

Signal Desk

A Signal-Based ABM Operating Model

Account tiering and campaign orchestration for B2B teams that need to prioritize without a third-party intent data platform.

Diana Pantchev · dianapantchev.com · 2026

1. The problem
2. Engine loop
3. Scoring lanes
4. Tier logic
5. Suppression rules
6. Campaign cohorts
7. Implementation options
8. Prototype screenshots
9. Final takeaway

1. THE PROBLEM

ABM without intent data requires a different instrument.

Enterprise ABM programs are built on a premise: that signals purchased from a third-party co-op can tell you which accounts are actively researching your category. The leading platforms cost \$50,000–\$200,000 per year. When the budget disappears, the list disappears with it.

Buyers of corporate social impact platforms announce their triggers. A new ESG leader is a public hire. A foundation is a public filing. An impact report is a press release. The hard part was never the data. It was knowing what to look for and what it is worth.

No shared prioritization	The account list is whoever sales happens to be working.
Field knowledge trapped	A champion departure is invisible to marketing until the deal is cold.
ABM spend by enthusiasm	The most internally-advocated accounts are not the closest to buying.
Nothing demotes	Almost every ABM tool promotes accounts. None demote them automatically.
Open deals inflate pipeline	Spend on accounts sales is actively working inflates pipeline numbers.

2. THE ENGINE LOOP

The campaign produces the next cycle's evidence.

Signals flow into a score. The score determines a tier. The tier determines a play. The play produces engagement which lands back as CRM behaviour. Every tier change is logged with a date and a cause.



The weights are a hypothesis until two quarters of closed-won data are available to re-fit them. The quarterly refit is when this stops being a framework and becomes a model.

Each lane has its own decay rate. Each cap reflects what it can honestly carry.

Lane 1: Public signals — no cap, replaced each weekly sweep

Signal	Wt	Why
New social impact or ESG leader hired	+25	New leader has a mandate to show progress inside 12 months.
Corporate foundation launched or restructured	+25	Creates a grants management requirement spreadsheets cannot carry.
Incumbent platform named publicly	+22	Displacement window opens near renewal.
Merger or acquisition	+20	Two giving programs, two contracts, one survivor.
First impact or ESG report published	+20	Public reporting commitment. Manual collection breaks in year two.
Grants program scaling or RFP activity	+20	Volume has outgrown the process.
B Corp certification pursued or achieved	+18	Certification requires evidence. Evidence requires tooling.
Public giving or volunteering commitment	+18	A number said out loud creates an obligation to measure it.
Rapid headcount growth	+15	A manual program that worked at 500 fails at 2,000.
Disaster or crisis response moment	+15	Urgent matching gift need. Short window.
ERG or DEI investment announced	+12	Adjacent budget holder, warm route to the people team.
New market or office expansion	+10	Country-specific programs need to go global.
Layoffs or restructuring	-20	Discretionary programs freeze first.
Recently signed with a competitor	-25	Locked for the term. Long-cycle nurture only.
Social impact leader departed, no replacement	-15	No internal owner means no purchase.

Negative signals are first-class. An account that renewed with a competitor scores below one with no signals at all.

Lane 2: Field verdicts — full at 30 days, zero at 180 days

Signal	Wt	Why
In an active evaluation	+35	The strongest signal available. No public source will ever show this.
Budget confirmed	+25	Moves the account from possible to timed.

Named champion engaged	+20	Someone inside is doing work on your behalf.
Existing relationship at executive level	+15	Warm entry no data provider can see.
Asked about us unprompted	+12	Intent expressed directly rather than inferred.
No traction after sustained effort	-15	Effort here has a poor record. Reallocate.
Champion left or went quiet	-20	The internal engine stopped.
Renewed with a competitor	-40	Confirmed lock. Overrides the public sources.

Nothing is logged without human approval. Reps review extracted verdicts before they affect the score.

Lane 3: CRM behaviour — full at 30 days, zero at 90 days (Phase 2)

Signal	Wt	Why
Inbound contact or demo request	+25	Unsolicited. Full weight.
Two or more buying committee roles engaged in 90 days	+18	Breadth predicts a real evaluation better than depth with one contact.
Pricing, security, or RFP page viewed	+15	Late-stage reading. Someone is building a case internally.
Meeting held and logged	+12	Calendar time is the scarcest thing a buyer spends.
Gated form submitted in 90 days	+10	Traded contact details for content.
3+ email clicks in 60 days	+8	Solicited engagement. Half weight by design.
No activity in 180 days while tiered high	-12	Tiered spend with nothing coming back.
Account-level unsubscribe or hard bounce	-10	The channel is closing.

Cap: +40 / -20. Campaign-driven activity at half weight. Unsolicited at full weight. Without that split the score measures your own marketing spend, not buyer intent.

Editable signal weights

Signal DeskAccount tiering for corporate social impact

The book (8)Find accountsMeeting notesThe modelCadence and data

Weights are a hypothesis, not a finding. Change one and the book re-ranks immediately, so a rep who disagrees can argue with a number instead of distrusting the list.

Reset to defaults

Activation thresholds

Tier	Score at or above	Play
1:1	70	Named account plan. Personalized page referencing the specific trigger. Executive direct mail. AE and BDR sequenced behind the physical touch.
1:few	40	Cohort by trigger type, not industry. Shared narrative, personalized opener. Roundtable as the conversion moment.
1:many	15	Programmatic coverage. Nurture on the trigger theme. Re-score monthly.

Decay

Field verdicts hold full weight for	Field verdicts reach zero at
30	180
CRM behaviour holds full weight for	CRM behaviour reaches zero at
30	90

Days. Behaviour goes stale faster than a confirmed budget, so it decays faster. A budget confirmed in March is not worth what it was in March, and every CRM treats it as permanently true.

Public signals

Signal Desk scoring model. Weights are editable in the interface. Change one weight and the entire account book re-ranks immediately.

Account book with tier chips and composite scores

Signal DeskAccount tiering for corporate social impact

The book (8)Find accountsMeeting notesThe modelCadence and data

Sample accounts. The company names are invented, the profiles and triggers are the real shape of the market. Use an ICP preset under Find accounts to build a researched list.

Build a real listClear sample data

Sweep is due. The book has never been swept. Run it now

4One to one2One to few1One to many1Monitor8Moved0Disqualified

Re-score public signalsAdd meeting notes

Search the bookScore

1:1Kestrel Financial Group95+37

1:1Halcyon Software93+35

1:1Northgate Advisory80+25

1:1Arbor Health Systems77+37

1:fewThornbury Retail Group57+20

1:fewVantage Logistics53+8

Halcyon SoftwareSample

Invented name, real profile. B2B SaaS, roughly 1,900 employees, strong ERG culture and no central giving tool.

Re-score

1:193One to one+35 since last change

Public 45, field +48. The tick shows where public signals alone would place it.

Named account plan. Personalized page referencing the specific trigger. Executive direct mail. AE and BDR sequenced behind the physical touch.

1:1 at 93, public 45, field +48. b corp certification pursued or achieved, rapid headcount growth. no traction after sustained effort (Dev); named champion engaged (Sarah); in an active evaluation (Dev)

Account book view. Tier chips, composite scores, trigger labels, and the most recent tier-change event per account.

An account enters at the tier its score earns. It moves in both directions.

Every ABM tool promotes accounts. Almost none demote them automatically. Decay makes demotion automatic and dated, so the list is honest on a Monday without anyone having to be.

Tier	Entry	Play	Owner	Promotes when	Demotes when
1:1	70+	Named account plan. Personalized page. Executive direct mail. AE + BDR sequenced.	ABM + named AE	Meeting booked or deal opens	Below 65 and 30 days in tier, or champion departs
1:few	40–69	Cohort by trigger type. Shared narrative, personalized opener. Roundtable as conversion.	ABM + BDR pool	Crosses 70 on any lane	Below 35 and 30 days in tier
1:many	15–39	Programmatic coverage. Nurture on the trigger theme, not the product.	Demand generation	Crosses 40, or any field verdict logged	Below 10, or no activity 180 days
Monitor	below 15	No spend. Weekly sweep watches for a leadership hire, commitment, or expiring lock.	Nobody	Crosses 15 on public lane	Terminal until a signal appears
Out	gate	Not ICP, open deal, or competitor renewal inside term. No ABM spend at all.	Sales or nobody	Requalified, deal closes lost, or lock decays	Terminal while gate holds

The gate that saves the most money: an open deal means the account leaves tiering entirely. ABM spend on accounts sales is already working inflates every influenced-pipeline number reported.

Promote at the threshold. Demote five points below it.

The five-point gap prevents tier-flipping on accounts near a boundary. Thirty days minimum in tier means a play gets to finish before the engine changes its mind.

Rule	Trigger	Effect
Tier dwell (30 days)	30 days minimum in tier	A play runs to completion before the engine can change it.
Play cooldown (21 days)	21 days between play changes	Prevents conflicting campaigns on the same account.
Opt-out hold	Account-level unsubscribe or hard bounce	Channel suppressed. Rest of program continues.
Layoffs hold	Layoffs or restructuring signal	Email suppressed. Direct mail and executive programs continue.
Champion departure	Champion left or went quiet verdict	Account demoted one tier, held 30 days, score recalculates next sweep.
Competitor lock	Renewed with a competitor (-40)	Exits active tiering. Returns when the verdict decays at 180 days.

Cohorts are grouped by why an account is in market, not what it sells.

Cohort	Entry signal	Argument	Conversion moment	Tier
New leader	ESG leader hired +25	New leader has a mandate to show progress inside 12 months. Lead with peer benchmarking.	30-min benchmarking session	1:1
Displacement	Incumbent named +22	Roadmap unresolved post-acquisition. Lead with migration risk, not features.	Migration assessment	1:1
Restructure	Foundation +25, Grants +20	Companies rethinking giving vehicles. Lead with structure and compliance.	Roundtable with foundation lead	1:few
Post merger	M&A; +20	Two programs, two contracts, one renewal. Lead with cost of running both.	Consolidation workshop	1:few
Reporting obligation	Impact report +20, B Corp +18, Pledge +18	Year one off spreadsheets. Year two needs an audit trail. Lead with what breaks.	Reporting readiness audit	1:few
Outgrown the process	Headcount +15, Expansion +10	Manual program fails at a specific size threshold. Lead with the threshold.	Program review	1:many
ERG adjacent	ERG/DEI +12	Different budget holder. Warm route where CSR has no owner yet.	ERG teardown	1:many
Crisis response	Disaster +15	Matching inside 48 hours. Time-boxed to 30 days.	Fast-track matching setup	1:many

The model, not the delivery mechanism, is the reusable part.


The same scoring logic can run in several configurations. Choose based on existing stack, budget, and engineering capacity.

Option	Stack	Best when	Limitations
React app + CSV (MVP)	Browser-only, no backend	Proving the model with a small team. Zero infrastructure.	No HubSpot API from browser. CSV ceiling. Sweeps on demand.
Clay + HubSpot	Clay tables + HubSpot workflows	Team already in Clay. No engineering required.	Decay logic needs manual setup. Less transparent scoring.
Supabase + React + HubSpot	Postgres, Edge Functions, pg_cron, HubSpot API	Production deployment with AE verdicts and nightly CRM pull.	40–56 hrs for usable version. Requires engineering.
Full build + auth + Clay	All above + Clay enrichment + auth	Full team with AE self-serve verdicts and quarterly refit.	76–110 hrs total. Refit needs two quarters of closed-won data.

Build the HubSpot read before the write. Same auth, no risk of pushing junk into the CRM, and the three-lane scoring view stabilises before the push depends on it.


React tool. No backend. Data lives in the browser in the MVP.

Transcript analysis and field-verdict workflow



Meeting notes tab. Paste a Fathom or Gong transcript. Statements that map to a verdict come back with the quote and the speaker. Nothing is logged without approval.

Tier-change history log



Every tier change logged with a date, a cause, and who triggered it. Makes scoring auditable and grounds the conversation with sales.

This is not a product. It is a worked example of how to approach the problem.

The reusable part of Signal Desk is the taxonomy and the tiering logic, not the delivery mechanism. The same model could run as Clay tables plus a HubSpot workflow, which would be cheaper and fit the existing stack better. Choosing between those is the actual judgment call.

What Signal Desk documents is the decision: which signals matter and why, what evidence decays and how fast, where the ownership boundaries sit, and what suppression rules make the list honest without human maintenance. That documentation survives the tool.

Want to adapt this model to your category?

Book a call at dianapantchev.com · dpantchev@gmail.com