

TECH & AI CASE STUDY - E-COMMERCE - 2026

One Snowflake Warehouse That Cut Reporting Time from Days to Minutes

Based on a delivery engagement with Harbor & Pine Home Goods

5 days to <15 min
REPORT TURNAROUND

23
SOURCES UNIFIED

55/week
ANALYST HOURS FREED

78%
DASHBOARD ADOPTION

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 Harbor & Pine Home Goods 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 Harbor & Pine Home Goods - 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

Harbor & Pine Home Goods is a 140-employee e-commerce retailer selling furniture, decor and seasonal home accents through its own storefront and major online marketplaces. The business lives and dies on fast merchandising cycles: paid media is tuned weekly, category pages are re-merchandised constantly, and inventory decisions on bulky, high-cost items like sofas and dining sets carry long cash consequences when they go wrong. An annual ad budget of \$6.5M funds growth across Meta, Google and email, and the finance team runs a tight monthly close to keep the board informed.

For a company this size, the analytics ambition had always outrun the analytics plumbing. Growth had been fast, systems had multiplied, and the data layer had been assembled reactively — one export at a time — until the act of simply knowing last week's revenue had become a project.

02 Where Things Stood

Analytics ran across a patchwork: Shopify exports, Google Sheets maintained by hand, a legacy Postgres replica nobody wanted to touch, and three marketing platforms each with their own attribution logic. Twenty-three sources in total, none authoritative. Every Monday, two analysts spent up to 30 combined hours reconciling figures that rarely matched — finance reported one revenue number, merchandising reported another, and both were derived from different exports with different refund handling and different timezone cutoffs. The reconciliation was done in spreadsheets with color-coded cells, and its output was stale by Tuesday.

Leadership decisions on inventory, paid media and pricing waited on reports that took five days to assemble, so category and marketing calls were routinely made on data old enough to be misleading. Quarter-end close stretched to nine days because every reported figure had to be traced back to source exports by hand — an audit trail of emailed spreadsheets and tribal memory. Disagreements over which dashboard was correct consumed senior meetings, and the \$6.5M ad budget was optimized against stale, unreliable data: campaigns were paused or scaled on numbers that were days old and occasionally wrong. Gartner's estimate that poor data quality costs the average organization \$12.9 million a year, with the figure rising for firms running fragmented data estates, read like a description of Harbor & Pine's operating reality.

03 The Challenge

The challenge was not that Harbor & Pine lacked data; it had too much of it, in too many shapes, with no agreed way to combine it. Each source told a partial story — Shopify knew orders but not ad spend, the ad platforms knew spend but not margin, the ERP knew inventory but not returns — and the only place the pieces met was an analyst's spreadsheet, assembled by hand, on a schedule the business had outgrown. The warehouse-of-spreadsheets approach also had a fragility problem: the reconciliation process lived in two analysts' heads and their personal workbooks, and either resignation would have halted reporting for weeks.

There was a strategic cost as well. McKinsey's research reports that data-driven organizations are 23 times more likely to acquire customers and 19 times more likely to be profitable than their peers — and Harbor & Pine was competing in paid-media auctions against rivals repricing and retargeting daily while its team worked from last week's exports. Any solution had to deliver three things at once: one trusted definition of revenue and margin, freshness measured in hours rather than days, and a way for non-analysts to answer their own questions without recreating the spreadsheet economy under a new name.

04 Discovery and Diagnosis

Emerson Technology began with a two-week source-mapping exercise, on the principle that a warehouse project succeeds or fails on the completeness of this map rather than on vendor choice. The team inventoried all 23 systems holding customer, order, marketing and finance data, and for each one documented access method, update frequency, data quality and the business questions it answered alone. The map surfaced uncomfortable specifics: refund records were split between Shopify and the ERP with different timestamps; marketplace channels reported in platform-specific taxonomies; and three sources had no native connectivity at all. The diagnosis phase then followed the money and the meetings. The team reconstructed the Monday reconciliation line by line to identify why finance and merchandising numbers diverged — the root causes were refund timing, marketplace fee treatment and a timezone mismatch in daily cutoffs, none of them visible in any spreadsheet. A reporting inventory showed that the five-day turnaround existed not because analysis was hard but because data assembly consumed the analysts' week. Forrester's Total Economic Impact studies of cloud data platforms, which typically report triple-digit returns over three years driven mainly by analyst productivity and infrastructure savings, gave the CFO a framework for the business case; the diagnosis gave the engineering plan its sequence — ingestion first, definitions second, self-serve third.

05 The Solution

Emerson Technology consolidated everything into a single Snowflake warehouse with dbt transformations and automated pipelines, replacing the export economy with one source of truth. The architecture has three layers.

Landing the data automatically

Fivetran connectors land raw data from Shopify, Klaviyo, Meta Ads, Google Ads, the ERP and the fulfillment systems on a fixed schedule, while custom Python jobs handle the three sources without native connectors. Nothing depends on an analyst remembering to export; ingestion is monitored, alerted and owned. This was deliberately sequenced first and automated from day one, because hand-built, unmonitored pipelines are where freshness and trust quietly break down first. Incremental models keep warehouse spend predictable as volume grows, so the platform's cost curve follows the business rather than spiking at every peak season.

One transformation layer, one truth

dbt builds the transformation layer, and every metric has exactly one tested, documented definition — from net revenue to contribution margin — with refund timing, marketplace fees and timezone cutoffs resolved once, centrally, in code. Automated data-quality tests run on every load: freshness checks, uniqueness checks and row-count checks that block bad batches from publishing, so a broken feed fails loudly in the pipeline instead of quietly in a board meeting. Zero-copy environment cloning lets analysts and engineers test schema changes and new models safely against production-like data, guaranteeing that experimentation never interrupts the numbers leadership depends on. Governance is enforced in one place — role-based access, row-level security and automated lineage documentation — rather than scattered across tools.

Serving two audiences deliberately

On top sits a semantic layer feeding two deliberately different experiences. Executives get Power BI dashboards: a handful of stable, polished charts of revenue, margin, ad efficiency and inventory position, refreshed automatically and identical for everyone. Category managers get a self-serve exploration workspace with governed access to the underlying models, so they can slice a category, a channel or a promotion without filing a ticket — and without being able to redefine revenue along the way. Designing the two experiences separately, instead of forcing one interface to serve both, avoided the most common way BI rollouts satisfy neither audience.

06 The First 90 Days

Weeks one and two completed the 23-source map and the root-cause analysis of the finance-versus-merchandising split. Weeks three through six stood up the Snowflake foundation and landed the first six sources through Fivetran, with the custom Python connectors for the three orphan systems built in parallel. The team chose an ingestion-first rollout for a reason: every week the raw data flowed reliably, the Monday reconciliation ritual shrank, and the two analysts who owned it began reviewing pipeline alerts instead of coloring cells. Their feedback loop became the project's early-warning system — they knew every historical quirk in the data, and catching a misload in week four bought credibility that no launch announcement could.

Weeks seven through ten delivered the dbt transformation layer and the metric definitions. Finance and merchandising sat in the same room to agree net revenue, contribution margin and ad-spend attribution once; the arguments were real but finite, and each resolution was written into code with tests attached, never into a slide. The old spreadsheets were not simply switched off — each one received a published sunset date, and the dates were enforced, because silent migration never happens and visible deadlines are what force adoption of a new source of truth. Weeks eleven and twelve shipped the executive dashboards, opened the self-serve workspace to category managers with 45-minute onboarding sessions, and established a weekly operating review that tracked adoption — weekly active users per dashboard — alongside pipeline health.

By the end of the quarter the Monday ritual was gone. The two analysts spent their hours building cohort analyses and margin models instead of reconciling exports, and for the first time, a surprising number on a dashboard could be traced to source data in minutes using automated lineage.

07 The Results

Reporting turnaround collapsed from five days to under 15 minutes for standard views, and the Monday reconciliation ritual disappeared entirely. The 55 analyst-hours previously lost each week now go to analysis instead of copy-paste work. With one agreed definition of revenue and margin, the marketing team reallocated \$800K of underperforming ad spend within a quarter — moving budget toward channels and audiences the reliable data actually supported. Finance shortened its monthly close by four days, ending the quarter-end scramble that had consumed nine days of tracing and re-tracing. Dashboard adoption reached 78% of staff who previously waited on emailed spreadsheets.

METRIC	BEFORE	AFTER
Report turnaround	5 days	Under 15 minutes
Sources unified	-	23
Analyst hours freed	-	55/week
Dashboard adoption	-	78%

The second-order effects compounded. Paid media decisions now happen mid-week on fresh data, which matters in auctions where competitors adjust daily. The finance-versus-merchandising revenue argument ended, and senior meetings moved from litigating numbers to acting on them. The platform also survived its first peak-season traffic spike without manual intervention — the incremental models and automated tests absorbed the volume that would previously have broken the export process.

08 What Your Business Should Do

The sequence Harbor & Pine followed is the template, and it begins before any platform decision. Map every source before buying anything: inventory the systems holding customer, order, marketing and finance data, because the warehouse project succeeds or fails on the completeness of that map. Then pick one metrics layer with strict definitions — define each metric once in dbt or a semantic layer, with tests and named owners, since conflicting definitions of revenue are the fastest way to lose executive trust. Automate ingestion from day one with managed connectors for standard SaaS sources and custom code only for edge cases; hand-built, unmonitored pipelines are where freshness and trust quietly die first.

Add tests that block bad loads — freshness, uniqueness and row-count checks on every critical table, failing loudly before dashboards publish rather than after an executive spots a wrong number. Clone environments before touching production, using zero-copy clones to test schema changes safely so experimentation never interrupts the numbers leadership depends on. Ship executive and self-serve experiences separately, because a handful of stable charts for leaders and raw governed exploration for analysts are different products, and designing one interface for both is why most BI rollouts satisfy neither. Retire legacy reports on a public schedule with enforced sunset dates — silent migration never happens — and track adoption, not just uptime, by reporting weekly active users per dashboard alongside pipeline health. A technically perfect warehouse that nobody opens has failed, however green the monitoring looks.

Close by pressure-testing your own operation with the questions this program answered: if two teams quoted last quarter's revenue today, would they give the same number, and how long would finding out take; how many analyst hours does your organization burn each month reconciling data that should reconcile itself; would your current stack survive a 3x volume increase, or does every peak season become an outage. McKinsey's finding that data-driven organizations are 23 times more likely to acquire customers and 19 times more likely to be profitable is not about owning a warehouse — it is about decisions made on fresh, agreed numbers. Work through the eight actions above with your data and finance leads together, and revisit the reflection questions in this practice book each quarter to keep the single source of truth single.

REAL-WORLD PROOF & BENCHMARKS

Gartner estimates that poor data quality costs the average organization \$12.9 million per year, a figure that rises for firms with fragmented, siloed data estates. - Gartner, 2021

McKinsey research reports that data-driven organizations are 23 times more likely to acquire customers and 19 times more likely to be profitable than their peers. - McKinsey & Company

Forrester's Total Economic Impact studies of cloud data platforms have typically reported triple-digit returns over three years, driven mainly by analyst productivity and infrastructure savings. - Forrester TEI

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. Map every source before buying anything**

Inventory the systems that hold customer, order, marketing and finance data. A warehouse project succeeds or fails on the completeness of this map, not on vendor choice.

OWNER

TARGET DATE

STATUS

--	--	--

☐**2. Pick one metrics layer with strict definitions**

Define each metric once in dbt or a semantic layer, with tests and owners. Conflicting definitions of revenue are the fastest way to lose executive trust.

OWNER

TARGET DATE

STATUS

--	--	--

☐**3. Automate ingestion from day one**

Use managed connectors for standard SaaS sources and reserve custom code for edge cases. Hand-built, unmonitored pipelines are where freshness and trust quietly break down first.

OWNER

TARGET DATE

STATUS

--	--	--

☐**4. Add tests that block bad loads**

Set freshness, uniqueness and row-count tests on every critical table. Fail loudly before dashboards publish — not after an executive spots a wrong number in a board meeting.

OWNER

TARGET DATE

STATUS

--	--	--

☐**5. Clone environments before touching production**

Use zero-copy clones to test schema changes and new models safely. This keeps experimentation fast while guaranteeing the production numbers leadership depends on are never interrupted by mistakes.

OWNER

TARGET DATE

STATUS

--	--	--

☐

6. Ship executive and self-serve separately

Executives need a handful of stable, polished charts; analysts need raw exploration and SQL access. Designing one interface for both audiences is why most BI rollouts satisfy neither.

OWNER

TARGET DATE

STATUS

--	--	--

☐

7. Retire legacy reports on a public schedule

Publish a sunset date for each old spreadsheet and dashboard, then enforce it. Silent migration never happens; visible deadlines force adoption of the new source of truth.

OWNER

TARGET DATE

STATUS

--	--	--

☐

8. Track adoption, not just uptime

Report weekly active users per dashboard alongside pipeline health every week. A technically perfect warehouse that nobody opens has failed, no matter how green the monitoring looks.

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. If two teams quoted revenue for last quarter today, would they give the same number — and how long would finding out take?

Q2. How many analyst hours does your organization burn each month reconciling data that should reconcile itself?

Q3. Which five metrics actually drive your business decisions, and does everyone agree on how they are calculated?

Q4. What percentage of your reports are still assembled by hand, and what would happen if their owner resigned tomorrow?

Q5. When a dashboard shows a surprising number, can your team trace it to source data in minutes?

Q6. Would your current data stack survive a 3x increase in volume, or does every peak season become an outage?

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 Harbor & Pine Home Goods'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