Database Providers

RESEARCH REPORT No. 05

Intent Data + Contact Data

Combining signals for pipeline acceleration, and why the bottleneck is almost never the signal

59%

of accounts surfaced by intent data cannot be acted on, because contact coverage of the buying committee is inadequate

MODELED FROM OBSERVED OUTBOUND PROGRAMMES

2026 EDITION

Research by Database Providers

July 2026 · thedatabaseproviders.com

SIGNALS IDENTIFY THE BUILDING; CONTACT DATA GETS YOU INSIDE

ABOUT THIS REPORT

This report examines what happens when intent data and contact data are combined, and where the resulting motion actually breaks. It models the pipeline yield of a signal-driven outbound programme as a chain of four factors, quantifies how much each one contributes, and ranks the three levers a team can pull against it.

The finding that matters is a ranking rather than a recommendation to buy anything. Of the three ways to improve a signal-driven programme, the cheapest one is worth roughly four times the most expensive one. Teams consistently invest in the lever with the smallest modeled return.

The report covers intent signal taxonomy and what each type can and cannot tell you, the decay profile of a signal measured against the decay profile of contact data, a pipeline yield model with stated inputs, a three-axis scoring approach that adds reachability to the usual fit and intent axes, and the operating controls that determine whether any of it works in practice.

PUBLISHED July 2026

EDITION 2026 Edition

RESEARCH BY Database Providers

AUDIENCE CMOs, demand generation and ABM leaders, RevOps and sales operations teams, SDR managers

REPORT CODE DP-WP-2026-05

CONTENTS

FOREWORD

Intent data is sold on the promise of timing. Somewhere out there an account is researching your category, and if you knew which one it was you could arrive at exactly the right moment instead of interrupting people at random. The promise is real, the technology largely works, and the market for it has grown accordingly.

What the promise leaves out is that timing is only half a targeting decision. Knowing that an account is in market tells you nothing about who inside it to contact, and a signal you cannot act on has no value at all. When we model signal-driven programmes end to end, the point of failure is rarely the signal. It is the moment after the signal, when an account surfaces and nobody can find two people at it worth emailing.

The second thing the promise leaves out is speed, and this is where our modeling produced a result we did not expect. We assumed the dominant lever would be contact coverage, since that is the part of the problem we work on. It is not. Cutting the time between a signal appearing and the first outbound touch is worth roughly four times as much as improving coverage by a comparable amount, and it is the only one of the three levers that costs nothing to buy. A team can improve it this quarter with a routing rule.

We have published that finding as we found it, including the part where the biggest available improvement is not something we sell. Our own product sits on the second lever, and the second lever is worth about a quarter of the first. A report that ranked the levers in the order convenient to us would be easy to write and easy to disbelieve.

Research Team, Database Providers
Fremont, California

EXECUTIVE SUMMARY

A signal-driven outbound programme converts surfaced accounts into opportunities through a chain of four factors: how many accounts surface, whether the buying committee at those accounts is reachable, how quickly the first touch happens, and how well the outreach converts. Modeling the chain shows the constraint sits in the middle two factors, not at either end.

FINDING 01

Only 41% of intent-surfaced accounts have adequate contact coverage to be acted on.

Adequate means at least two relevant buying-committee roles are reachable. A modeled 33% of surfaced accounts have exactly one reachable contact and 26% have none. The signal is generated, paid for, and then discarded for want of somebody to send to.

FINDING 02

An intent signal loses half its value between the second day and the eighth.

Modeled conversion falls from a relative index of 0.86 for a touch within 48 hours to 0.42 for a touch between eight and fourteen days. Most programmes operate at a median latency inside the second range, which means they routinely capture under half the available value of signals they have already paid for.

FINDING 03

Cutting signal-to-touch latency to 48 hours is worth a modeled 105% increase in opportunities.

Against a baseline programme surfacing 1,000 accounts per quarter at a median latency of eight to fourteen days, moving to a 48-hour median more than doubles modeled opportunity volume. No additional data purchase is required to achieve it.

FINDING 04

The latency lever is worth 4.3 times the coverage lever and both beat buying better signal.

Improving coverage by ten percentage points yields a modeled 24% increase; improving signal precision by ten percentage points yields 18%. Teams typically spend on the third lever, occasionally on the second, and rarely on the first, which is the reverse of the modeled returns.

FINDING 05

Director-level contacts decay at 36% annually, and they are the layer signal routing depends on.

Intent signals surface accounts, but outreach lands on individuals, and the mid-seniority roles that carry most evaluation work are the fastest-decaying seniority band. Coverage is therefore not a one-time acquisition problem but a maintenance problem.

KEY FINDINGS AT A GLANCE

Six numbers that summarise the research findings.

41%

SURFACED ACCOUNTS ACTIONABLE

Modeled share with at least two reachable buying-committee roles. See Chapter 1.

8 days

WHEN SIGNAL VALUE HALVES

Modeled conversion index falls by half between a 48-hour touch and an 8-to-14-day touch. See Chapter 2.

+105%

FROM CUTTING LATENCY

Modeled opportunity increase from moving median signal-to-touch latency to 48 hours. See Chapter 3.

4.3x

LATENCY VS COVERAGE LEVER

Relative modeled value of the speed lever against the coverage lever. See Chapter 3.

36%

DIRECTOR-LEVEL ANNUAL DECAY

The seniority band that carries evaluation work decays fastest. See Chapter 5.

7

ROLES IN A BUYING COMMITTEE

Typical committee size that coverage must span, not one named champion. See Chapter 1.

CHAPTER

01

Intent tells you when; contact data tells you who

Two datasets answering different questions, and a motion that requires both to be current

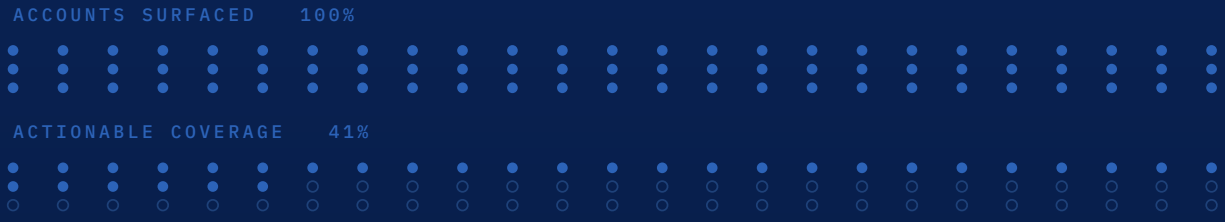

CHAPTER 01

Intent tells you when; contact data tells you who

Intent data and contact data are frequently discussed as though they were competing purchases, or as though one were an upgrade of the other. They answer different questions and neither substitutes for the other. Intent data answers a timing question: which accounts are showing behaviour consistent with an active evaluation. Contact data answers an addressing question: who inside a given account can be reached, in what role, at what address.

A signal-driven motion needs both to be true simultaneously. An account can be unmistakably in market and completely unreachable, in which case the signal has cost money and produced nothing. A contact file can be immaculate and the account not in market, in which case the outreach is well-addressed and badly timed. The combination is what produces a meeting, and the combination is only as good as its weaker half at the moment of action.

That last qualification is the one this report is built around. Both datasets degrade, but they degrade on completely different clocks, and the practical consequence is that a programme can be well-supplied with both and still fail because the two were never current at the same time.

Fewer than half of surfaced accounts have enough reachable committee members to run a credible multi-threaded sequence.

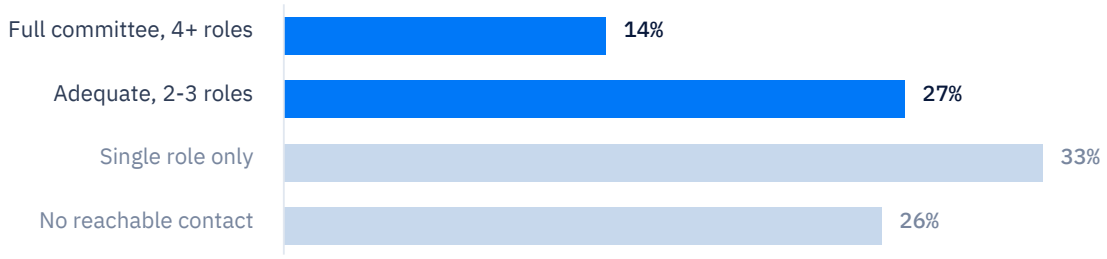

Source: Database Providers, modeled buying-committee coverage of intent-surfaced accounts, 2026. Adequate-or-better totals 41%.

What the signal types can and cannot tell you

Not all intent data carries the same information, and the differences matter for how a signal should be routed. First-party signals come from your own properties: pricing page visits, documentation reads, repeat sessions, form abandonment. They are the highest-confidence signals available because they are unambiguous about the subject of interest, and they usually identify a person rather than only an account.

Second-party signals come from platforms where buyers research categories, including review and comparison sites. They are strong on category intent and typically resolve to an account rather than an individual. Third-party signals, often assembled from content consumption and bidstream data across publisher networks, cover the widest surface and carry the most noise; they usually resolve to an account or even to a location, and the topic attribution can be loose.

Engagement signals sit slightly apart: opens, clicks, and replies on your own outbound. They are frequently ignored in intent discussions because they are not bought, which is precisely why they are underused. They are free, they are person-level, and they are the most recent signal any programme has access to.

The routing consequence is that signal type should determine the shape of the response. A person-level first-party signal justifies an immediate individual touch. An account-level third-party signal justifies research and a multi-threaded approach to a committee, because the signal does not tell you who moved. Treating both the same

way is a common and expensive simplification.

> *An account can be unmistakably in market and completely unreachable. The signal has then cost money and produced nothing.*

Why committees rather than champions

The coverage measure used throughout this report counts reachable buying-committee roles rather than reachable contacts, and the distinction is deliberate. A typical B2B purchase involves on the order of seven roles: an economic buyer, a technical evaluator, an end-user representative, a procurement or finance gatekeeper, a security or compliance reviewer where relevant, an internal champion, and at least one influencer with no formal role in the process.

A programme holding one contact at an account is exposed in a specific way. If that contact is not the person the signal reflects, the outreach lands on someone with no current interest. If the contact forwards internally, the thread leaves your visibility. And if the contact has left, which at director level happens to roughly a third of records in a year, the account looks covered in the CRM and is not.

This is why the coverage threshold used here is two or more relevant roles rather than one. Two reachable roles allow a genuinely multi-threaded approach in which the failure of any single contact does not end the attempt. On the modeled distribution, 41% of surfaced accounts clear that bar, and 14% have the four or more roles that would support a full committee motion.

WHAT THIS MEANS FOR LEADERS

— Measure coverage as reachable committee roles per account, not as contacts held.
— Match the response shape to the signal type: person-level signals get individual touches, account-level signals get committee research.
— Use your own engagement data as an intent source, since it is free, person-level, and the most recent signal you hold.
— Treat an account with one reachable contact as uncovered rather than covered.

CHAPTER

02

Signals decay in days, contact data decays in months

Two clocks running at different speeds, and the one most programmes fail to measure

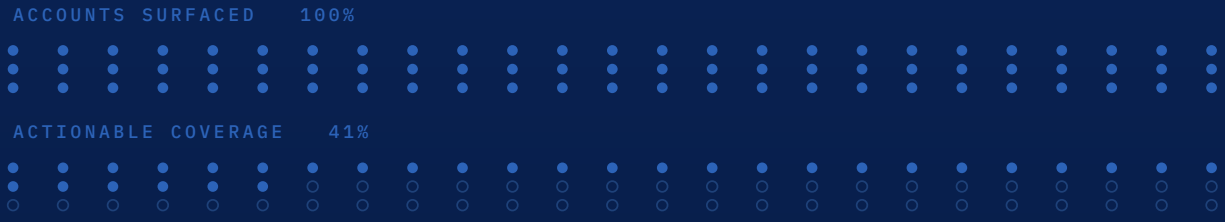

CHAPTER 02

Signals decay in days, contact data decays in months

The reason a signal-driven programme can hold good data on both sides and still underperform is that the two sides expire on incomparable timescales. Contact data decays at a modeled 3.05% per month, which compounds to roughly 31% over a year. That is slow enough to be managed with a quarterly verification cadence and slow enough that a team can be complacent about it for a long time before consequences appear.

Intent signals decay in days. An account researching a category this week may have shortlisted by next week and selected within the month. The signal is not a description of a durable state; it is a snapshot of a process already in motion, and the process does not pause while a signal sits in a queue.

Modeled against time-to-first-touch, the conversion index falls steeply and early. A touch inside 24 hours sets the reference point. By 24 to 48 hours the index is at 0.86. Between three and seven days it is 0.61. Between eight and fourteen days it is 0.42, and beyond fourteen days it is 0.29. The signal loses half its value between the second day and the eighth.

Half the value of an intent signal is gone by the eighth day, which is inside most programmes' normal operating latency.

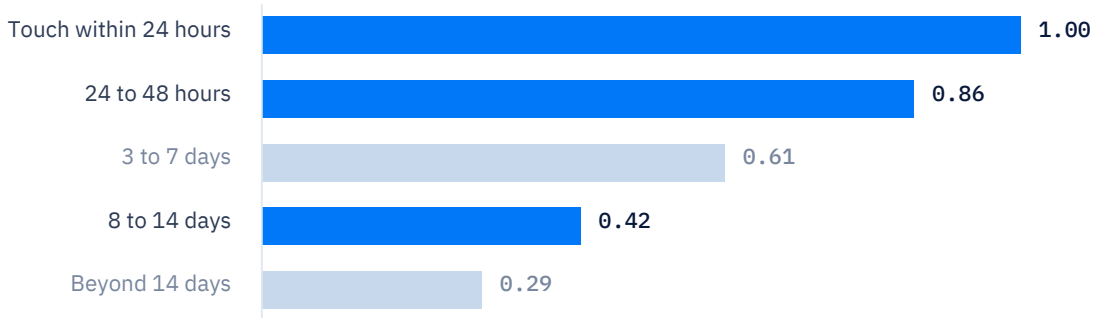

Source: Database Providers, modeled relative conversion index by signal-to-touch latency, 2026. Indexed to a touch within 24 hours.

Signal-to-touch latency, and why nobody measures it

The metric implied by that curve is the interval between a signal being generated and the first outbound touch reaching a person at the account. We refer to it as signal-to-touch latency, and in our experience it is almost never instrumented. Teams measure signal volume, because the vendor reports it. They measure reply and meeting rates, because the sequencing tool reports them. The interval between the two is owned by nobody and appears on no dashboard.

The reasons it stretches are mundane and cumulative. Signals arrive on a vendor's batch schedule rather than in real time. They land in a platform that someone reviews weekly. Accounts are checked against ownership rules and routed to a rep, who is working a queue prioritised by something else. Contacts are missing, so the account goes to a research step. Each step is defensible and the sum is often two weeks.

Because the curve is steep in exactly that range, small process improvements produce disproportionate returns. Moving a weekly review to a daily one, and routing accounts that already have coverage straight to sequence rather than through research, typically removes the largest share of the delay without changing headcount or buying anything.

The measurement itself is straightforward: timestamp signal receipt, timestamp first touch, report the median and the ninetieth percentile per signal source. The ninetieth percentile matters more than the median here, because a small tail of very slow accounts is usually where the highest-value signals are sitting, having been routed to research precisely because they looked important.

> *Teams measure signal volume and reply rates. The interval between the two is owned by nobody and appears on no dashboard.*

How the two clocks interact

The interaction is what makes coverage a speed problem as well as a completeness problem. When a signal surfaces an account with no reachable contacts, the response is research: find the committee, verify addresses, then sequence. That research is not free in time, and the time it consumes is spent against the signal decay curve rather than alongside it.

So the cost of thin coverage is not only the accounts that are never actioned. It is also the delay imposed on the accounts that eventually are. An account that takes six days to research and then converts at the 0.61 index has lost nearly forty percent of its value to a data gap, and that loss is invisible because the eventual outcome looks like a normal outbound result.

The design conclusion is that coverage should be built ahead of signals rather than in response to them. A programme that maintains committee coverage across its target account list can route a signal to sequence the same day. A programme that acquires contacts reactively pays the decay penalty on every signal it receives, however good the signal was.

WHAT THIS MEANS FOR LEADERS

- Instrument signal-to-touch latency and report the median and ninetieth percentile by source.
- Move signal review from weekly to daily before considering any additional data spend.
- Route accounts that already have coverage straight to sequence, bypassing the research step.
- Build committee coverage ahead of signals rather than researching accounts after they surface.

CHAPTER

03

The three levers, ranked

A pipeline yield model with stated inputs, and why the cheapest lever is worth the most

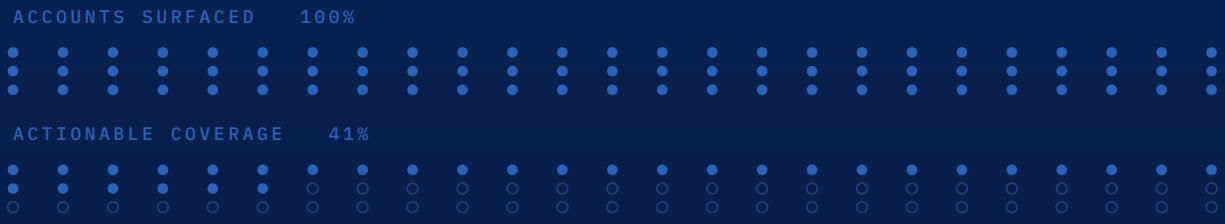

CHAPTER 03

The three levers, ranked

Ranking the available improvements requires a model, and the model can be simple enough to argue with. Modeled opportunity volume from a signal-driven programme is the product of four factors: the number of accounts surfaced in the period, the share of those accounts with adequate committee coverage, the latency multiplier from Chapter 2, and the base rate at which a well-timed, well-addressed sequence produces an opportunity.

The baseline used here is a programme surfacing 1,000 accounts per quarter, with 41% adequate coverage, a median signal-to-touch latency in the eight-to-fourteen-day band giving a multiplier of 0.42, and a 6.0% conversion from actioned account to opportunity. That produces a modeled 10.3 opportunities per quarter.

Three interventions are then applied one at a time, each sized to be roughly comparable in ambition. Cut median latency to the 24-to-48-hour band, raising the multiplier to 0.86. Improve coverage by ten percentage points, from 41% to 51%. Or improve signal precision by ten percentage points, which raises the conversion rate proportionally on the reasoning that a higher share of surfaced accounts are genuinely in market.

Speed is the dominant lever by a wide margin, and it is the only one of the three that costs nothing to buy.

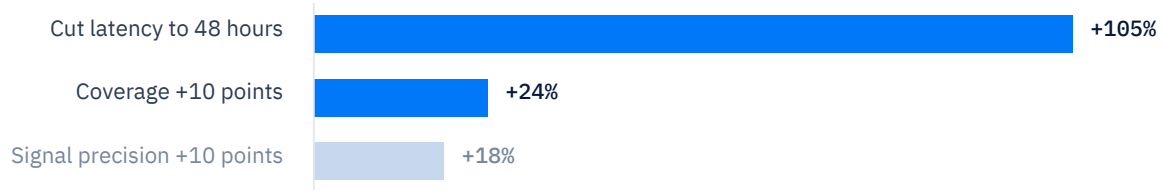

Source: Database Providers, modeled change in quarterly opportunities from a 1,000-account baseline. Inputs stated in the methodology.

Reading the ranking

The latency intervention takes modeled opportunities from 10.3 to 21.2, an increase of 105%. The coverage intervention takes them to 12.9, an increase of 24%. The precision intervention takes them to 12.2, an increase of 18%. The latency lever is worth 4.3 times the coverage lever and just under six times the precision lever.

The reason is structural rather than a quirk of the inputs. Latency enters the model as a multiplier that more than doubles across the range in question, because the decay curve is steep exactly where most programmes sit. Coverage and precision each move by ten percentage points on bases of 41% and an implied 55%, which are proportionally much smaller changes. Any model with a steep early decay curve and mid-range operating latency will produce this ordering.

The ordering is also the inverse of typical spending. Signal precision is what vendors compete on and what procurement conversations are about. Coverage is occasionally addressed, usually by buying more contacts. Latency is a process variable with no purchase attached, no vendor advocating for it, and no line in a budget, which is a reasonable explanation for why it is the one left untouched.

A reader who disputes the inputs can rerun the model in a spreadsheet in a few minutes, and we would encourage it. The inputs most worth substituting are the conversion base rate, which varies widely by category and deal size, and the baseline latency, which a team can measure directly rather than assume. The ordering is robust to plausible changes in both; what changes is the size of the gap.

> *Latency is a process variable with no purchase attached and no vendor advocating for it, which explains why it is the lever left untouched.*

Where the levers compound

The three interventions are not alternatives, and applying them together produces more than the sum of the individual gains because they multiply rather than add. A programme that cuts latency to 48 hours and improves coverage by ten points reaches a modeled 26.4 opportunities against a 10.3 baseline, an increase of 156%, against 129% if the two gains were merely additive.

Sequencing still matters, because the interventions differ in cost and lead time. Latency improvements are process changes achievable within a quarter at no data cost. Coverage improvements require either acquisition or verification spend and take effect as the file is built. Precision improvements usually mean a vendor change or an additional source, which carries the longest lead time and the highest cost per unit of modeled return.

The recommendation that follows is unusual for a report published by a data supplier, so we will state it plainly. Fix routing speed first, because it is free and worth the most. Address coverage second, because it is worth a quarter as much and is where our own product sits. Consider signal precision last, because it is the most expensive route to the smallest modeled gain.

WHAT THIS MEANS FOR LEADERS

- Measure your current signal-to-touch latency before authorising any additional data spend.
- Sequence improvements by cost and lead time: routing speed, then coverage, then signal precision.
- Rerun the model with your own conversion rate and measured latency rather than accepting these inputs.
- Expect combined interventions to compound, since the factors multiply rather than add.

CHAPTER

04

Scoring on three axes, not two

Adding reachability to fit and intent turns a score into a work queue

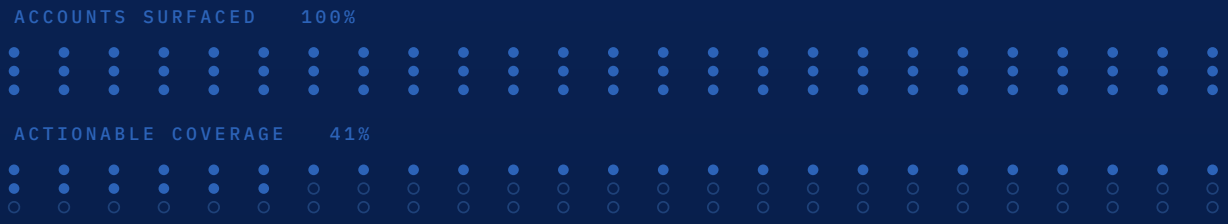

CHAPTER 04

Scoring on three axes, not two

Most account scoring models combine two axes. Fit describes whether an account resembles the accounts you win: firmographics, technographics, size, sector, sometimes funding stage. Intent describes whether the account is showing evaluation behaviour now. The product of the two produces a prioritised list, and the list is then handed to a team to work.

The omission is that a score built on fit and intent alone ranks accounts by desirability rather than by actionability. An account can score at the top of both axes and be unworkable this week because nobody at it is reachable. When that account reaches the top of a queue, the rep does not make a call; the rep starts a research task, and the signal begins decaying while the research proceeds.

Adding reachability as a third axis fixes the mismatch between what the score promises and what the queue can deliver. Reachability is straightforward to compute from data a team already holds: how many relevant committee roles are reachable at this account, and how recently were those records verified.

A workable score weights reachability alongside fit and intent, because an unreachable account cannot be worked today.

Source: Database Providers, indicative scoring weights for a signal-driven outbound queue, 2026. Weights should be tuned against observed outcomes.

What each axis should actually contain

Fit is the most stable of the three and the easiest to over-engineer. It rewards resemblance to closed-won accounts, and the components that carry most of the predictive weight are usually few: size band, sector, and one or two technographic or lifecycle markers relevant to the category. Elaborate fit models tend to encode historical sales coverage rather than genuine propensity.

Intent needs both strength and recency, and recency should be weighted more aggressively than most implementations do. Given the decay curve in Chapter 2, a strong signal from twelve days ago is worth materially less than a moderate signal from yesterday, and a score that treats a signal as a durable attribute of the account will systematically misrank the queue. Signal type belongs here too, since a person-level first-party signal carries more information than an account-level third-party one.

Reachability has two components. Coverage counts distinct relevant committee roles with a reachable address, capped so that ten contacts in one function does not outscore three across three functions. Freshness discounts records by age using the monthly decay rate, so an account verified last month scores above an identical account verified a year ago. The second component is what stops reachability from silently degrading into a historical claim.

The output that makes this operationally useful is not a single ranked list but a split. Accounts scoring high on all three go straight to sequence today. Accounts high on fit and intent but low on reachability go to a research or enrichment queue with an explicit deadline set against the decay curve. Accounts high on fit and reachability but

low on intent go to nurture. Separating those paths is what prevents a research task from sitting in the same queue as a ready-to-send account.

> *A score built on fit and intent alone ranks accounts by desirability. Adding reachability ranks them by what can be worked today.*

Tuning rather than adopting

The weights above are a starting point rather than a recommendation, and the honest position is that any published weighting is a guess about someone else's business. The 40/35/25 split reflects a programme where signals are plentiful and coverage is the constraint, which is the common case in our observation but not a universal one.

Two adjustments are worth making early. A programme with abundant first-party signal from a high-traffic site should weight intent higher, because the signal is both more plentiful and more reliable. A programme selling into a market where committees are large and mobility is high, which describes several of the Gulf and technology segments Database Providers covers, should weight reachability higher, because that is where the variance sits.

The tuning method is the same in either case: hold the weights fixed for a quarter, record which scoring band each closed opportunity came from, and adjust toward the bands that produced outcomes. This is unglamorous and it beats any published weighting, including this one.

WHAT THIS MEANS FOR LEADERS

- Add reachability as a third scoring axis, computed from committee coverage and record freshness.
- Weight intent recency aggressively rather than treating a signal as a durable account attribute.
- Split the output into send-now, enrich-now, and nurture queues rather than one ranked list.
- Tune weights against your own closed-opportunity data instead of adopting a published split.

CHAPTER

05

Operating the combined motion

The controls that determine whether the model holds, and the anti-patterns that break it

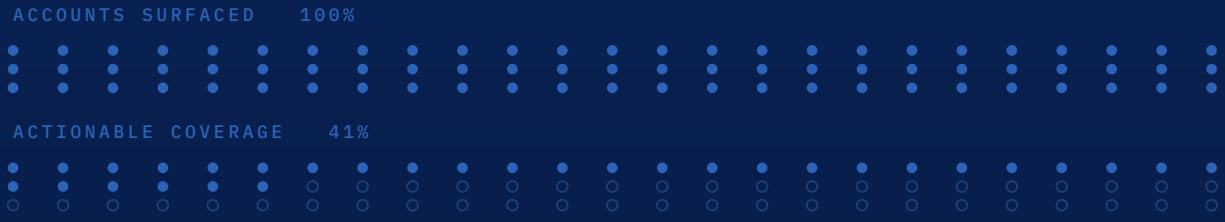

CHAPTER 05

Operating the combined motion

The model in Chapter 3 assumes a programme that can act on a signal quickly and has somebody to act on. Both assumptions rest on operating controls rather than on data purchases, and in our observation most programmes have few of them in place. That is the gap between the modeled ranking and what actually happens in a quarter.

Five controls carry most of the weight. A routing service level agreement that commits to first touch within a stated window. A committee refresh cadence that keeps coverage current on the target list rather than acquiring contacts reactively. Reachability included in scoring, as set out in Chapter 4. Signal-to-touch latency measured and reported. And deduplication across signal sources, so an account appearing in three feeds generates one coordinated response rather than three uncoordinated ones.

Modeled adoption of each is low, and the pattern is instructive: the controls that require a process decision are more common than the controls that require measurement, which suggests teams are willing to change how they work but have not established what to watch.

Deduplication is common; the measurement and maintenance controls that the model depends on are rare.

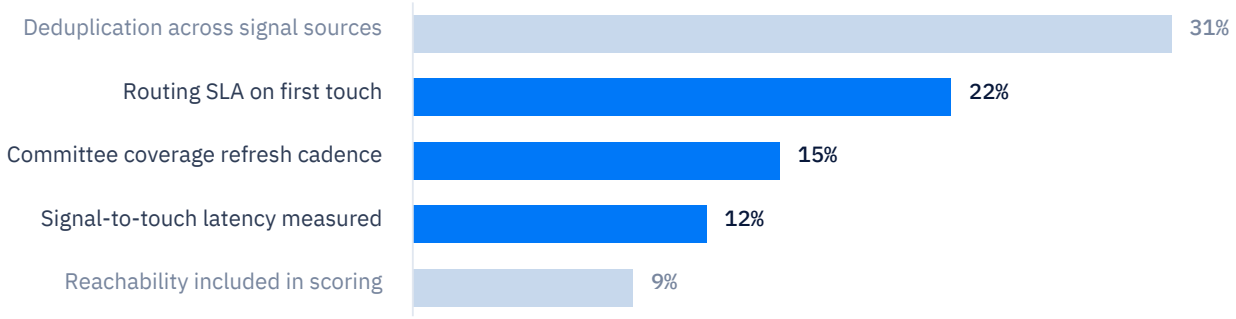

Source: Database Providers, modeled share of signal-driven outbound programmes with each control in place, 2026.

Coverage as maintenance, not acquisition

The control with the widest gap between importance and adoption is committee refresh cadence, and the reason it gets neglected is that coverage looks like a solved problem once contacts have been acquired. The CRM shows contacts at the account. The account appears covered. Nothing reports that the coverage has expired.

The decay figures make the expiry concrete. Contact records decay at a modeled 31% annually overall, and director-level records, which carry most of the evaluation work in a buying committee, decay at 36%. A committee map assembled twelve months ago has lost roughly a third of its mid-seniority members, and those are the members most likely to reply.

For a signal-driven programme this failure has a particular shape. A signal surfaces an account, the queue reports adequate coverage, the sequence goes out, and it lands on people who have left. The programme records a low reply rate and concludes the signal was weak. The signal was fine and the coverage was twelve months stale.

The fix is a refresh cadence applied to the target account list rather than to the whole database: quarterly for the mid-seniority band, and timed against sector volatility rather than the calendar. That is a much smaller volume of verification than refreshing everything, and it is applied exactly where the motion depends on it.

Anti-patterns worth naming

Four failure patterns recur often enough to be worth stating directly. The first is buying a second intent source to fix a coverage problem, which adds signal volume to a programme that is already discarding most of the signal it has. The second is routing every surfaced account to research by default, which imposes the full decay penalty on the accounts that needed no research at all.

The third is treating a signal as a durable state, most visibly when accounts surfaced weeks earlier remain in a queue at their original priority. The fourth is measuring the programme on reply rate alone, which conflates signal quality, coverage accuracy, latency, and message quality into a single number that cannot diagnose any of them. A programme reporting a declining reply rate has learned nothing about which of the four factors moved.

Diagnosing properly requires separating them, and the separation is not difficult. Report coverage adequacy as a share of surfaced accounts. Report latency median and ninetieth percentile. Report reply rate only among accounts that were both adequately covered and touched inside the target window, because that is the only population from which message quality can be inferred. Three numbers instead of one, each of which points at a specific owner.

> *A programme reporting a declining reply rate has learned nothing about which of the four factors actually moved.*

WHAT THIS MEANS FOR LEADERS

- Set a routing SLA on first touch and measure compliance, since speed is the dominant lever.
- Refresh committee coverage quarterly on the target list, concentrating on the mid-seniority band.
- Never buy a second signal source while most of the first source is going unactioned.
- Report coverage, latency, and reply rate separately so a decline can be attributed to a cause.

ACTION FRAMEWORK

The Signal Activation Model

Signal-driven programmes sit at one of four levels, defined by whether a signal can actually be converted into a timely, well-addressed touch. Each level names the specific capability that unlocks the next, and none of them requires a new data purchase to reach.

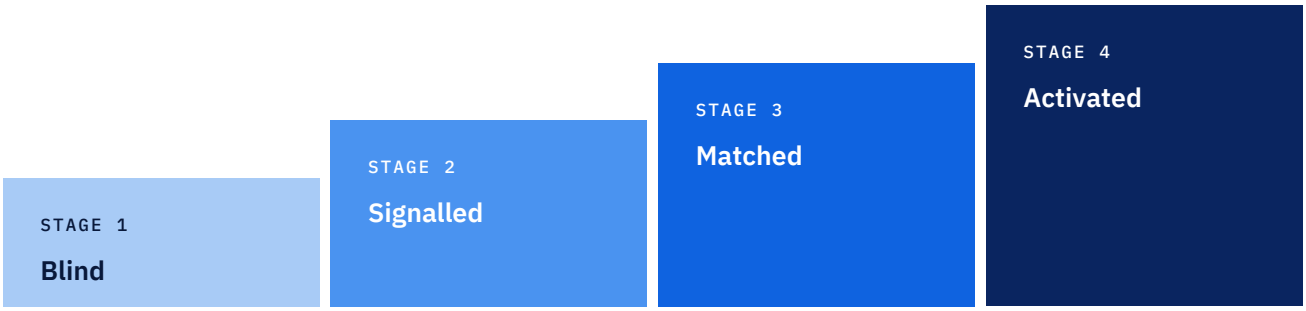

STAGE 01

Blind

Outbound runs on static target lists with no intent input, so timing is arbitrary and every account is approached as though it were equally ready. Reply rates are treated as a function of message quality alone, because no other variable is visible. The step up is not buying intent data; it is instrumenting the engagement signals the programme already generates, which are free and person-level.

STAGE 02

Signalled

Intent signals arrive from one or more sources and are reviewed on a schedule, usually weekly. Accounts are routed to owners and worked, but median signal-to-touch latency sits in the range where most of the signal's value has already decayed, and thin coverage sends a large share of accounts to research. The step up is measuring latency and committing to a first-touch window.

STAGE 03

Matched

Committee coverage is maintained ahead of signals on the target account list, reachability is part of the score, and accounts with existing coverage route straight to sequence rather than through research. Latency is measured and reported by source. Most of the modeled gain available to a signal-driven programme is captured at this level.

STAGE 04

Activated

Signals are deduplicated across sources into one coordinated response per account, refresh cadence is tuned to the volatility of each segment rather than the calendar, and the programme reports coverage, latency, and conversion separately so a change can be attributed to a cause. Scoring weights are tuned against closed-opportunity data rather than adopted from a published split.

METHODOLOGY

Coverage and control-adoption figures are modeled from Database Providers observations of signal-driven B2B outbound programmes and are directional rather than measured. Coverage counts distinct relevant buying-committee roles with a reachable, verified address at an account, where adequate means two or more such roles. The seven-role committee description is a typical composition rather than a fixed structure and varies by category and deal size.

The latency multipliers in Chapter 2 are modeled relative conversion indices by time from signal generation to first outbound touch, indexed to a touch within 24 hours: 1.00 within 24 hours, 0.86 at 24 to 48 hours, 0.61 at three to seven days, 0.42 at eight to fourteen days, and 0.29 beyond fourteen days. The statement that a signal loses half its value between the second and eighth day follows from the 0.86 and 0.42 values.

The pipeline yield model in Chapter 3 is the product of four factors: accounts surfaced per quarter, share with adequate coverage, latency multiplier, and conversion from actioned account to opportunity. The stated baseline is 1,000 accounts, 41% coverage, a 0.42 multiplier, and a 6.0% conversion rate, giving 10.3 modeled opportunities. Interventions are applied singly: latency to the 24-to-48-hour band (0.86) gives 21.2; coverage to 51% gives 12.9; signal precision improved by ten percentage points, modeled as raising conversion proportionally from an implied 55% precision base to 65%, gives 12.2. The combined latency and coverage case gives 26.4. All percentages are computed against the 10.3 baseline.

Contact decay figures are carried over from the 2026 B2B Data Decay Report, where they are modeled from aggregate verification outcomes: 3.05% per month compounding, approximately 31% annually across all segments and 36% annually for director-level records. Scoring weights in Chapter 4 are indicative for planning and are explicitly not a recommendation to adopt without tuning against observed outcomes.

NOTES AND LIMITATIONS

— All figures described as modeled are directional estimates rather than measured results from a controlled study, and will not match the experience of any individual organisation.
— The pipeline model is deliberately simple so that readers can substitute their own inputs. The conversion base rate and baseline latency are the two inputs most worth replacing with measured internal figures.
— The ranking of the three levers is robust to plausible changes in the inputs; the size of the gaps between them is not.
— Signal precision is modeled as a proportional effect on conversion rather than measured directly, since precision definitions differ substantially between intent vendors.
— Database Providers supplies B2B contact data, which is the second of the three levers examined here and not the highest-value one. Readers should weigh the findings accordingly.
— Scoring weights and committee compositions are indicative and should be tuned against closed-opportunity data for the reader's own category.

HOW DATABASE PROVIDERS CAN HELP

Fix routing speed first. When you get to coverage, that is where we help.

This report ranks routing speed above contact coverage, and coverage is the lever we sell. We would rather say that plainly than have you spend on the second-best lever first. When coverage is the constraint, the requirement is committee-level breadth on the target list and a refresh cadence that keeps the mid-seniority band current.

— Committee-level contact coverage across 120+ verticals, built to span multiple buying roles per account rather than a single named contact.

— Quarterly re-verification of existing CRM records, concentrated on the director and manager band where decay runs fastest.

— Free sample data in your target vertical and seniority band, so you can measure your own coverage adequacy before deciding whether it is your constraint.

ABOUT DATABASE PROVIDERS

Database Providers is a USA-based B2B data services company headquartered in Fremont, California. For more than a decade, the firm has helped over 750 organizations across 120+ business verticals connect with verified decision-makers through segmented email databases, data appending, verification, and enrichment services spanning healthcare, technology, manufacturing, banking and financial services, education, and real estate markets in the USA, India, the UAE, Singapore, and Malaysia.

Database Providers publishes benchmarks and practitioner guidance drawn from its verification operations, which maintain a global database of 140 million+ business contacts. Reports in this series are intended to be usable by practitioners rather than promotional, and are published with their methodology, inputs, and limitations stated in full.

10+	750+	120+	140M+
YEARS	CUSTOMERS	VERTICALS	CONTACTS

Request a free data sample

Measure your own coverage adequacy before you buy anything. Request a sample in your target vertical and seniority band, check how many committee roles it reaches per account, and compare it against your current file.

DATABASE PROVIDERS

39109 Guardino Dr, Fremont, CA 94538, USA

sales@thedatabaseproviders.com

www.thedatabaseproviders.com