

TECH & AI CASE STUDY - HOSPITAL SYSTEM - 2026

Readmission Risk ML That Cut 30-Day Readmissions 19%

Based on a delivery engagement with MercyPoint Health

-19%

30-DAY READMISSIONS

2.6x

CARE-TRANSITION ENROLLMENT

\$6.1M/yr

AVOIDED READMISSION COST

90%

HIGH-RISK FLAGS ACTED ON

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 MercyPoint Health 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 MercyPoint Health - 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

MercyPoint Health is a nonprofit community hospital system operating four acute-care hospitals across a three-county region: a large urban flagship that absorbs the system's most complex cases, supported by three smaller community hospitals embedded in surrounding towns. The system handles tens of thousands of inpatient discharges a year, and its payer mix leans heavily toward Medicare, which places federal reimbursement policy on the board agenda. A care-transitions program already existed on paper: pharmacist-led medication reconciliation, post-discharge nurse calls and scheduled follow-up appointments for patients heading home with significant conditions. Its capacity, however, covered only a small fraction of the discharges that clinically qualified for it, and the program had never been able to answer the question that mattered — which patients needed it most.

For a system with that payer mix, readmissions sit at the exact intersection of clinical quality and financial survival. The Agency for Healthcare Research and Quality reports that roughly one in five Medicare patients is readmitted within 30 days of discharge, at a cost of billions annually. CMS's Hospital Readmissions Reduction Program can reduce a hospital's Medicare payments by up to 3% for excess readmissions, and MercyPoint's exposure had been drifting in the wrong direction for several consecutive reporting cycles. Internally, the finance team estimated each unplanned return stay at an average of \$15,200 once clinical costs were tallied, which turned every preventable readmission into a number leadership could feel.

02 Where Things Stood

Deciding who received scarce follow-up resources was a weekday morning ritual built on instinct and a crude rules-based screen. Each day, case managers worked through a discharge list flagged by a rule built on prior admissions within a lookback window and a short list of index diagnoses; beyond that net, selection depended on whichever case manager happened to know the patient. The process rewarded memory over evidence: veteran staff knew certain wards intimately, but that knowledge lived in their heads and retired when they did. The blind spot showed up plainly in the system's own records: roughly 60% of patients who returned within 30 days had never been flagged as high risk before they walked back through the door. Post-discharge follow-up calls were spread evenly across all discharges, so the highest-risk patients were reached no faster than routine surgical recoveries — sometimes slower, because complex patients are the hardest to reach by phone. Pharmacists spent reconciliation hours double-checking straightforward medication lists while patients juggling multiple changed prescriptions went home without any review at all — allocation was flat by design.

The consequences compounded quietly. Penalty exposure under the readmissions program kept growing, and each avoidable return carried the estimated \$15,200 cost on top of the reimbursement risk. Meanwhile, five years of Epic-derived records — admissions, comorbidities, medication changes, utilization history, social-risk indicators — sat effectively unused for prediction. The data needed to see these patients coming already existed inside the organization; nothing was reading it.

03 The Challenge

The challenge was not that readmission prediction is exotic; the clinical literature has described the risk factors for years. It was the compound effect of structural problems that reinforced one another. The predictive signals were fragmented across the EHR, so no single view of a patient's risk existed at the moment of discharge. The screening rule was too blunt to catch first-time returners, which meant the flagging process itself capped the program's reach. And resource allocation was evenly spread, converting every prediction error into a missed opportunity for a patient who needed the pharmacist or the seven-day follow-up far more than someone else did. Underneath all three sat a feedback gap: nobody systematically tracked which flags had been correct, so the process could not learn from its own misses.

There was a human cost as well. Case managers knew they were guessing, and they said so in huddles; several described the daily selection exercise as the least defensible part of their week. Any solution therefore had to meet two hard constraints: it had to live inside the EHR workflow, because case managers had no spare hours for another portal, and it had to earn clinical trust, because a score clinicians silently override is decoration rather than infrastructure. Emerson Technology's brief was to change who gets reached after discharge, not merely to produce a more sophisticated number nobody acts on.

04 Discovery and Diagnosis

Emerson Technology began with a two-week discovery sprint across MercyPoint's data landscape and care-transitions operations. The data audit confirmed that five years of Epic-derived history was available and unusually complete: every inpatient admission, comorbidity coding, medication-change events, prior utilization and a set of social-risk indicators. Nearly everything a modern readmission model needs was already being collected; it simply had never been joined, engineered into features or pointed at the 30-day outcome question. The privacy review ran in parallel, sketching the HIPAA-aligned access model the build would later adopt. The diagnostic analysis rebuilt two years of return history and measured exactly where the existing process failed. It confirmed the 60% miss rate and showed the misses were not random: they concentrated among patients with recent medication changes, heavy polypharmacy, limited caregiver support and weekend discharge timing — profiles the rules screen was structurally blind to. The screen reliably caught frequent utilizers while missing the first-time returner whose discharge plan quietly unraveled at home. Peer-reviewed work in JAMA and Health Affairs has repeatedly reported that targeted care-transition interventions measurably reduce 30-day readmissions versus usual care. MercyPoint's problem was targeting: the interventions worked, but they were pointed at the wrong patients, or at everyone equally.

Call logs confirmed the flat allocation of follow-up effort, and case manager interviews surfaced the trust conditions any model would have to satisfy: scores explainable at the bedside, visible inside the existing EHR screens, and subject to override without punishment. Those three requirements shaped the architecture more than any algorithmic choice did.

05 The Solution

Emerson Technology designed and built a discharge-time risk system in Python around a gradient-boosted readmission model, installed inside the EHR workflow rather than beside it. The architecture has three layers.

The clinical data foundation

dbt pipelines pull and refresh the feature set nightly from the Epic-derived extracts: admission history, comorbidity profiles, medication-change counts, utilization patterns and social-risk indicators. Features were engineered alongside the system's own clinicians, who pushed for variables the analytics team would never have prioritized, such as whether a medication change was one the patient actually understood. Every signal is normalized to the

same patient-encounter grain so the model sees one coherent record per admission instead of disconnected reports. Automated validation checks flag missing feeds and coding drift before they can contaminate a score, and all access runs under HIPAA-aligned controls with de-identified data for training and identified data only at the point of care.

The model and its reason codes

At the core, the gradient-boosted model scores every inpatient daily throughout the stay and produces a final risk tier at discharge, when the care-transition decision actually happens. Each score ships with reason codes — the specific factors pushing the estimate, such as a cluster of recent admissions, a comorbidity pattern, a spike in medication changes or a utilization history that resembles past returners. A case manager can interrogate the score at the bedside and get an answer in clinical language rather than a probability with no provenance. MLOps monitoring tracks feature and prediction drift on a fixed cadence, so the analytics team sees when the model needs retraining instead of discovering staleness through a bad quarter of outcomes.

The handoff into care transitions

High-risk patients route automatically into the expanded care-transitions program: pharmacist medication reconciliation before or immediately after discharge, a nurse call within 72 hours, and a scheduled follow-up appointment within seven days. The tier and its reason codes surface directly on the EHR screens where discharge planning already happens, not on a separate dashboard that requires a second login and a second habit. Case managers approve the handoff or override it with a logged reason, and the override log feeds a weekly clinical review that doubles as model-quality feedback. Flag-to-handoff concordance is tracked as a first-class operational metric, because a prediction only prevents a readmission when it changes what the transitions team does that week.

06 The First 90 Days

The rollout was deliberately narrow at the start. Weeks one and two established governance and the data foundation: a clinician council of case managers, hospitalists, a senior pharmacist and the privacy officer met twice weekly to define what a good flag looked like, agree on the override rule and sign off on the access model. That council mattered more than any technical milestone: the clinicians who would live with the scores authored the acceptance criteria themselves.

Weeks three through six built and validated the model. The gradient-boosted model was trained on the five years of history, backtested against known 30-day outcomes, then run in shadow mode at the flagship hospital, scoring discharges daily without touching anyone's workflow. The backtest made the case to skeptics: the model caught a substantial share of returning patients the rules screen had missed, concentrated exactly in the medication-change and social-risk profiles the diagnostic phase had identified.

Weeks seven through ten delivered the EHR integration and the expanded program capacity. Pharmacist reconciliation hours were extended and nurse-call staffing increased so the routing had somewhere real to send patients; a risk tier with no receiving program is just a banner. Pilot routing went live at the flagship, and every override was logged with a reason and reviewed in the weekly huddle. The weekly review caught genuine model blind spots early and reassured clinicians that their judgment still governed the decision rather than being replaced by it.

Weeks eleven through thirteen extended the system to the remaining three hospitals. Each case-management team received hands-on training and a one-page playbook for the handoff workflow, and drift monitoring went live before volume arrived rather than after. By the end of the quarter, all four hospitals ran one selection process, and the morning list had changed character: a short, explained, ranked queue of who needed help first.

07 The Results

The first year of operation settled the argument the pilot had started. Thirty-day readmissions fell 19% across the four-hospital system, with the reduction concentrated exactly where the model directed resources — the high-risk tiers the old screen never caught. Enrollment in care transitions rose 2.6 times as case managers stopped spending scarce hours on low-risk patients and started spending them where predicted risk justified the cost. Nine in ten high-risk flags now trigger a care-transition handoff, a concordance rate that turned the model into an operational habit. The system estimates \$6.1M in avoided readmission costs annually, and its CMS penalty exposure improved two percentile bands — movement finance can now point to on the same page as the avoided-cost figure.

METRIC	BEFORE	AFTER
30-day readmissions	-	-19%
Care-transition enrollment	-	2.6x
Avoided readmission cost	-	\$6.1M/yr
High-risk flags acted on	-	90%

The qualitative changes compounded the quantitative ones. Reason codes changed the tenor of bedside conversations: a case manager who disagrees with a tier can ask why and receive an answer in clinical terms, which converts opposition into calibration. Overrides became rare and, more usefully, informative — each logged override is either a model blind spot to fix or a clinical factor the model cannot see yet. Pharmacists working the right patients began catching medication problems that previously surfaced as return visits days later. Because drift monitoring was in place from the first month, model health is a reviewed metric rather than an assumption — which is what lets leadership fund it as permanent infrastructure instead of a pilot that quietly decays.

08 What Your Business Should Do

MercyPoint's playbook transfers to any health system, and none of it requires exotic machinery. Start by mining your own readmission history: analyze which patients returned over the last two years before buying any tool, because your data already holds the patterns a model will need to learn from. Build the model with clinicians rather than for them — involve case managers and physicians in feature design from day one, since models built with the ward, for the ward, are the ones that actually get used at the bedside. Surface the scores inside the EHR, on the screens where discharge decisions already happen; a risk tier parked in a separate dashboard dies quietly within a quarter. Demand explainability as a requirement, not a preference — adopt models that produce reason codes clinicians can interrogate and challenge, because an unexplained high-risk flag gets overruled while an explained one changes a discharge plan on the spot.

Then fix the allocation problem the model exists to solve. Target scarce programs at high risk: stop spreading follow-up calls evenly across every discharge and concentrate pharmacists, nurses and social workers where predicted risk justifies the spend, because scarcity is the whole design problem. Track flag-to-action concordance weekly — measure how often high-risk flags actually become a care-transition handoff within seven days, since a technically strong model that clinicians ignore delivers exactly zero prevented readmissions. Monitor drift and retrain on a schedule, so the model flags its own staleness before your outcomes reveal the problem; patient populations shift, and a frozen model slowly stops being true.

Finally, count dollars and percentiles together. Report avoided readmission costs alongside CMS penalty percentile movement every quarter, because boards fund what finance can defend, and those two numbers on one page defend the program better than either alone. Before you commission any of this work, sit with the

questions MercyPoint's team had to answer honestly: what share of your readmitted patients were flagged as high risk before they returned; whether your resources are allocated today by intuition, by rules or by evidence; what a single avoidable readmission is worth to your system once penalties are included; whether your clinicians would trust a risk score and what it would take to earn that trust; where your discharge workflow breaks for patients with complex medication changes at home; and what you would reinvest if readmissions fell 19% next year. Work through the eight actions in this practice book with your clinical and analytics leads, and return to those reflection questions each quarter to keep the program honest.

REAL-WORLD PROOF & BENCHMARKS

The Agency for Healthcare Research and Quality reports that roughly one in five Medicare patients is readmitted within 30 days, costing billions annually. - AHRQ

CMS's Hospital Readmissions Reduction Program can reduce a hospital's Medicare payments by up to 3% for excess readmissions, making prevention financially material. - CMS

Peer-reviewed studies in JAMA and Health Affairs have repeatedly reported that targeted care-transition interventions can measurably reduce 30-day readmissions versus usual care. - JAMA / Health Affairs

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. Mine your own readmission history first

Analyze which patients returned in the last two years before buying anything. Your data already holds the patterns any model will need to learn from.

OWNER

TARGET DATE

STATUS

--	--	--

☐

2. Build the model with clinicians

Involve case managers and physicians in feature design from day one. Models built with the ward, for the ward, are the ones that actually get used at the bedside.

OWNER

TARGET DATE

STATUS

--	--	--

☐

3. Surface scores inside the EHR

Risk scores parked in a separate dashboard die quietly. Put the tier and its reason codes directly where discharge decisions already happen, inside the existing workflow.

OWNER

TARGET DATE

STATUS

--	--	--

☐

4. Demand explainable scores

Adopt models that produce reason codes clinicians can interrogate and challenge. An unexplainable high-risk flag gets overruled; an explained one changes a discharge plan on the spot.

OWNER

TARGET DATE

STATUS

--	--	--

☐

5. Target scarce programs at high risk

Stop spreading follow-up calls evenly across every discharge. Concentrate pharmacists, nurses and social workers where predicted risk justifies the spend — scarcity is the whole design problem.

OWNER

TARGET DATE

STATUS

--	--	--

**6. Track flag-to-action concordance**

Measure how often high-risk flags actually trigger a care-transition handoff within the week. A technically strong model that clinicians ignore delivers exactly zero prevented readmissions.

OWNER

TARGET DATE

STATUS

--	--	--

**7. Monitor drift and retrain on schedule**

Set drift monitoring so the model flags its own staleness before outcomes reveal the problem. Patient populations shift, and a frozen model slowly stops being true.

OWNER

TARGET DATE

STATUS

--	--	--

**8. Count dollars and percentiles**

Report avoided readmission costs alongside CMS penalty percentile movement every quarter. Boards fund what finance can defend, and those two numbers together defend the program well.

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. What share of your readmitted patients were flagged as high risk before they returned?

Q2. How are scarce care-transition resources allocated today — by intuition, by rules, or by evidence?

Q3. What is a single avoidable readmission worth to your system once penalties are included?

Q4. Would your clinicians trust a risk score, and what would it take to earn that trust?

Q5. Where does your discharge workflow break for patients with complex medication changes at home?

Q6. If readmissions fell 19%, what could your system reinvest the savings in next year?

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 MercyPoint Health'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