

TECH & AI CASE STUDY - PAYMENTS PLATFORM - 2026

Fraud Scoring That Cut False Declines 61%

Based on a delivery engagement with Clearport Payments

-61%
FALSE DECLINES

-38%
FRAUD LOSSES

P99 85ms
DECISION LATENCY

-52%
REVIEW QUEUE

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 Clearport Payments 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 Clearport Payments - 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

Clearport Payments is a payment service provider built for marketplaces: it processes \$2.8B a year on behalf of 4,100 merchants, ranging from single-person vintage shops to multi-category sellers doing eight figures annually. Marketplace payments are structurally harder than ordinary card-not-present processing, because Clearport sits between two sets of customers it must keep happy at once — the buyers whose cards get charged and the merchants whose payouts depend on approvals. The company's product promise to merchants is simple: maximize legitimate approvals, keep chargebacks inside card-network thresholds, and never make a seller explain a mysterious decline to their own customers.

Keeping that promise is an economics problem as much as a technology problem. Every fraudulent transaction approved costs money and chargeback ratios; every legitimate transaction declined costs revenue, merchant trust and, in most cases, the customer permanently, because shoppers who hit a false decline rarely retry. The payments industry knows this imbalance well — Javelin Strategy & Research has estimated that false declines cost US merchants around \$118 billion in lost sales, many multiples of actual fraud losses. Clearport's leadership understood the theory. What they lacked was a decisioning system capable of acting on it in real time.

02 Where Things Stood

Clearport's fraud defenses ran on a static rules engine accumulated over years: block this BIN range, flag that country pair, throttle velocity above a threshold. The engine blocked anything unusual, and unusual is a poor proxy for fraudulent on a marketplace where sellers spike legitimately around product launches and holiday demand. Roughly 4.3% of legitimate volume was being declined as suspected fraud. Merchants felt every one of those declines, because their customers rarely came back — several large sellers had begun openly threatening to switch processors, which turned fraud policy from a back-office concern into a retention crisis.

The engine was not even keeping fraud contained in exchange for the friction. Chargebacks ran at 0.74% — uncomfortably close to card-network program thresholds — and fraud losses reached \$11M a year. Rule changes were the operational nightmare: every new rule risked breaking an entire merchant category overnight, as when a policy tweak intended to catch resale bots quietly began declining a legitimate merchant's entire drops-and-limited-edition line. Analysts reviewed 2,100 queued transactions daily at a 40% false-positive rate, which meant the manual safety net was itself half noise.

The organizational cost was as corrosive as the financial one. Risk and growth operated as the same weekly meeting with different dashboards and no shared number: the risk team optimized for losses, the commercial team optimized for approval rates, and every threshold debate ended in a compromise neither side believed. Fraud had become the tax both teams paid for the absence of a decisioning system that could separate the two objectives.

03 The Challenge

The challenge was to build real-time transaction scoring accurate enough to loosen the rules without loosening the controls — and to do it at marketplace scale, where a single flash sale can multiply traffic in minutes. The Nilson Report has estimated global card fraud losses at roughly \$33 billion in a recent year, with projections exceeding \$40 billion within a decade, and Juniper Research has projected cumulative online payment fraud losses exceeding \$340 billion globally across 2023-2027, driven heavily by account takeover and social engineering. Attack pressure was rising on one side while merchant patience was falling on the other; standing still meant losing on both fronts simultaneously.

There was also a latency constraint that shaped everything. Checkout is a real-time experience, and a fraud engine that adds even a few hundred milliseconds of decision time taxes every legitimate transaction to guard against a fraudulent few. Any replacement for the rules engine had to grade transactions fast enough that shoppers never feel it — and it had to explain itself, because analysts needed to understand why transactions were flagged before anyone would trust the system enough to lower a threshold.

04 Discovery and Diagnosis

Emerson Technology began with a three-week discovery across Clearport's transaction pipeline, rules inventory and analyst workflow. The audit reconstructed the economics first: by matching declined transactions against subsequent chargeback and refund outcomes, the team sized the false-decline problem precisely — 4.3% of legitimate volume, worth far more than the fraud it was preventing. This confirmed the pattern Javelin's research describes at industry scale, and it reframed the project internally: this was a revenue-recovery program as much as a fraud program. Framing mattered, because it gave risk and growth their first shared number.

The technical audit then mapped what signals existed and where they lived. Card velocity, device fingerprints, session behavior and merchant-category history were all being captured somewhere, but they were computed in batch or not at all at decision time — the rules engine saw a thin slice of the context a modern model needs. The team also reviewed the manual queue and found analysts spending their hours on cases with obvious answers, while genuinely ambiguous cases waited behind them. The diagnosis was blunt: the company had static rules where it needed streaming features, blunt thresholds where it needed calibrated scores, and opaque flags where it needed explanations.

05 The Solution

Emerson Technology designed and built a streaming decisioning platform on AWS that scores every transaction as it happens and enforces only what the business intends. The architecture has three layers.

The real-time feature layer

Transaction events flow through Kafka into a Redis feature store that computes velocity, device-fingerprint, behavioral and merchant-category features in real time. Card-testing bursts, bust-out patterns and account-takeover sequences show up as feature movements within minutes rather than surfacing days later in batch reports. Features are versioned and shared across models, so a velocity definition agreed once is reused everywhere — the end of five teams computing the same rate five different ways. This layer is what made the models possible; without fresh features at decision time, no model can outperform the rules it replaces.

Scoring with guardrails

An XGBoost model paired with a neural sequence model scores every transaction, holding decision latency at 85 milliseconds at P99 even at peak load. The gradient-boosted model captures feature interactions — a device seen across many cards, an amount unusual for the merchant category — while the sequence model reads the

session's behavioral trajectory, catching the hesitation-and-retry signature of manual card testing. Around both models sit explicit rules-as-code guardrails for sanctions, policy and edge cases the business must decide deterministically; the models operate inside legal and commercial boundaries they cannot see or learn on their own. Threshold changes are promoted through the same experiment pipeline as the models themselves, so no one adjusts production risk behavior by editing a config file on a Friday.

Explainable review

A Next.js analyst console shows each flag's top contributing features via SHAP explanations, so every queued case opens with a reason rather than a bare score. Case management stays integrated with Clearport's existing queues, so analysts kept their workflows and gained context instead of learning a new tool. The console matters commercially as well as operationally: explained decisions speed reviews, and they built the internal trust that later made loosening thresholds politically possible. Model retraining runs monthly with refreshed labels from chargeback outcomes, and the deployment passed through six weeks of shadow mode before enforcement, then champion-challenger, so new models prove themselves on live economics before they get a vote.

06 The First 90 Days

Weeks one and two built the baseline: decline reasons, approval rates and chargebacks instrumented by merchant category and geography, giving the team the quarter-by-quarter visibility it had never had. Weeks three through eight ran shadow mode — the new models scored live traffic without enforcing anything, and their hypothetical decisions were compared nightly against the incumbent rules engine. Shadow mode surfaced blind spots while revenue stayed fully protected: it caught, for example, a behavioral feature that misread one merchant's subscription billing pattern, which would have declined a healthy category at scale had the model gone straight to enforcement.

Weeks nine through twelve moved the models into champion-challenger and then into graded enforcement, starting with the merchant categories where shadow results showed the largest gap between model and rules. Every promotion required measured economics — approval lift and fraud delta together — not offline accuracy alone. Analysts were trained on the SHAP console during the same window, and the review queue's 40% false-positive rate became the visible scoreboard of progress. By the end of the first quarter, risk and growth were reviewing one dashboard in one weekly meeting, and threshold changes had moved from hallway arguments to experiment readouts.

07 The Results

Within two quarters the platform changed the company's economics. False declines fell 61%, releasing legitimate volume the old rules had been quietly refusing — approvals recovered several points of transaction volume, and merchants noticed before the sales team could even announce it. Fraud losses dropped 38% to \$6.8M annualized, proof that the lift came from better decisions rather than looser ones. The analyst review queue shrank 52%, decision latency holds at P99 85 milliseconds even at peak, and merchant complaints tied to declines fell by roughly half.

METRIC	BEFORE	AFTER
False declines	4.3% of legitimate volume	-61%
Fraud losses	\$11M a year	\$6.8M (-38%)
Decision latency	-	P99 85ms
Review queue	2,100 transactions daily	-52%

The qualitative changes compounded the quantitative ones. Risk and growth now share one score, one dashboard and one weekly review, and the recurring argument about whose numbers are right has been replaced by arguments about which experiment to run next. Analysts spend their time on genuinely ambiguous cases instead of obvious noise, and the large sellers who had been threatening to switch processors became the deployment's loudest advocates. The false-decline number now appears in the same monthly report as the fraud number, which is the single clearest sign the culture shifted.

08 What Your Business Should Do

The pattern behind Clearport's result is repeatable, and it starts with measurement rather than modeling. Quantify false declines first: measure approvals, declines and chargebacks by merchant category and geography for a full quarter, because most platforms discover that false declines cost multiples of the fraud they were preventing. Shadow-mode new models before enforcement — score live traffic without acting on it and compare against the incumbent rules, surfacing blind spots while revenue stays fully protected. Wrap models in rule guardrails so sanctions, policy and edge cases stay deterministic, and engineer velocity features across card, device and merchant windows, because speed-based features catch card-testing and bust-out patterns hours before static models notice anything at all.

Budget latency like revenue: set a P99 target and test it under peak load before launch, because fraud tools that add checkout friction trade real revenue for marginal accuracy. Explain decisions to analysts with top contributing features shown inside the review console — explained decisions speed reviews and build the trust needed to loosen thresholds safely. Deploy champion-challenger continuously rather than replacing models annually, since static fraud models decay within months as attack patterns shift underneath them. And close the loop with merchants by sharing decline reasons and feedback channels every month, because merchants who understand decline logic fix checkout fraud signals at the source instead of calling support.

Ask the questions Clearport now asks quarterly: do you know your false-decline rate as precisely as your fraud rate, and who reports it monthly; how much revenue disappears when a legitimate customer is declined, and does anyone own that number; which review decisions could a model explain well enough for your analysts to trust today; what is your decision latency at peak measured from the customer's side of the checkout; how quickly could you deploy a new fraud model without risking a merchant category overnight; and if fraud losses fell 38%, where would you reinvest the difference first. Work through the eight actions in this practice book with your team, and revisit those reflection questions each quarter to keep the decisioning honest.

REAL-WORLD PROOF & BENCHMARKS

The Nilson Report estimated global card fraud losses at roughly \$33 billion in a recent year, with industry projections exceeding \$40 billion within a decade. - The Nilson Report

Juniper Research projected cumulative online payment fraud losses exceeding \$340 billion globally across the 2023-2027 period, driven heavily by account takeover and social engineering. - Juniper Research

Javelin Strategy & Research has reported that false declines cost US merchants an estimated \$118 billion in lost sales — many multiples of actual fraud losses. - Javelin Strategy & Research

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. Quantify false declines first

Measure approvals, declines and chargebacks by merchant category and geography for a full quarter. Most platforms discover false declines cost multiples of the fraud they were preventing.

OWNER

TARGET DATE

STATUS

--	--	--

☐

2. Shadow-mode new models

Score live traffic without enforcing decisions and compare results against the incumbent rules engine. Shadow mode surfaces model blind spots while revenue stays fully protected throughout.

OWNER

TARGET DATE

STATUS

--	--	--

☐

3. Wrap models in rule guardrails

Keep explicit rules for sanctions, policy and edge cases around model outputs. Guardrails hold the model inside business and legal boundaries it cannot see or learn.

OWNER

TARGET DATE

STATUS

--	--	--

☐

4. Engineer velocity features

Track card, device and merchant velocity windows in real time. Speed-based features catch card-testing and bust-out patterns hours before static models notice anything at all.

OWNER

TARGET DATE

STATUS

--	--	--

☐

5. Budget latency like revenue

Set a P99 latency target and test it under peak load before launch. Fraud tools that add checkout friction trade real revenue for marginal accuracy.

OWNER

TARGET DATE

STATUS

--	--	--

**6. Explain decisions to analysts**

Show the top contributing features for every flag inside the review console. Explained decisions speed reviews and build the trust needed to loosen thresholds safely.

OWNER

TARGET DATE

STATUS

--	--	--

**7. Deploy champion-challenger continuously**

Run new models alongside production and promote only on measured economics. Static fraud models decay within months as attack patterns shift underneath them every quarter.

OWNER

TARGET DATE

STATUS

--	--	--

**8. Close the loop with merchants**

Share decline reasons and feedback channels with merchant-facing teams every month. Merchants who understand decline logic fix checkout fraud signals at the source instead of calling support.

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	PRIORITY 2	PRIORITY 3
NUMBER IT SHOULD MOVE	NUMBER IT SHOULD MOVE	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. Do you know your false-decline rate as precisely as your fraud rate — and who reports it monthly?

Q2. How much revenue disappears when a legitimate customer is declined, and does anyone own that number?

Q3. Which review decisions could a model explain well enough for your analysts to trust today?

Q4. What is your decision latency at peak, measured from the customer's side of the checkout?

Q5. How quickly could you deploy a new fraud model without risking a merchant category overnight?

Q6. If fraud losses fell 38%, where would you reinvest the difference first, and why?

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 Clearport Payments'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