

TECH & AI CASE STUDY - B2B SAAS - 2026

Usage-Signal Churn Prediction That Protected \$2.3M of ARR

Based on a delivery engagement with SignalStone Software

91% to 96%
NET REVENUE RETENTION

\$2.3M
ARR PROTECTED

14 to 51 days
WARNING LEAD TIME

12% to 34%
WIN-BACK RATE

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 SignalStone Software 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 SignalStone Software - 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

SignalStone Software is a B2B analytics vendor with 140 employees, selling usage-based subscription products to mid-market and enterprise customers across logistics, retail and financial services. Its commercial model depends on expansion: new logos are expensive in a crowded category, so the growth story the board tells investors rests on net revenue retention — existing customers upgrading seats, adding modules and renewing without drama. That story had been weakening. The company ended last year at 91% net revenue retention, a leak its CFO sized at roughly \$2.8M of churned and downgraded ARR, and the board demanded a credible plan to defend the expansion narrative before the next fundraise.

Customer success was the team asked to fix it, and the team was structurally blind. Account reviews ran quarterly, long after usage had decayed, and customer success typically learned an account was at risk when a renewal lapsed or a champion resigned — events that look sudden in a CRM but are usually the last step of a slow decline that was visible in the data for months. Save plays were generic, late and built on gut feel, which meant the CS team's genuine effort was being spent on the wrong accounts at the wrong time.

02 Where Things Stood

The signals of churn existed; they simply never met. Logins and feature adoption lived in the product database, API volume in infrastructure metrics, support tickets in the helpdesk, and invoice history and seat counts in billing — four systems nobody had joined into one account-level view. A customer success manager reviewing an account saw the CRM notes and the renewal date, but not the API volume that had fallen 60% over two quarters, the support ticket with an unresolved escalation, or the feature whose usage quietly stopped after an internal champion changed roles. Each system held a piece of the story, and no human could assemble the pieces across hundreds of accounts every week.

The consequences compounded in both directions. Accounts that looked quiet were actually declining, and accounts that looked noisy — heavy ticket volume — were sometimes the healthiest, because customers who file tickets are customers still trying to succeed. Quarterly account reviews meant the median account got attention months after the moment it mattered, and the save conversations that did happen opened with apologies rather than evidence. Bain & Company research famously found that improving customer retention rates by just 5% can increase profits by 25% to 95% depending on the industry; SignalStone's problem was not arguing with that arithmetic, it was having no operational way to act on it.

03 The Challenge

The challenge was to convert scattered telemetry into a weekly, actionable view of account health that customer success would actually trust and use. The technical half required joining product telemetry, support history, billing data and CRM notes into features a model could learn from — and defining churn precisely enough to train on three years of renewal outcomes. The organizational half was harder: scores only matter if someone changes their behavior because of them, so every flag needed a suggested play, and the scores needed to appear where account executives already work rather than in another portal. McKinsey research reports that companies applying predictive analytics to churn management have cut churn by 10 to 15% in subscription-based businesses, which gave the board a credible ceiling — but reaching even the low end required the model, the workflow and the plays to arrive together.

There was also a trust constraint specific to this company's history. An earlier internal attempt at a health score had failed quietly: a red-amber-green dashboard nobody believed, built on weights two product managers argued

about and no one could explain. CSMs remembered that failure, and several said openly they would not act on another opaque percentile. Whatever Emerson built had to show its reasoning on every flag, or it would join the previous dashboard in the graveyard.

04 Discovery and Diagnosis

Emerson Technology began with a three-week discovery sprint across the four data systems and the CS workflow. The data audit confirmed the join was feasible: account IDs could be reconciled across product, support, billing and CRM, and three years of renewal outcomes provided labels — renewals, downgrades, lapses — clean enough to train on. The team then interviewed every CSM and account executive about how risk actually surfaced today, and the answers converged on the same pattern: accounts were reviewed quarterly, champions' departures were discovered socially rather than systematically, and save plays were chosen by whoever owned the account based on instinct.

The diagnosis surfaced two findings that shaped the design. First, decay preceded churn by months: in account after account, API volume and feature breadth began sliding two to three months before any commercial signal appeared, which meant a weekly scoring cadence could realistically buy the CS team a full quarter of runway. Second, the failed health score had failed for absence of reasons, not absence of accuracy — its flags were usually directionally right, but CSMs could not see why, so they ignored it. PwC's Future of CX survey found roughly one in three consumers will leave a brand they love after a single bad experience; in B2B, the equivalent moment is the unanswered support escalation that confirms a champion's doubts. The model would need to see support friction as clearly as usage.

05 The Solution

Emerson Technology designed and built a weekly churn-risk system on AWS that scores every account, explains every score and pushes both into the CRM. The architecture has three layers.

One account view

Product telemetry, support history, billing data and CRM notes are consolidated into a warehouse on AWS, with dbt models defining account-level health features on a single, agreed account grain. Login frequency, breadth of feature adoption, API volume trends, ticket recency and severity, seat utilization and invoice history all land as comparable signals per account. Defining the features in dbt mattered as much as computing them: each metric has one definition, one owner and one test suite, which ended the era of three teams quoting three different health numbers for the same customer. Churn itself was defined operationally before any modeling — contract lapse, seat collapse or a usage floor — so the model, customer success and finance aligned on the same problem.

Scoring and explanation

An XGBoost classifier trained on three years of renewal outcomes scores every account weekly, paired with a survival model that estimates days to renewal risk, giving CSMs both a probability and a runway. SHAP values generate plain-language risk reasons for each flag — API volume down 40% over eight weeks, two unresolved escalations, seat utilization at 51% — so every score arrives with its evidence attached. Weekly scoring runs on Airflow with drift monitoring, and threshold changes are versioned like code, so no one can quietly move the definition of "at risk" between board meetings. Finance receives a monthly reconciliation of flagged versus realized churn, which keeps the model accountable to outcomes rather than to vibes.

Working where reps live

A Next.js success console shows CSMs ranked queues, trend lines and suggested plays per account, but the system's most important integration is invisible: reverse ETL pushes scores and reasons into the CRM, so

account executives see risk in the tool they already open every morning. Each risk tier carries a concrete play — an executive check-in, targeted training, a renewal restructure — because a score without a next action only spreads anxiety across the CS org. The console exists for depth; the CRM exists for habit.

06 The First 90 Days

Weeks one and two finalized the operational churn definition with finance, product and CS in the room, then built the warehouse joins and backfilled three years of labeled outcomes. Weeks three through six trained and backtested the classifier and survival model, scoring historical accounts blind and comparing flags against what actually happened — the backtest showed the model catching decay a median of roughly seven weeks before the old quarterly-review process would have noticed. Weeks seven through ten piloted the console with five CSMs on the book of business with the highest ARR exposure, logging every disagreement between CSM judgment and model flags for weekly review; the disagreement log caught a billing-feature bug early and, more importantly, gave skeptical CSMs visible evidence the flags earned their confidence.

Weeks eleven and twelve completed the rollout: reverse ETL live for every AE, risk-tier playbooks published, a 45-minute onboarding per CSM, and the first monthly finance reconciliation of flagged versus realized churn delivered to the CFO. The team also agreed on the measurement discipline that would matter most later — tracking saves, downgrades and losses on flagged accounts and pricing each outcome, because flag accuracy flatters models while save economics tell you whether the program is actually paying.

07 The Results

Across four quarters, the retention curve bent. Net revenue retention climbed from 91% to 96%, and the CS organization now flags at-risk accounts a median 51 days before renewal — versus 14 days of warning under the old quarterly review — which converts save plays from last-minute discounts into planned interventions. Win-back on flagged accounts improved from 12% to 34%, and finance attributes \$2.3M of retained ARR to the program in its first year. The CS team plans capacity from the risk queue instead of the renewal calendar, and churn conversations now start with evidence instead of apologies.

METRIC	BEFORE	AFTER
Net revenue retention	91%	96%
ARR protected	-	\$2.3M
Warning lead time	14 days	51 days
Win-back rate	12%	34%

The qualitative changes compounded the quantitative ones. The board deck for the fundraise carried a retention chart with a working machine behind it, not a slide of intentions. Executive check-ins triggered by the model surfaced two expansion conversations that pure account reviews would have missed, because the model flags attention-worthiness, not only danger. The failed red-amber-green dashboard was retired without nostalgia: CSMs had not stopped caring about risk — they had stopped trusting scores without reasons, and the reasons changed that.

08 What Your Business Should Do

The pattern behind SignalStone's result is repeatable, and it begins with plumbing rather than prediction. Join usage, support and billing data into one account-level view, because most churn models fail because the signals live in disconnected systems, not because the math is wrong. Define churn operationally — contract lapse, seat collapse or a usage floor — so the model, customer success and finance argue about the same problem. Score accounts weekly with trend lines, since risk builds daily through usage decay and support friction and weekly scoring catches deteriorating accounts while there is still runway to intervene meaningfully. Explain every risk score with feature attributions such as SHAP values, because success teams trust and act on reasons, never on opaque percentiles.

Attach playbooks to risk tiers — an executive check-in, targeted training or a renewal restructure per band — so a flag always carries a next action. Route scores into your CRM through reverse ETL so reps work where they already live, because tools that demand new logins quietly die within a quarter. Track saves, not just flags: measure whether flagged accounts were saved, downgraded or lost, and price each outcome, since flag accuracy flatters models while save economics tell you whether the program is actually paying. And review model drift quarterly — product changes shift usage patterns and silently degrade scores, so re-validate feature importance and recalibrate thresholds against realized renewals, not against last year's test set.

Ask the questions SignalStone's leadership now reviews quarterly: which ten accounts would hurt most if they left next quarter, and what evidence tells you they are healthy; how many days of warning your team actually gets today before a renewal goes bad; what share of customer success time goes to accounts that were never at risk; which save play your team would run first if churn signals arrived 50 days earlier; whether finance, product and customer success share one definition of churn or three; and what one point of net revenue retention would be worth to your business this year. Work through the eight actions in this practice book with your team, and revisit those reflection questions each quarter to keep the retention engine honest.

REAL-WORLD PROOF & BENCHMARKS

Bain & Company research famously found that improving customer retention rates by just 5% can increase profits by 25% to 95%, depending on the industry. - Bain & Company / HBR

McKinsey research reports that companies applying predictive analytics to churn management have cut churn by 10 to 15% in subscription-based businesses. - McKinsey & Company

PwC's Future of CX survey found roughly one in three consumers will leave a brand they love after a single bad experience. - PwC, 2018

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. Join usage, support and billing data

Combine product telemetry, support tickets and invoices into one account-level view. Most churn models fail because the signals live in disconnected systems, not because the math is wrong.

OWNER

TARGET DATE

STATUS

--	--	--

☐

2. Define churn operationally

Decide exactly what counts as churn: contract lapse, seat collapse or a usage floor. One precise definition keeps the model, customer success and finance aligned on the same problem.

OWNER

TARGET DATE

STATUS

--	--	--

☐

3. Score accounts weekly

Risk builds daily through usage decay and support friction. Weekly scoring with trend lines catches deteriorating accounts while there is still enough runway to intervene meaningfully.

OWNER

TARGET DATE

STATUS

--	--	--

☐

4. Explain every risk score

Use feature attributions such as SHAP values to show why an account was flagged. Success teams trust and act on reasons, never on opaque percentiles.

OWNER

TARGET DATE

STATUS

--	--	--

☐

5. Attach playbooks to risk tiers

Pair each risk band with a concrete play: an executive check-in, targeted training or a renewal restructure. A score without a next action only spreads anxiety across the CS org.

OWNER

TARGET DATE

STATUS

--	--	--



6. Route scores into your CRM

Push health scores and risk reasons into the CRM through reverse ETL so reps work where they already live. Tools that demand new logins quietly die within a quarter.

OWNER

TARGET DATE

STATUS

--	--	--



7. Track saves, not just flags

Measure whether flagged accounts were saved, downgraded or lost, and price each outcome. Flag accuracy flatters models; save economics tell you whether the program is actually paying.

OWNER

TARGET DATE

STATUS

--	--	--



8. Review model drift quarterly

Product changes shift usage patterns and silently degrade scores. Re-validate feature importance and recalibrate thresholds every quarter against realized renewals, not against last year's test set.

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. Which ten accounts would hurt most if they left next quarter, and what evidence tells you they are healthy?

Q2. How many days of warning does your team actually get today before a renewal goes bad?

Q3. What share of your customer success time goes to accounts that were never at risk?

Q4. If churn signals arrived 50 days earlier, which save play would your team run first?

Q5. Do finance, product and customer success share one definition of churn, or three different ones?

Q6. What would one point of net revenue retention be worth to your business this 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 SignalStone Software'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