

TECH & AI CASE STUDY - GROCERY RETAIL - 2026

Edge Compute Across 200 Stores That Enabled Real-Time Shelf Operations

Based on a delivery engagement with Orchard Lane Markets

99%+

PRICE-LABEL ACCURACY

-22%

FRESH-FOOD WASTE

-75%

SHELF-AUDIT LABOR HOURS

-31%

OUT-OF-STOCKS, TOP SKUS

A fillable practice edition: the full story, 8 actions,
6 reflections and a 90-day plan - work it with your team.

Free - no email required
emersontechnology.com

MAKE THIS PRACTICE BOOK YOURS

Work it, do not just read it

This practice book turns the Orchard Lane Markets case study into a working session for your team. Read the full story first - what happened, what was built and what it changed. Then tick the 8 actions you can commit to, write short answers to the 6 reflections, and finish with a 90-day plan. Every field in this PDF is fillable: open it in Adobe Acrobat Reader, Chrome, Edge or macOS Preview, type directly into the boxes and save.

YOUR NAME

ROLE

COMPANY

TEAM / DEPARTMENT

DATE

EMAIL (OPTIONAL)

MY GOAL FOR THIS PRACTICE SESSION

- 01

The story - What happened at Orchard Lane Markets - before, during and after the project.
- 02

The results - The numbers that changed, plus industry benchmarks from public research.
- 03

The 8-action checklist - Tick boxes as your organization completes each step. Owner and date fields included.
- 04

The 6 reflections - Fillable questions that pressure-test your own roadmap.
- 05

My 90-day plan - Three fillable phase cards that turn this case study into your roadmap.

RUN IT AS A 90-MINUTE TEAM SESSION

25 min READ Read the story together. Ask everyone to note the one number that surprised them most.	40 min PRACTICE Work the 8-action checklist and the 6 reflections. Assign an owner and a date to every action you commit to.	25 min COMMIT Fill in the 90-day plan. One goal per phase, one number that tells you it worked. Book a review date before you leave the room.
--	--	---

THE STORY

What happened, and why it worked

01 The Company

Orchard Lane Markets is a 200-store regional grocery chain with roughly \$1.9B in annual revenue, concentrated in dense suburban trade areas where competition is a short drive away in every direction. Grocery is a business of small margins defended by execution: price integrity, shelf availability and fresh-department waste are the levers that decide whether a store makes money, and all three are decided at the shelf, hourly, not in a weekly report. Orchard Lane's merchandising and store-operations teams were experienced and hard-working, but the systems underneath them treated time in units of days, while the shelf moves in units of minutes.

The chain's scale made the problem structural rather than local. Two hundred stores means two hundred pricing surfaces, two hundred sets of shelves audited by hand, and two hundred opportunities per day for the label on the shelf to disagree with the price in the point-of-sale system. Every one of those disagreements is either a compliance risk during a state pricing inspection, a margin leak when the label is wrong in the customer's favor, or a trust leak when it is wrong the other way. Orchard Lane's leadership knew the cost was real; what they lacked was an architecture capable of acting at store speed.

02 Where Things Stood

Before Emerson Technology engaged, prices and inventory updated once a night in a batch job from the central data center. Between those nightly runs, the store was running on yesterday's decisions: shelf labels drifted from POS prices, out-of-stocks on high-velocity items often went unnoticed for a full day, and fresh departments marked down produce on guesswork rather than on sell-through data. Store teams spent an estimated 1,400 labor hours a week walking the aisles and auditing shelves by hand — skilled hours spent looking for problems that a camera could find, while shelf-condition problems still slipped through between passes.

The fragility of the batch architecture had already been demonstrated expensively. A pricing-file corruption pushed wrong markdowns to 60 stores, and because detection depended on the next nightly cycle and human reports, the error took two days to find and correct — an estimated \$300,000 in margin and goodwill, and a lasting credibility wound with the store managers who had to face confused customers. Worse, the same architecture meant that when the wide-area network link to a store was down, pricing operations at that store simply stopped. Real-time fixes were impossible on an architecture where every decision waited for the next nightly run.

The waste problem carried both financial and strategic weight. Fresh-food waste ran at a level the finance team had learned to treat as structural, and markdowns were applied by feel in produce and bakery. IHL Group research has estimated that out-of-stocks and overstocks cost retailers worldwide more than \$1 trillion annually, with grocery among the most exposed segments — Orchard Lane was not an outlier, just a well-run chain on an architecture that guaranteed yesterday's information.

03 The Challenge

The challenge was to move decision-making from the data center to the shelf without losing central governance — a tension that defeats most retail technology programs from one direction or the other. Fully centralized systems are too slow and too fragile: they fail when links fail, and they turn seconds-long problems into overnight ones. Fully autonomous stores create chaos: two hundred stores each inventing their own markdown logic is a compliance and margin disaster. The architecture had to let the center set policy while the store executes in real time, even — especially — when the connection between them is down.

The constraints were practical and unforgiving. Stores are harsh environments for hardware: power blips, heat, dust and five-year equipment lifecycles. Camera-based shelf auditing touches shopper privacy, so raw video could never leave the store. Two hundred stores also means two hundred deployment targets, and any update mechanism that involved a technician's visit would be obsolete before it finished rolling out. Finally, the program had to prove itself on a metric executives already tracked — the fastest route to killing a promising retail technology project is a business case nobody can verify.

04 Discovery and Diagnosis

Emerson Technology began with a three-week discovery sprint combining data-center archaeology with store-floor observation. The team traced the nightly batch pipeline end to end, mapped every dependency that made a store helpless when the WAN dropped, and interviewed store managers about where their hours actually went. The audit found the shelf-audit labor figure — roughly 1,400 hours a week chain-wide — was conservative, because it excluded the untracked time managers spent reconciling label disputes at the register. It also confirmed that the \$300,000 pricing incident had a specific root cause the new architecture would eliminate by design: a central file with no local validation and no rollback path. Gartner has reportedly predicted that by 2025, 75% of enterprise-generated data will be created and processed outside traditional data centers or centralized cloud, and Orchard Lane's own telemetry agreed — nearly all of its operationally urgent data originated on the sales floor.

The diagnosis ranked candidate workflows by pain and feasibility: shelf auditing, price-label synchronization and fresh markdowns surfaced first, in that order, and shelf auditing won the pilot slot because its value was measurable in existing reports. IDC's Worldwide Edge Spending Guide reportedly projects global edge computing spending of roughly \$230 billion in 2024 with double-digit growth through the decade, which the team cited not as a reason to act but as evidence that the required building blocks — hardened small-form compute, container orchestration at the edge, and store-deployable ML models — had matured past the early-adopter stage. The discovery concluded with a pilot design: ten stores, one workflow, one executive metric.

05 The Solution

Emerson Technology deployed an edge platform to every store: a compact Kubernetes cluster (K3s) running on on-prem servers that executes shelf-scan models, electronic-shelf-label synchronization and a local rules engine even when WAN links drop. The architecture has three layers.

The in-store platform

Each store runs a small, standardized cluster with GPU nodes for vision workloads, designed for harsh environments and remote management, with hardware refresh planned from day one rather than improvised at year five. The cluster hosts three services: a computer-vision shelf-scan model that flags gaps and mispriced labels in seconds, an electronic-shelf-label synchronization service that keeps shelf prices aligned with POS prices continuously, and a local rules engine that executes markdown and pricing policy without the data center. Camera-equipped shelf audits process images on-store and forward only aggregated events — never raw video — to the central platform, which simultaneously cut bandwidth costs, preserved shopper privacy and simplified compliance review to a policy statement rather than a legal analysis.

Central platform and policy control

On the central side, a Kafka event backbone streams aggregated store events into an AWS data platform built on Kinesis, S3 and Athena, giving merchandising a chain-wide, near-real-time view of shelf condition, price integrity and fresh sell-through for the first time. Markdown and pricing rules are authored centrally — corporate governance stays corporate — but they execute locally at the store where the data lives, under central version

control. The nightly batch was not merely accelerated; it was redefined into a continuous synchronization problem, and the pricing-file corruption class of incident was engineered away with signed deployments, local validation and instant rollback. Store managers get a tablet app that turns anomalies into prioritized fix lists, so a detection becomes a named task on today's shift rather than a chart on a wall.

Fleet operations and delivery

Because 200 stores means 200 deployment targets, the platform is treated as software, not as installations. Terraform and GitOps pipelines roll out model and configuration updates store by store in staged waves, with automated health checks and instant rollback — the same deployment discipline a software company applies to its own fleet. Every store is now a software-deployable unit: a pricing rule change, a new vision model or a markdown policy revision ships chain-wide in days, with per-store health visible centrally. When WAN links fail, stores keep pricing, auditing and marking down locally and synchronize automatically on reconnection, which converts network outages from operational events into non-events.

06 The First 90 Days

Weeks one and two stood up the pilot: ten stores selected for volume, layout variety and — deliberately — one store with a chronically unreliable WAN link, because designing for disconnected stores was a stated requirement rather than a hope. Weeks three through six deployed the first K3s clusters and the shelf-scan model, running the vision system in parallel with the human audit routine so detection quality could be compared honestly; the models found label mismatches and gaps the walking audits missed, and the walking audits found a modest rate of model false positives that fed two weeks of threshold tuning. Weeks seven through ten connected electronic shelf labels to the local sync service, ending the drift between shelf and register in the pilot stores, and wired the tablet fix-list app into the store teams' shift routines.

Weeks eleven and twelve added the fresh-department markdown rules executing locally under central policy, and ran a controlled comparison: pilot stores' produce waste against matched non-pilot stores, measured on the same reports finance already used. The result was unambiguous enough to fund the chain-wide rollout — fresh waste down double digits, audit hours visibly collapsing, and the WAN-outage store in the pilot operating a full pricing cycle with the link down. Store-manager adoption, the risk the team had feared most, resolved itself once the fix lists started arriving as tasks rather than as findings: the app told the night crew exactly which aisle and which label, and managers stopped auditing and started delegating.

The rollout to 200 stores proceeded in staged waves through the following quarters, each wave carrying its own health checks and rollback criteria. A second pricing-incident attempt occurred during the rollout — a malformed rule caught by the local validation layer and rolled back in minutes, visible to exactly ten stores for one hour — and the quiet handling of that event did more to sell the architecture internally than any dashboard had.

07 The Results

A year into the chain-wide rollout, the scoreboard spoke clearly. Price-label accuracy rose above 99% across the chain, ending the drift between shelf and register and with it the compliance exposure during state pricing inspections. Shelf audits that consumed 1,400 labor hours a week now take a quarter of the time, executed as targeted tablet tasks instead of aisle-long walking inspections. Fresh-food waste fell 22% as markdowns triggered from real-time sell-through data rather than guesswork, and out-of-stocks on top-selling SKUs dropped 31%. The edge clusters kept stores operating through multiple WAN outages that previously would have halted pricing operations entirely.

METRIC	BEFORE	AFTER
Price-label accuracy	-	99%+
Fresh-food waste	-	-22%
Shelf-audit labor hours	-	-75%
Out-of-stocks, top SKUs	-	-31%

The operational culture shifted alongside the numbers. The chain now treats each store as a software-deployable unit, so a merchandising decision ships chain-wide in days instead of quarters. Store managers describe the difference simply: they learn about shelf problems from the tablet at the start of the shift rather than from a customer at the register. The reclaimed audit hours have been redirected to service and fresh presentation, and the \$300,000 pricing incident now serves as the internal argument for the architecture — the same class of failure, replayed under the new platform, was caught and rolled back before it became a story.

08 What Your Business Should Do

Orchard Lane's playbook generalizes to any multi-site operator, and it begins with restraint. Start with one painful workflow — shelf audits, markdowns or receiving — and digitize it end to end in ten pilot stores before scaling chain-wide, because a single verified win on a metric executives already track funds the rest of the rollout. Design for disconnected stores from the first sprint: assume WAN links will fail at the worst moment, and require the edge to price, audit and operate locally for days, then sync automatically when connectivity returns. Keep raw video at the edge, transmitting only aggregated events to the cloud — that single decision slashes bandwidth costs, preserves shopper privacy and reduces compliance review to a policy statement. Make anomalies actionable by converting detections into prioritized fix lists on devices staff already carry, since insight that requires a new workflow gets quietly ignored within weeks.

Then operate the estate like software. Centralize policy and localize execution: set markdown and pricing rules centrally, but execute them at the store where the data lives, balancing governance with shelf-level speed and resilience. Roll out with GitOps, store by store, in staged waves with automated health checks and instant rollback — 200 stores means 200 deployment targets, and they should be treated that way. Prove ROI on one number leadership already reviews, such as fresh-food waste or out-of-stocks, and let that verified win carry the funding case. Plan the hardware refresh from day one: standardize on replaceable edge nodes with remote management, because stores are harsh environments and fan failures, power blips and five-year refresh cycles are certainties, not risks. The macro trend is not subtle — Gartner has reportedly predicted that 75% of enterprise data will be created and processed outside traditional data centers by 2025, and IDC reportedly projects roughly \$230 billion of global edge spending in 2024 — but the business case does not depend on the trendline; it depends on how long after a price error hits your shelves anyone detects it.

Before the first pilot store, sit with the questions Orchard Lane's leadership had to answer honestly. How long after a price error hits your shelves does anyone actually detect it? What would your stores do tonight if the wide-area network link went down for a day? Which decisions in your operations wait hours or days purely because data moves in batches? How many staff hours a week go to manually checking things a camera could check? If every store were a software deployment target, how fast could you ship a fix chain-wide? Where does fresh inventory lose value fastest in your chain — and could a local rule capture it? Use the action checklist in this practice book to answer these with your team, and refresh the answers each quarter as your store formats, your labor market and your competitors evolve.

REAL-WORLD PROOF & BENCHMARKS

Gartner has reportedly predicted that by 2025, 75% of enterprise-generated data will be created and processed outside traditional data centers or centralized cloud. - Gartner

IDC's Worldwide Edge Spending Guide reportedly projects global edge computing spending of roughly \$230 billion in 2024, continuing at double-digit growth through the decade. - IDC, 2023

IHL Group research has reportedly estimated that out-of-stocks and overstocks cost retailers worldwide more than \$1 trillion annually, with grocery among the most exposed segments. - IHL Group, 2020

THE PRACTICE PAGES

The 8-action checklist

Work through the list with your team. Tick each box as your organization completes the step, and fill in an owner and a target date so every action has a name and a deadline.

☐**1. Start with one painful workflow**

Pick the highest-cost store process, whether shelf audits, markdowns or receiving, and digitize it end to end in ten pilot stores before you scale chain-wide.

OWNER

TARGET DATE

STATUS

--	--	--

☐**2. Design for disconnected stores**

Assume WAN links will fail at the worst moment. Edge clusters must price, audit and operate locally for days, then sync automatically when connectivity returns.

OWNER

TARGET DATE

STATUS

--	--	--

☐**3. Keep raw video at the edge**

Process camera feeds on-store and transmit only aggregated events to the cloud, never raw footage. This slashes bandwidth costs, preserves shopper privacy and dramatically simplifies compliance review.

OWNER

TARGET DATE

STATUS

--	--	--

☐**4. Make anomalies actionable tasks**

Convert detections into prioritized fix lists on the devices staff already carry. Insight that requires a brand-new workflow or screen gets quietly ignored within weeks.

OWNER

TARGET DATE

STATUS

--	--	--

☐**5. Centralize policy, localize execution**

Set markdown and pricing rules centrally, but execute them at the store where the data lives. This balances corporate governance with shelf-level speed and resilience.

OWNER

TARGET DATE

STATUS

--	--	--

☐

6. Roll out with GitOps, store by store

Ship software updates like code: staged waves with automated health checks and instant rollback. 200 stores means 200 deployment targets — treat them that way.

OWNER

TARGET DATE

STATUS

--	--	--

☐

7. Prove ROI on one metric

Tie the pilot to one number executives already track, such as fresh-food waste or out-of-stocks. A single verified win funds the rest of the rollout.

OWNER

TARGET DATE

STATUS

--	--	--

☐

8. Plan hardware refresh from day one

Standardize on replaceable edge nodes with remote management. Stores are harsh environments; plan for fan failures, power blips and five-year refresh cycles up front, not mid-rollout.

OWNER

TARGET DATE

STATUS

--	--	--

AFTER THE CHECKLIST: OUR TOP 3 PRIORITIES

From the eight actions above, write down the three you will start first - and the single number you expect each one to move. Three priorities with owners beat eight without any.

PRIORITY 1 NUMBER IT SHOULD MOVE	PRIORITY 2 NUMBER IT SHOULD MOVE	PRIORITY 3 NUMBER IT SHOULD MOVE
--	--	--

THE PRACTICE PAGES

Reflections for your team

Use these questions in your next planning session and type short answers into the boxes. The gaps you find are usually the fastest wins.

Q1. How long after a price error hits your shelves does anyone actually detect it?

Q2. What would your stores do tonight if the wide-area network link went down for a day?

Q3. Which decisions in your operations wait hours or days purely because data moves in batches?

Q4. How many staff hours a week go to manually checking things a camera could check?

Q5. If every store were a software deployment target, how fast could you ship a fix chain-wide?

Q6. Where does fresh inventory lose value fastest in your chain — and could a local rule capture it?

THE PRACTICE PAGES

My 90-day plan

Finish the session by committing to a plan. Three phases, one goal each, one number that tells you it worked. Bring this page to your next leadership review - or to a free consultation with Emerson Technology, and we will pressure-test it with you.

Days 1-30 Focus: audit and alignment

PRIMARY GOAL**KEY ACTIONS****HOW WE MEASURE SUCCESS****OWNER**

Days 31-60 Focus: build and prove

PRIMARY GOAL**KEY ACTIONS****HOW WE MEASURE SUCCESS****OWNER**

Days 61-90 Focus: scale and measure

PRIMARY GOAL

KEY ACTIONS

HOW WE MEASURE SUCCESS

OWNER

Tip: save a copy per team so each group keeps its own answers. Revisit the reflections every quarter - the questions stay the same even as the numbers move.

BOOK THE REVIEW BEFORE YOU CLOSE THE BOOK

Plans without a review date drift. Before you close this practice book, agree on when you will come back to these pages, who will run that session, and when you will check the numbers again. If you want an outside pair of eyes, bring this page to a free consultation with Emerson Technology and we will pressure-test the plan with you.

REVIEW DATE	REVIEWED BY	NEXT CHECK-IN
------------------------------------	------------------------------------	--------------------------------------

Ready for a transformation like Orchard Lane Markets's?

Emerson Technology designs and builds the data, AI, cloud and automation systems behind results like these. On a free consultation call we will map this exact playbook to your business - your data, your constraints, your numbers.

Book a free consultation

support@emersontechnology.com

or start online: emersontechnology.com/contact