



Intent Signal Delivery

How to build it on the stack you
already own.

CONTENTS

- 01 What a signal actually is, and where yours are sitting
- 02 The eight signals
- 03 Resolution
- 04 How we wire it into your systems
- 05 How it stays honest
- 06 What lands in Slack
- 07 How we run the engagement

What a signal actually is, and where yours are sitting.

A signal is any observable event that moves the odds a company buys in the near term. A pricing page opened three times in a week. A seat count crossing a threshold. A competitor named on a call. An account you lost last year reading your docs again.

Most revenue teams already collect all four. The data is present, licensed, and paid for. It sits in the intent platform, in the CRM, in the conversation tool, in marketing automation, in whatever captures product events, and each of those systems is complete and internally consistent on its own terms.

The gap is what happens next. Nobody has built the layer that reads them together, so the job lands on the rep. To work out whether an account deserves a call this morning, a rep opens ZoomInfo, then Salesforce, then Gong, then the marketing timeline, then Slack, and holds all of it in their head long enough to reach a judgment. Then they do it again for the next account, and the one after that.

That is an integration task. We are asking a salesperson to perform it by hand, dozens of times, before their first call of the day. What they build instead is a shorter list from memory, which is the rational response to being handed an impossible amount of reading. Reps are good at this work. Doing it manually is simply slower than the market moves.

So the proposition is narrow. We become the integration layer. One system reads every source you own, resolves them to the same company and the same person, weighs them against each other, and hands a rep a short list with the reason already attached.



EVERYTHING NEEDED TO MAKE THE MEASUREMENT, AND NOTHING JOINING IT UP


The eight signals.

Sources are not interchangeable inputs to one score. Two properties separate them, and both are measurable before you weight anything.

Refresh latency sets the ceiling on how useful a source can be. A weekly file describes a world that was true up to six days ago. **Resolution** sets what a rep can do on receipt. A signal that names a person can be acted on now. A signal that names only a company needs a research step first, and research steps are where lists quietly die.

Plotted against each other, the eight sources most teams own separate cleanly, and the separation is usually the opposite of the spending. Third party intent, the source most often bought first, is both the slowest and the least precise. Site behavior, which almost nobody routes anywhere, is instant and names a person.

One pattern sits underneath all of it. Two weak signals inside a short window generally beat one strong signal on its own, because coincidence in time is itself evidence.



● Site behavior Your own analytics, reverse IP	Instant Person	● Open pipeline Salesforce, HubSpot	Instant Account
● Marketing engagement HubSpot, Marketo, Pardot	Instant Person	● Closed lost history Salesforce, HubSpot	Instant Account
● Conversation data Gong, Chorus, Clari Copilot	Post call Person	● Review site activity G2, Capterra, TrustRadius	Weekly Inferred
● Product usage Pendo, Amplitude, your own events	Daily Person	● Third party intent ZoomInfo, Bombora, G2, Clay	Weekly Inferred

ARTIFACT 1

Eight sources by refresh latency and resolution. The useful quadrant is the top left, where a signal arrives immediately and already names a person. Third party intent, typically the largest line item of the eight, sits in the opposite corner.


Resolution.

It reads like a data hygiene footnote. It is where most of these projects are won or lost, which is why it gets its own section here.

Suppose one company exists in your systems four times. Northwind Freight Inc. in the CRM, Northwind Freight in marketing automation, NW Freight LLC in billing, and a bare domain in the intent feed. A scoring engine computes four partial scores. Each one sits under the threshold. The account that should have surfaced never does, and nobody can point to a failure, because every component did exactly what it was built to do.

Resolution is the layer that collapses those four into one, at the company level and at the person level, and then keeps them collapsed while people change jobs and companies get acquired. Doing it well requires judgment calls specific to your business. Whether subsidiaries roll up to the parent. How much of a name match is enough. What you are willing to do with a domain that has no person attached to it.

Those are business decisions, so we make them with you before anything gets scored. A plain model on resolved records outperforms a sophisticated one on split records, and the margin is not close.



ARTIFACT 2

Four partial scores, none of them over the line, describing one company that should have been called weeks ago. The same problem repeats at the contact level every time somebody changes jobs.

How we wire it into your systems.

Four layers, each one replaceable without touching the others. That separation is what keeps the system alive when you swap a vendor, and swapping a vendor is a question of when.


Ingest. We connect to what you already license and take events as they happen rather than on a schedule somebody else chose. Where a tool publishes webhooks we use them. Where it does not, we poll at the fastest interval your contract permits. Everything that arrives is normalized into one event shape, so nothing downstream has to know or care which vendor produced a given fact.

Resolve. Every incoming event gets attached to a resolved company and, where the evidence supports it, a resolved person, following the rules agreed in the previous section. Events that cannot be resolved yet are held rather than discarded, because a great many of them resolve later when a form fill or an email click supplies the missing identity.

Score. Weights are applied per signal type and aged on a decay curve, so a page view from three weeks ago counts for less than one from this morning. The result is a single number per account. Every score keeps its own inputs, which means any number on any list can be explained back to the events that produced it, months later, to a skeptical rep.

Deliver. Thresholds are set per segment, because a hot enterprise account and a hot mid-market account look nothing alike. Anything that crosses gets routed to the owner with the evidence attached.

All of it runs on infrastructure we stand up inside accounts in your name, and hand over. There are no new seats for your team to learn and no license that has to be renewed to keep the thing running.



ARTIFACT 3

The four layers, and the loop on the right that keeps the third one honest. Sources enter at the top without any of them being a dependency. Losing one degrades the score rather than stopping the system.

How it stays honest.

Every scoring model begins as a set of guesses about what matters. Ours does too. What separates one from another is what happens to those guesses after launch.

Once a quarter we close the loop. We take the scores the system produced, look at which of those accounts a rep actually contacted, look at what came of it, and set that against the accounts the system scored low and nobody touched. Then we move the weights to match what the outcomes say. Every change is versioned, so a weight set that performs worse than the one it replaced can be put back the same afternoon.

That last part carries more weight than it sounds like it should. Most scoring models running in production have never been checked against a result. They were configured once, by somebody who has since left, and they have been quietly wrong ever since while everybody assumed the number meant something.

The loop also catches signals that stop working. A source that predicted well in the first quarter can predict nothing by the third, because a vendor changed methodology or your market moved underneath it. The back-test surfaces that inside a quarter, well before a rep has spent a year learning to ignore the alerts.

We run this for as long as we are engaged, and we teach whoever inherits the system to run it after we are gone.



ARTIFACT 4

The back-test. The dashed return is the only step most scoring models never take, and it is the one that turns a set of opening guesses into something you can defend to a board.

What lands in Slack.

Everything above exists in order to produce this.

A digest rather than an alert stream. Alerts train people to dismiss them inside a fortnight. A fixed morning list gets read, because it arrives at the same time, in the same place, and it ends.

Three to seven accounts, ordered by score. Each one carries the number, the signals behind it, a plain sentence describing what actually happened, and a recommended next action. It lands in the channel or the direct message a rep already lives in, so nothing new has to be opened and no new habit has to be formed.

None of this is technically hard, and it is the piece clients push back on most, because it looks too plain to be the output of a serious build. It is also the only part a rep will ever see. Get it wrong and the other three layers may as well not exist.

Northwind Freight 97 ①

● Third party intent ● Open pipeline ②

Researching route optimization for eleven days.
Two open opportunities, both stalled at
evaluation since June. ③

Re-engage the open opportunity ④

Bellamy Health 93

● Closed lost ● Site behavior

Renewed reading on the original evaluation
topic. The champion who ran that deal is still in
seat.

Reopen with the prior champion

Corvid Analytics 89

● Product usage ● Conversation

Seat utilization past threshold three weeks
running. Expansion raised on the last two calls.

Open the expansion conversation

Halden Materials 84

● Site behavior ● Marketing

Pricing page four times in six days from two
people, one of them new to the account.

Route to the account owner today

READING A ROW

- ① **The score**
One number per account,
recalculated the moment anything
moves.
- ② **The signals behind it**
Which sources fired, so a rep can
weigh the evidence themselves rather
than trusting a number.
- ③ **A plain reading**
What happened, in a sentence,
written for somebody who has thirty
seconds.
- ④ **The next action**
The single recommended move. Reps
override it often, and that override is
itself an input to the back-test.

ARTIFACT 5

A digest as a rep receives it. Companies invented. The format is deliberately boring, because the goal is a decision made in ten seconds and no reason to open a second tab.

How we run the engagement.

Eight to twelve weeks, depending on how many sources are in scope and how much resolution work the records need. We work in the open the whole way through, and you can see the system running from the third week rather than at a reveal in the twelfth.

What we need from you is access and decisions, in that order. Access early, so that week three is not spent waiting on an admin who is on leave. Decisions when they arise, mostly on resolution rules and thresholds, because those are business calls and we will not quietly guess at them on your behalf.

What you receive is the running system, the documentation, the full weight history, and working sessions with whoever will own it afterward. Everything sits in accounts in your name, on infrastructure you control.

Many engagements continue on a retainer, to run the quarterly back-test and adjust when the stack changes. Plenty do not, and that is a perfectly good outcome. Nothing stops working when we leave.

WEEKS 1 TO 2	WEEKS 3 TO 4	WEEKS 5 TO 7	WEEKS 8 TO 9	WEEKS 10 TO 12
Access and inventory	Resolution	Scoring	Delivery	Handover
WE NEED Admin access to each source, and an hour with a rep who builds their own list.	WE NEED Decisions on subsidiary rollup and how tight a name match has to be.	WE NEED Your read on which of last year's deals were the ones worth having.	WE NEED Channel choice, segment thresholds, routing rules, and a pilot group.	WE NEED Whoever will own this, for two working sessions.
YOU RECEIVE A written inventory of what each system holds and what it can actually emit.	YOU RECEIVE A resolved account and contact set, with every duplicate it collapsed listed out.	YOU RECEIVE An opening weight set, scored back against the last four quarters of closed business.	YOU RECEIVE The digest running live to the pilot group every morning.	YOU RECEIVE Documentation, the weight history, and the first back-test run together.


Cadence. A standing thirty minute call, twice a week.

Fridays. A written note: what moved, what did not, what we need.

Never. A status deck.

ARTIFACT 6

The engagement, with what it asks of you at each stage. The two weeks of inventory before anything is built are the reason the later stages hold their dates.



If any of this is worth a conversation, we are easy to reach.

Thirty minutes, no charge, no deck. Tell us what you already own and what your reps do at eight in the morning. We will tell you on the call whether we would build this for you, or whether something else in your stack ought to be fixed first.

hello@sapwood.io

SAPWOOD

INTENT SIGNAL DELIVERY

SEPTEMBER 2026

DENVER, COLORADO