balboa. × pendo

A JOINT PERSPECTIVE · FIELD GUIDE

Predictive Customer Success

Building the Data Foundation for an Agentic CS Motion

AUTHORED BY

Jonathan Huhn, Chief of Staff, Balboa
Jay Nathan, CEO, Balboa
Jenna McLaughlin, Sr. Director, Pendo Predict, Pendo

PUBLISHED

July 2026

balboa.

44% of CSMs miss target.

When nearly half the role falls short, the problem lives in the architecture they are asked to work within.

REPVUE · JUNE 2026

The signals that tell you whether a customer will renew, expand, or churn already exist inside your business. They sit scattered across product analytics, the CRM, support systems, and billing records. Today the work of turning them into a decision falls to individual CSMs working on instinct, and nearly half miss target.

A better system is now within reach for any company with a product and a customer base. A predictive model ingests every signal, learns from your own history of churn and renewal, and produces one continuously updated answer for each account. That answer lands in the workflow your team already uses, explained in plain language, paired with a recommended next step, and acted on for you.

This is the move from **reactive**, **dashboard-driven** customer success to **predictive**, **agentic** customer success. The fuel is behavioral data you already collect in product analytics platforms like Pendo, Amplitude, Mixpanel, or Heap. This paper walks through which data sources to use, how to make the data model-ready, how the model works, how predictions become action, and who you need to run it.

You do not need a data science team, a platform migration, or a year of build time. **You need** the data you already have, the right architecture, and a few strong people who know what they are doing.

Why the current state of CS is broken

To understand the predictive CS leap, it helps to lay out why the current state underdelivers. More often than not, the problem is the architecture.

The CRM was out of its depth. First off, customer relationship management tools like Salesforce were never built for analytics. A CRM is a transactional system (effectively a sub-ledger to the general ledger), designed around opportunities, bookings, sales compensation, and pipeline management. It is very good at recording what happened to a deal, but it is poor at telling you what will happen next with a customer. Asking it to power forward-looking CS decisions is like asking a thermometer to tell you the time.

Customer success platforms tried to fill that gap, and created a new problem. Tools like Gainsight, ChurnZero, and Planhat layered an analytical view on top of customer data, and then bolted on adjacent capabilities, e.g. an NPS tool here, a timeline tool there, and a health score somewhere else. Since these capabilities never belonged outside the CRM, it often results in a sort of “shadow CRM”: two system-of-record databases trying to maintain two versions of the truth via a tenuous two-way sync, and a permanent tax on data integrity. Of course, the CSP did buy the CS team some much-needed autonomy. It was a system they controlled instead of always waiting in line behind every big sales initiative for attention from RevOps. But autonomy came at the cost of fractured data.

Meanwhile, the clearest signal was siloed away. Product analytics data (i.e. how customers use the product) is the most direct evidence of whether customers are getting value. It tends to be a firehose, pouring out so much behavioral data that the trick is to determine which signals matter. It has also historically lived in the product org, and the CS teams who retain and grow customer accounts have rarely had real access to it. As such, the data most predictive of churn became the least available to the people responsible for preventing it.

So the wrong kind of work fell on humans. Lacking a system that synthesizes these signals, CSMs do it in their heads. They glance at a health score, a last-login date, an open ticket count, a renewal date, and they connect the dots by intuition. Some are extraordinary at this, but for most, manual synthesis of fragmented data does not scale, no matter how skilled the person. According to a June 2026 report from RepVue, nearly half (44%) of all CSMs in the United States miss their annual target. A number that high indicates the role itself is broken.

This is the one-dimensional analytics problem. Every CSP gives you attributes, one at a time, and then asks a human brain to combine them into 1.) a judgment and 2.) an action. That does not fit a growing, multi-product business with complex customer journeys. The fix is a different kind of system entirely.

Predictive + agentic customer success

First, the predictive half. Instead of a dozen attributes for a CSM to weigh, consider a single predictive model that takes in every signal you have (product usage, CRM history, support activity, billing patterns) and produces one continuously updated score for each account, with the reasoning behind it. The model does the data synthesis while the human does the judgment and the work.

In simplest terms, this model is like having an experienced CSM who has reviewed every account your company has ever lost and can now look at any current account and tell you whether it is on the same path. But this "CSM" doesn't forget anything important, it updates its read every night, and it goes off data science instead of intuition. That is the predictive component of the new CS motion.

Throughout this paper, we examine Pendo Predict as the predictive layer for this motion. We do so for a few reasons: 1.) it is purpose-built for the customer success use case rather than a general machine learning tool adapted for it, 2.) it is source-agnostic, meaning it does not require Pendo analytics and can ingest data from whatever stack you already run, and 3.) it automates the underlying data science infrastructure (ETL, normalization, feature extraction, model training) that would otherwise require a dedicated data science team to build and maintain.

The agentic half comes next. A score that lives inside a modeling tool is just a smarter dashboard. The real value comes when the prediction lands where the team already works: in the CRM record, explained in the company's own taxonomy, paired with a recommended next step, and capable of triggering an action on its own. Increasingly, the model and its underlying data become a queryable intelligence layer that an internal CS agent can sit on top of: forecasting outcomes, benchmarking performance, and answering "where is this account going?" on demand.

Combining the predictive and agentic components of the new CS motion, it becomes the analytical intelligence layer your CRM was always missing, living natively in the system your team already uses, without the overhead of a separate shadow platform.

> The CRM stays the system of record, and the predictive model becomes the brain on top of it.

HOW ACCURATE CAN IT GET?

You do not need a data science background to trust the model, only honest expectations. A good model catches far more at-risk accounts than any manual read, early enough to act.

65%
Industry baseline for models of this type

70%+
Target for a well-configured model

85–90%
Strong, mature implementations

SECTION 03 · THE DATA FOUNDATION

What goes into the model

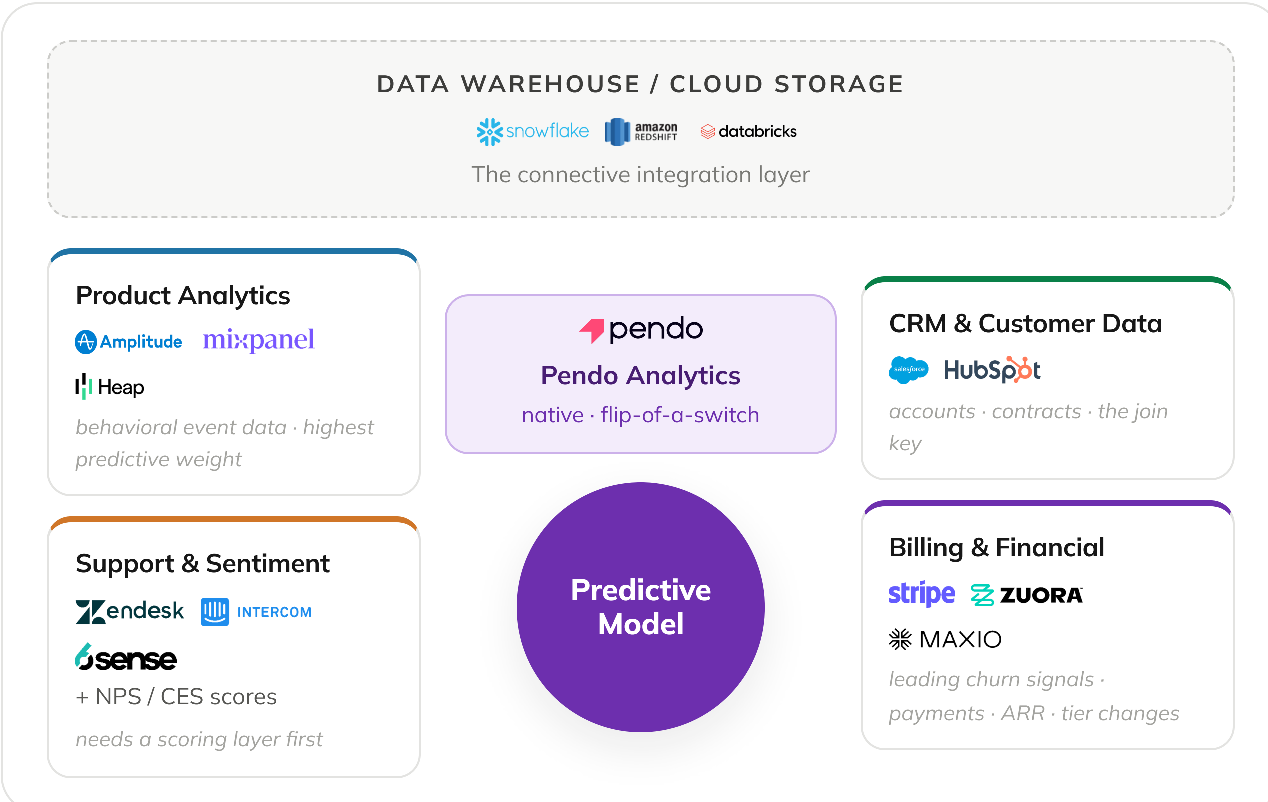

Figure 1. The data sources that feed a predictive model, with the warehouse as the common integration layer. At scale, most external sources route through the warehouse rather than connecting directly.

SOURCE	REPRESENTATIVE TOOLS	PRIMARY INTEGRATION PATH	PREDICTIVE WEIGHT
Product analytics	Pendo, Amplitude, Mixpanel, Heap	Native (Pendo) or via the warehouse	Highest
CRM & customer data	Salesforce, HubSpot	System-level account connection	Core
Support, sentiment & intent	Zendesk, NPS / CES, 6sense	Scoring / sentiment layer first	Supporting
Billing & financial	Payments, ARR, tier changes	Warehouse or billing system	Supporting
Data warehouse	Snowflake, Redshift, BigQuery, Athena	Native connectors; S3 / Azure fallback	Integration hub

Reading the sources, one by one

A predictive model is only as good as the data it learns from. Most companies already have the data they need, but it's scattered across a handful of systems. Here is how the sources stack up, starting with the behavioral data that carries the most predictive weight.

pendo Amplitude mixpanel Heap

Product analytics

Product usage is the most direct evidence that a customer is getting value, so it carries the highest predictive weight. Pendo is the flip-of-a-switch option, with years of tagged event data ready for Predict natively. Mixpanel connects natively too, though its instrumented event model means completeness depends on how thoroughly the team set it up. Amplitude enforces a stricter schema that produces cleaner data at scale and routes in through the warehouse. Heap autocaptures everything from install, trading governance for coverage, and also comes in cleanest through the warehouse.

salesforce HubSpot

CRM and customer data

The CRM is the structured backbone. It holds account records, contract history, and the join key that links every other source to a customer identity. Product usage tells you how engaged a customer is; the CRM tells you who they are and what they pay, and the two together drive most of a model's predictive power. Connect through a system-level account so the link survives role changes and permissions stay scoped.

zendesk INTERCOM 6sense

Support, sentiment, and external signal

Support tickets, NPS and CES scores, and product feedback all carry health signal, but raw text cannot feed a model. It needs a scoring or sentiment layer that turns it into a clean numeric feature first. And every added degree of aggregation moves the data further from its raw form, so its output earns more skepticism.

stripe ZUORA MAXIO

Billing and financial data

Payment history holds some of the most direct health signals available. Late payments, downgrades, and tier changes often lead churn by weeks. A field earns its place as a predictor only if it is dense and consistent across your history, no matter how meaningful it seems.

snowflake amazon REDSHIFT databricks

The warehouse as integration hub

For companies with a data warehouse, this is the best single integration point, a normalized view across every source above that is usually cleaner than a direct API extract. Pendo Predict connects to every major warehouse natively, and cloud storage buckets serve as a fallback. Connecting each source directly is the fastest path to proof of value, but routing everything through the warehouse is the more durable design as the program matures.

Making the data model-ready

Getting data to a model-ready state is about 90% of the work, most of it cleaning data and reconciling sources. The CRM is the messiest place to start.

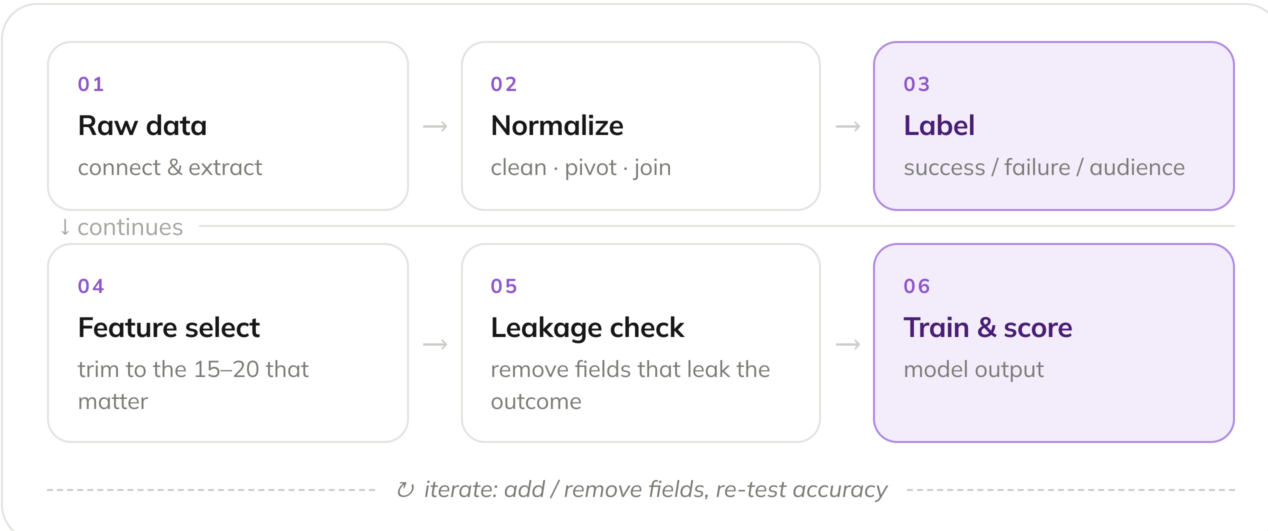

Figure 2. The iterative path from raw, scattered data to a trained, scored model.

Labeling comes first. Before a model can learn, every historical record needs one of three labels. Success is the outcome you predict, like a closed-lost renewal. Failure is its opposite, a renewal that happened. Audience is the records the model will score. Getting these definitions right is a human job that surfaces disagreements a company never resolved, like whether a downgrade or product migration counts as churn.

Data leakage is a common, sneaky failure. Fields that reflect an outcome that already happened, like account status, must be excluded, because they will not exist when the model runs in production. Leave them in and the model looks brilliant in testing but falls apart in the field.

Over-influential fields are also dangerous. When one field looks twenty points better than everything else, it is almost always a data artifact rather than a real predictor.

Predictor selection is ongoing. A model might start with several hundred candidate columns and get trimmed to the 15 to 20 that carry signal. Pendo Predict automates the first pass and recommends which to keep, but a practitioner makes the final calls, narrowing from there by hand.

Volume and seasonality round it out. Plan on 1.5 to 2 times as much history as the window you forecast (e.g. a three-month forecast wants about half a year of data). A technique called jitter keeps the model from misreading seasonal swings like heavy month-end usage.

SECTION 06 · THE AGENTIC LAYER

From prediction to action

A prediction that sits inside a modeling platform changes nothing. The model only becomes a customer success motion when its output lands in the workflow where the team already lives.

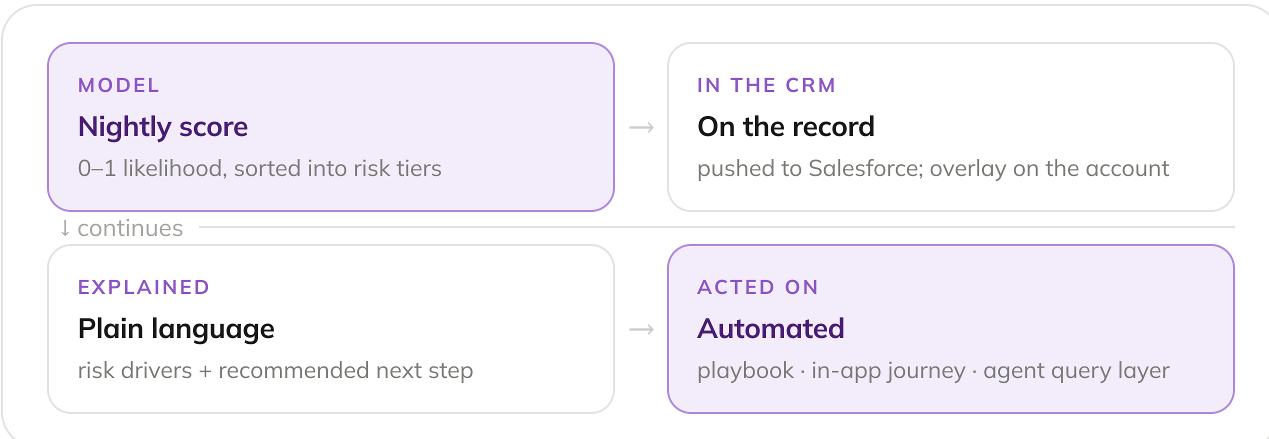

Figure 3. A prediction is only useful when it reaches the workflow and turns into an action. Synthesis moves from the human to the system.

It starts in the CRM. Health scores push into Salesforce as a custom field on a nightly basis, and a lightweight overlay surfaces the prediction directly on the opportunity or account record. There's no separate dashboard to log into, and no report to hunt down. When the CSM opens the record, the score is already there.

Then it gets explained. An AI layer translates the model's output into plain language. It surfaces what is driving the risk, phrased in the company's own terminology and combined with the relevant data from the CRM and product telemetry. Alongside the explanation sits a recommended next step, and optionally a playbook mapped to the specific risk profile and account. The CSM does not have to guess what to do, because the system surfaces it.

Then it gets acted on. Scores also become a segmentation input back in the product, so a risk tier can trigger an in-app message, an email journey, or another automated intervention. For customers without an assigned CSM, this enables autonomous engagement. And for those with one, it complements the human outreach.

And increasingly, it can be queried. The fuller expression of the agentic layer is an intelligence layer built on the combined model and data, which can forecast outcomes, benchmark the model against market norms, and dissect predictions to answer the "why." This is the infrastructure for an internal CS agent, and it operates continuously, surfaces the right information at the right moment, and lifts the burden of synthesis off the human. That is the destination this whole architecture is built toward.

The team you need

Technology is only the first half of predictive customer success. The companies that win here staff the motion intentionally, and ownership shifts between two phases.

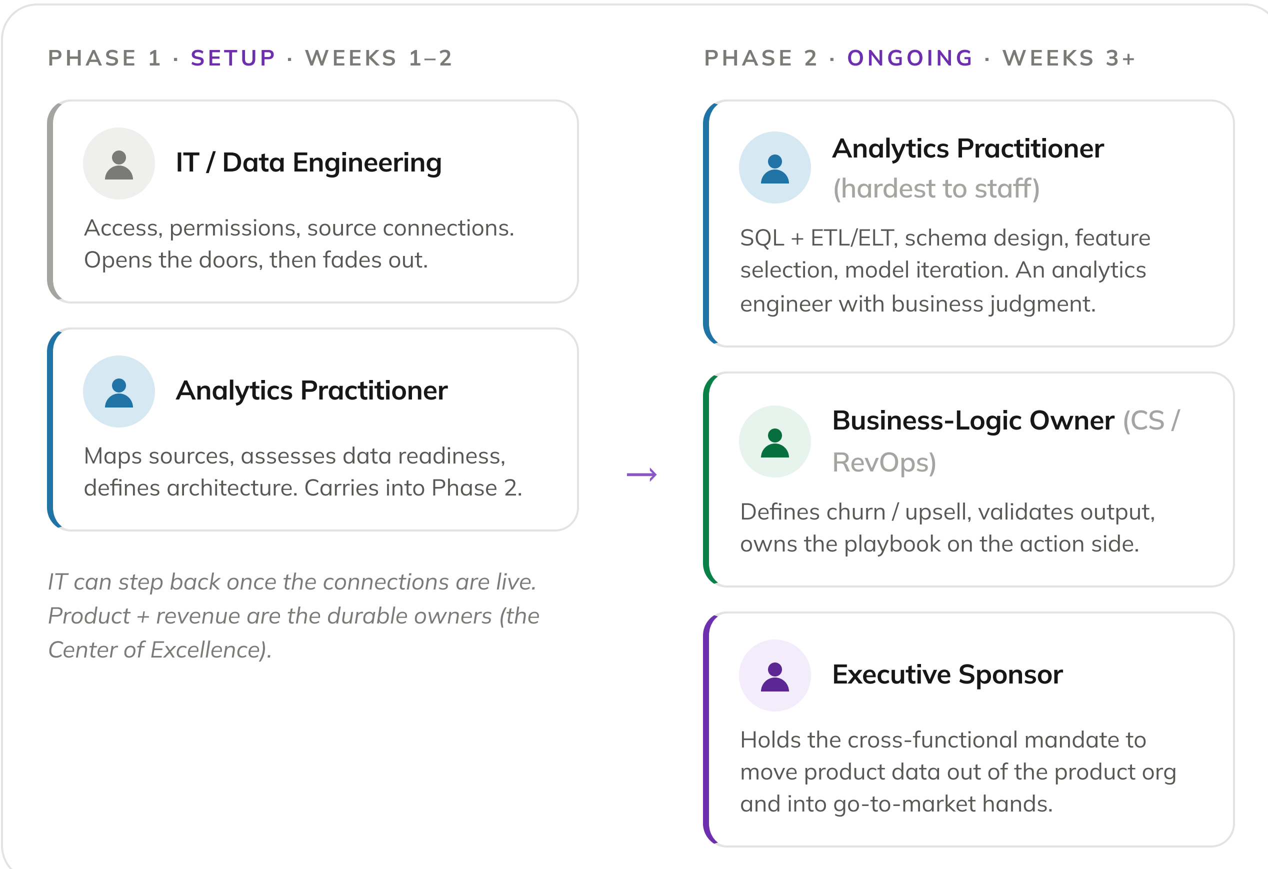

Figure 4. The roles required, and how ownership shifts from setup to ongoing operation.

Setup runs a couple of weeks. IT opens access, permissions, and source connections, then steps back, while an analytics practitioner maps the sources and judges whether the data has enough history and density to model.

Then the practitioner owns the data day to day, from SQL and pipelines to feature selection and model iteration, and they are the hardest role to staff. The business-logic owner in CS or RevOps supplies the definitions the model cannot derive and owns the playbook. The executive sponsor holds the mandate to move product data into go-to-market, and over time product and revenue own the motion together while IT steps back.

What this makes possible

The point of all this architecture is revenue expansion and cost reduction. A working predictive, agentic motion delivers both, and the shift below shows how.

	MANUAL / CSP-BASED CS	PREDICTIVE + AGENTIC CS
How risk is found	CSM reads scattered attributes, synthesizes by gut	Model synthesizes every signal into one score
When you know	Weeks after behavior shifts; often too late	Updated nightly; hours after behavior shifts
Coverage	~20% of CSMs catch most of the risk	The whole team operates on the same signal
Where it lives	A separate dashboard or report to go find	In the CRM, on the record, with next steps
What product data does	Sits in the product org, unused by GTM	Becomes a go-to-market asset

Figure 5. The shift from a manual, dashboard-driven motion to a predictive, agentic one.

The five outcomes below are modeled on a $50M ARR example business with 500 customers and $100K average ACV.

Gross revenue churn ↓

If 20% of 500 accounts show risk each year and you save 1 in 4 today, that is 15% gross churn, or $7.5M. Flip that to 3 in 4 and churn drops to 5%, saving about $5M.

Net revenue retention ↑

Expansion runs on the same engine, flagging accounts that look like customers who grow. Acting early can lift net revenue retention from 105% to 115%, adding $7.5M of net ARR in a year instead of $2.5M.

Accounts covered per CSM ↑

The model does the synthesis, so CSM effectiveness scales. At 25 accounts per CSM you need 20 people to cover 500; double that to 50 and you need 10, the same coverage with 10 fewer hires.

Time to insight ↓

A model that updates nightly replaces the quarterly, hand-built scoring exercise. A behavioral shift that once took until month three to surface now shows up within 24 hours.

Return on existing analytics spend ↑

Years of product analytics data already sit underutilized in a warehouse. At $50M ARR, a company might spend $250K on analytics and $150K on a standalone CSP. Predict delivers that inside the CRM, folding $400K of overlapping spend into one system.

SECTION 09 · GETTING STARTED

Where to begin

Most teams jump straight to "which model do we build?" The better starting point is the data. What do we have, where does it live, and does it have enough history to learn from? A data readiness assessment, mapping the sources, checking coverage, and identifying the join keys across systems, is the right first move for any company, whatever its Pendo journey or analytics stack.

Keep the first engagement narrow. One product, one use case, and one outcome. We recommend churn as the starting point, since its definition is the most consistent across companies and the record of closed-lost renewals gives the model clean labels. The first model does not need to be perfect to pay off. A model that catches any at-risk account you would otherwise miss returns value from day one. The data you need is already yours; a working model is now just a matter of data prep and a few of the right people.

ABOUT THE AUTHORS

Our contributors

Jonathan Huhn

CHIEF OF STAFF · BALBOA

Chief of Staff at Balboa, working across strategy and delivery to turn product analytics into repeatable, real-world customer outcomes.

jonathan.huhn@balboa.solutions

Jay Nathan

CEO · BALBOA

CEO of Balboa and a longtime customer success and go-to-market leader, focused on how companies modernize the way they retain and grow customers.

jay.nathan@balboa.solutions

Jenna McLaughlin

SR. DIRECTOR · PENDO

Senior Director at Pendo, leading the team bringing Pendo's predictive analytics and engagement intelligence offerings to market.

jenna.mclaughlin@pendo.io

If you want to launch or optimize your predictive customer success motion, **reach out to any of the contributors here.** We are happy to do an informal discovery session to learn your goals and help you plan your next steps.

balboa.