation itself the competitive moat, or is the moat in how fast a working model reaches the shelf?

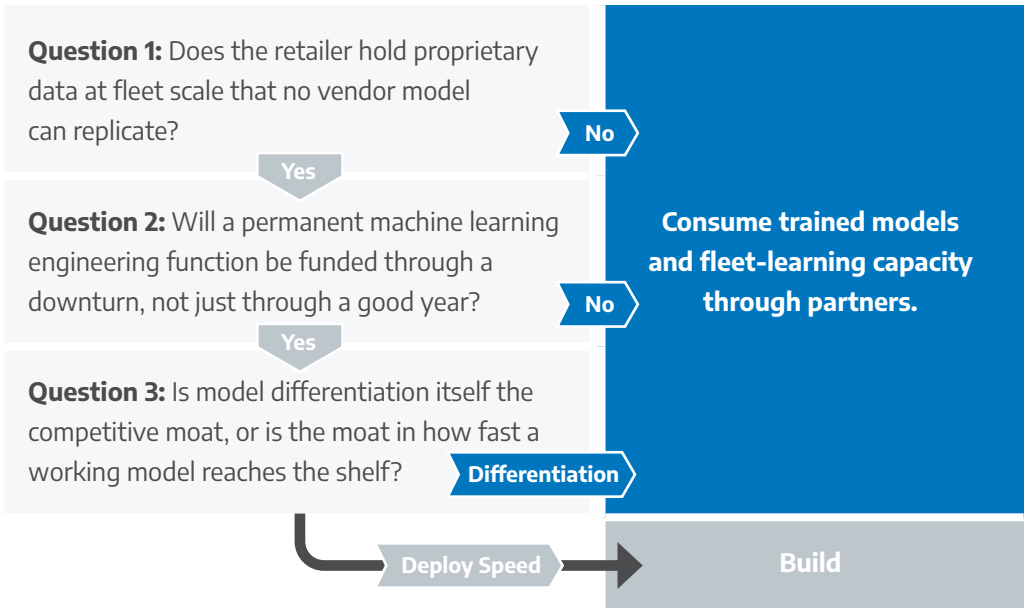
For most retailers, the honest answers are no, no, and the speed to the shelf. Consume trained models and fleet-learning capacity through partners and put the capital into the three tiers that actually touch the customer. The retailers who build their own AI factory should be able to name the proprietary dataset that justifies it before the first invoice.

Each phase of this roadmap has a procurement problem. Phase one needs assessment capability. Phase two needs network transformation and edge infrastructure across a device-to-data-center portfolio. Phase three needs managed device lifecycle so a pilot fleet does not become an operations burden the day it scales. Phase four needs integration services that connect device, edge, network, and AI factory into one managed pipeline.

The benefit a retailer is actually buying is a single accountable owner across four tiers, one refresh cycle instead of four, and one integration plan instead of a stack of vendor handoffs. Retailers who assemble this from four separate purchase decisions spend the eighteen months managing the seams rather than running the sequence. Most retailers cannot do this alone and need solid partners that can scale with you.

Build vs. Consume





The build path is correct for retailers who can name the proprietary fleet-scale dataset that justifies it before the first invoice.

Source: IHL analysis

What the sequence returns, phase by phase.

Run in order, this roadmap pays at every stage rather than at the end. Phase one recovers the accuracy that averages 130% more BOPIS revenue and 43% fewer lost sales. Phase two carries the network and edge benefits with a profit premium on its own. In IHL research, phase three delivers first-year computer vision ROI on cameras already installed. Phase four drives the recurring cost of inference toward zero by moving it onto hardware already inside the refresh budget.

Your next move.

The retailers pulling ahead are running this playbook in order. That is the whole advantage, and it is still available. Start with the phase-one test from Section 6. If you cannot state inventory accuracy store by store, the share of your software stack that is cloud or browser-based, and the share of your BI infrastructure modernized since 2020, you have just found your first ninety days.

Connection's retail team runs that assessment as a working session with your operations and IT leads, and the output is a sequenced plan rather than a proposal.

Book a readiness assessment conversation with one of Connection's retail experts.

Brian Gallagher, Retail Strategy and Business Development Director at Connection
brian.gallagher@connection.com

Connection
we solve IT

About IHL Group

IHL Group is a global research and advisory firm serving retailers and the technology companies that supply them. IHL's Sophia database tracks more than 500,000 technology installations across more than 6,000 retailers and 300-plus technology categories, giving IHL deployment-level visibility that survey-based research cannot replicate.

Methodology

Inventory figures throughout this paper are drawn from the 2026 IHL Inventory Distortion Study, fielded June 2026, built on IHL's WorldView IT Sizing Model, the Sophia Data Service, and IHL end-user surveys. Edge, network, and computer vision figures come from the 2026 IHL CSA Retail Transformation Study. AI impact and readiness projections come from the IHL AI Forecast Model and the IHL AI Readiness Index.

Connection and Dell Technologies

Connection, a Fortune 1000 company, in collaboration with Dell Technologies, helps retailers navigate the complexity of today's IT landscape by delivering innovative technology solutions that enhance growth, improve operational efficiency, and accelerate digital transformation. Together, we combine Connection's deep retail expertise, dedicated service model, and end-to-end technology capabilities with Dell Technologies' industry-leading infrastructure, devices, data, and AI solutions to help retailers modernize with confidence.

We are committed to helping retailers create smarter, more connected experiences that drive profitability, strengthen customer engagement, and support long-term growth. Through our joint solutions and services, retailers can advance smart store initiatives, optimize store and workforce productivity, leverage data-driven insights, and deliver seamless omnichannel experiences. Backed by retail industry experts, Dell Technologies specialists, in-house technical resources, and dedicated Account Managers, we help organizations address evolving business challenges, unlock innovation, and create the exceptional shopping experiences today's customers expect.

www.connection.com/retail

© 2026 PC Connection, Inc. All rights reserved. Connection® and we solve IT® are trademarks of PC Connection, Inc.
All other copyrights and trademarks remain the property of their respective owners. 3933072-0826