

TECH & AI CASE STUDY - E-COMMERCE - 2026

AI Sales Forecasting That Cut Stockouts by 47% for an E-commerce Brand

Based on a delivery engagement with NorthCrest Apparel

-47%

STOCKOUT RATE

31% to 12%

FORECAST ERROR

3.1 to 4.6

INVENTORY TURNS

\$1.9M

REVENUE RECOVERED

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

Free - no email requiredemersontechnology.com

MAKE THIS PRACTICE BOOK YOURS

Work it, do not just read it

This practice book turns the NorthCrest Apparel 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 NorthCrest Apparel - 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

NorthCrest Apparel is a direct-to-consumer apparel brand built around accessible everyday fashion: tees, knitwear, denim and seasonal collections sold through its Shopify storefront, two online marketplaces and a small network of retail partners. The company employs roughly 180 people across design, sourcing, marketing, customer care and operations, and it ships tens of thousands of orders per month with pronounced seasonal peaks around back-to-school and the holiday window. Its catalog is deceptively complex for a "simple" apparel brand: around 2,400 active SKUs once every style is multiplied by size, color and fit, and each season introduces hundreds of new variants while retiring old ones.

That combination — high SKU count, fast trend cycles, sharp seasonal spikes — makes inventory the single most financially sensitive part of the business. Every unit ordered too early ties up cash and risks a markdown; every unit ordered too late becomes a stockout, an apology email and a customer who buys from a competitor.

NorthCrest's leadership knew this intimately: the founding team had grown the brand from a two-person operation into a mid-market retailer, but the planning process had never evolved beyond the spreadsheets the founders built in the early years.

02 Where Things Stood

Before Emerson Technology engaged, demand planning at NorthCrest ran on a weekly spreadsheet ritual maintained by three analysts. Every Monday, the analysts exported Shopify order history into a workbook, layered on last year's numbers, adjusted for gut-feel seasonality, and circulated buy plans to the buying team by Wednesday. The process consumed most of each analyst's week, and it was already creaking at the seams: the workbook had grown to dozens of interdependent tabs, and no single person fully understood every formula chain.

Returns made the problem worse in a way the spreadsheets never captured. Apparel returns run high, and NorthCrest's planning workbook treated a returned unit as lost demand rather than future sellable stock, so the team systematically over-bought to compensate for units that were actually coming back. Reconciliation between the Shopify export and the 3PL's warehouse figures happened monthly, by hand, which meant the planners were always steering with at least two-week-old information. During peak season, when decisions compound daily, that lag turned small errors into expensive ones.

The consequences were measurable and painful. Hero SKUs — the top sellers that anchor each collection — went out of stock during nearly every seasonal spike, producing a 28% stockout rate that the finance team estimated at \$2.3M in lost sales annually. Meanwhile, slow movers accumulated quietly in the warehouse and forced margin-destroying end-of-season markdowns. Forecast error sat at 31% MAPE (mean absolute percentage error), which is low enough to look respectable on a dashboard and high enough to be commercially useless at SKU level. Because planners could not trust the numbers, the buying team over-ordered defensively, tying up roughly \$4.1M of working capital in the wrong sizes and colors — cash that could not fund marketing, product development or new hires.

03 The Challenge

The challenge was not simply "forecasting is hard." It was the compound effect of three structural problems that reinforced each other. First, the demand signals the company needed already existed but lived in separate systems: three years of Shopify orders here, web analytics there, paid-media spend in an ad platform, promotion calendars in a marketing planner, and nobody joining them together. Second, the forecasting workload was manual, so its cadence was weekly at best — far too slow for demand patterns that shift within days during a promotion or a weather-driven spike. Third, and most corrosive, the organization had lost confidence in its own numbers, so planners layered defensive padding onto every buy, which quietly converted forecast error into excess inventory.

There was a human cost as well. The three analysts who owned the spreadsheets were talented people spending the majority of their week on copy-paste data assembly. Buying meetings started from blank templates rather than system-generated plans, so discussion time went to reconciling whose number was right instead of deciding what to do about it. Any solution therefore had to change not just the math but the daily workflow of the buying team — or it would end up as another dashboard nobody opens.

04 Discovery and Diagnosis

Emerson Technology began with a two-week discovery sprint across NorthCrest's data landscape, systems and planning rituals. The audit inventoried every potential demand signal and found that the brand already captured an estimated 80% of the inputs a modern forecasting model needs: order history with SKU-level granularity, web traffic and conversion analytics, paid-media spend by channel and campaign, promotion calendars, and — because apparel demand is weather-sensitive — historical local weather for each major shipping region. The data was scattered, but it existed.

The diagnosis also surfaced less obvious blockers. SKU taxonomy was inconsistent: duplicate product codes, inconsistent size labels across channels, and roughly 14% of catalog entries that were orphaned variants no longer in active use. Catalog hygiene, not model sophistication, would cap forecast quality — a finding that shaped the entire rollout. The audit further showed that forecast error was not evenly distributed: it concentrated in promotion weeks and in the tails of the size curve, exactly the places where a naive historical average performs worst. Finally, interviews confirmed the ownership gap — marketing planned campaigns without informing the buying team, and buying planned inventory without visibility into upcoming promotions.

05 The Solution

Emerson Technology designed and built a unified forecasting pipeline on AWS that turns NorthCrest's fragmented signals into SKU-level demand forecasts and, critically, into decisions a planner can act on. The architecture has three layers.

The data foundation

An ingestion layer pulls three years of Shopify order history, web analytics, paid-media spend, promotion calendars and local weather data into a single feature store. Every signal is normalized to the same SKU-week grain, so a model sees one coherent row per product per week instead of five disconnected exports. The SKU taxonomy was cleaned first — deduplicated codes, canonical size and color labels, orphaned variants archived — because every downstream layer inherits catalog quality. Automated validation checks now flag new data-quality issues before they can contaminate a forecast.

The forecasting ensemble

At the core, an ensemble of gradient-boosted trees and seasonal decomposition models generates SKU-level demand forecasts at 13-week granularity, retrained nightly. The gradient-boosted models capture interactions — a promotion's lift differing by category, a heatwave pulling demand forward on summer weight knitwear — while seasonal decomposition provides a stable statistical baseline that keeps forecasts grounded when patterns drift. Ensembling matters commercially: no single algorithm survives promotion shocks and new-season launches equally well, but a weighted combination degrades gracefully. Every forecast ships with prediction intervals, so the buying team sees not just a number but a confidence band.

The planner experience

A Next.js planning dashboard translates raw model output into weekly buy plans with confidence bands, exception flags and drill-downs from category to SKU. Low-latency APIs push approved replenishment suggestions straight into NorthCrest's ERP, so the forecast becomes an operational instruction rather than a screenshot. Planners now approve exceptions instead of building spreadsheets — the system proposes, the human disposes. This inversion of the workflow, not the model choice, is what made adoption stick: within a month, buying meetings started from the system-generated plan, and the analysts who once assembled workbooks became the people who manage exceptions and tune the model's business rules.

06 The First 90 Days

The rollout was deliberately narrow at the start. Weeks one and two covered discovery and the SKU taxonomy cleanup. Weeks three through six launched the pilot: forecasts for hero SKUs only — the roughly 20% of products driving 80% of revenue — compared weekly against the legacy spreadsheet process, with both runs visible to the buying team. The pilot built trust the honest way: on hero SKUs, the ensemble's error was visibly lower, especially during a mid-season promotion that the model flagged and the spreadsheet missed.

Weeks seven through ten delivered the planner dashboard and the ERP integration, converting pilot forecasts into approved buy plans and automated replenishment suggestions. Weeks eleven and twelve extended coverage from hero SKUs to the full active catalog and included hands-on training sessions for every planner and buyer. Adoption was treated as an engineering requirement, not an afterthought. Every planner received a 45-minute onboarding session, a one-page playbook describing the exception workflow, and a named contact for questions during the first month. The team agreed on a simple rule for the transition period: when the model and a planner disagreed by more than 20%, the case was logged and reviewed in the weekly buying meeting. That review loop did two things — it caught genuine model blind spots early, and it gave planners visible evidence that the system earned its confidence rather than demanding it. Within six weeks, disagreement logs had thinned to a handful of cases a week, most concentrated on newly launched styles with no sales history, which became the next modeling focus.

By the end of the first quarter the weekly spreadsheet ritual was retired, the three analysts had new exception-management roles, and the organization had a single agreed number to argue with — which, in practice, meant far less arguing.

07 The Results

Within two quarters, the improvements were unambiguous. The stockout rate fell from 28% to 15% — a 47% reduction — with the largest gains precisely where they mattered most: hero SKUs during seasonal spikes. Forecast error dropped from 31% to 12% MAPE, crossing the threshold where planners stop hedging every buy. Inventory turns improved from 3.1 to 4.6, releasing roughly \$1.2M of working capital that had been sleeping in the wrong sizes and colors. The finance team estimates \$1.9M in previously lost revenue was recovered in the first two quarters as best-sellers stayed in stock through demand peaks.

METRIC	BEFORE	AFTER
Stockout rate	28%	15%
Forecast error (MAPE)	31%	12%
Inventory turns	3.1	4.6
Revenue recovered	-	\$1.9M

The qualitative changes compounded the quantitative ones. Buying meetings now start from system-generated plans with exceptions flagged, cutting meeting time and moving the conversation from "whose number is right" to "what do we do this week." Markdown pressure eased because slow movers are identified earlier in their lifecycle. The planning team's weekly rhythm shifted from data assembly to decision-making, and the company finally has a forecasting capability that improves every night while the business sleeps.

08 What Your Business Should Do

The pattern behind NorthCrest's result is repeatable, and none of the steps require a research lab. Start by auditing every demand signal you already own: point-of-sale and e-commerce history, marketing spend, promotion calendars, seasonality and weather. Most companies discover, as NorthCrest did, that 80% of what a model needs already sits in their systems — it simply never gets joined. Next, unify your SKU taxonomy before any modeling begins; deduplicate product codes, sizes and colors into one canonical catalog, because forecast quality is capped by catalog hygiene, not by model complexity.

Prove the concept on hero SKUs only — the 20% of products driving 80% of revenue — and let the narrow scope build organizational trust before widening. Choose an ensemble of methods rather than a single model, combining machine learning with seasonal baselines so the system survives promotion shocks. Retrain on a fixed nightly cadence with drift alerts, because demand patterns move weekly and a model refreshed quarterly is a model quietly going stale. Ship a planner dashboard rather than raw charts: weekly buy plans, confidence bands and exception flags designed around the buying decision. Connect the forecast to replenishment so approved suggestions flow into your ERP automatically — a forecast that requires manual copy-paste will quietly stop being used. Finally, track forecast error and stockout rate side by side in one monthly report leadership actually reviews; what the organization reviews monthly is what it improves quarterly. Work through these eight steps with your team using the action checklist in this practice book, and revisit the reflection questions each quarter to keep the capability honest.

REAL-WORLD PROOF & BENCHMARKS

McKinsey research reports that AI-driven supply-chain forecasting typically reduces forecasting errors by 20-50% and lost sales from stockouts by up to 65%. - McKinsey & Company

Amazon's anticipatory-shipping program showed years ago that predictive models can position inventory before demand peaks, a method now standard across large e-commerce operators. - Reuters, reported 2014

Inditex (Zara) credits AI-assisted inventory allocation for its ability to refresh fashion lines in weeks — a benchmark its retail competitors still chase. - Harvard Business Review case coverage

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. Audit every demand signal you already own**

List POS, e-commerce, marketing, promotion and seasonality data sources. Most brands already capture 80% of the inputs a forecasting model needs — they just sit in separate systems.

OWNER

TARGET DATE

STATUS

--	--	--

☐**2. Unify SKU taxonomy before modeling**

Deduplicate product codes, sizes and colors into one canonical catalog. Forecast quality is capped by catalog hygiene, not by model complexity.

OWNER

TARGET DATE

STATUS

--	--	--

☐**3. Start with hero SKUs only**

Pick the 20% of products driving 80% of revenue. Proving wins on a narrow scope builds trust and pays for the wider rollout.

OWNER

TARGET DATE

STATUS

--	--	--

☐**4. Choose an ensemble, not one model**

Combine gradient boosting with seasonal baseline methods. Ensembles survive promotion shocks and new-season launches far better than any single algorithm.

OWNER

TARGET DATE

STATUS

--	--	--

☐**5. Retrain on a fixed nightly cadence**

Demand patterns drift weekly. Automated retraining with drift alerts keeps error low without a data scientist on call.

OWNER

TARGET DATE

STATUS

--	--	--

☐

6. Ship a planner dashboard, not raw charts

Planners need weekly buy plans with confidence bands and exception flags — not model output. Design the interface around the buying decision.

OWNER

TARGET DATE

STATUS

☐

7. Connect forecasts to replenishment

Push approved suggestions into your ERP or 3PL automatically. A forecast that requires manual copy-paste will quietly stop being used.

OWNER

TARGET DATE

STATUS

☐

8. Track MAPE and stockouts monthly

Publish forecast error and stockout rate side by side in one report. What leadership reviews monthly is what the org improves quarterly.

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 much revenue did stockouts on your ten best-selling products cost last quarter — and do you actually have that number?

Q2. Which team owns demand forecasting today, and how many hours a week do they spend maintaining spreadsheets?

Q3. If your forecast error dropped by half, how much working capital could you release from slow-moving inventory?

Q4. Are your promotions and campaigns visible to whoever plans inventory, or do marketing and buying plan in isolation?

Q5. What fraction of your SKUs genuinely need ML-level forecasts versus a simple seasonal baseline?

Q6. When a stockout happens today, how long does it take your team to detect it — before or after customers complain?

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 NorthCrest Apparel'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