From signal to sales pipeline

Market, signals, the chain, the data layer, send-readiness

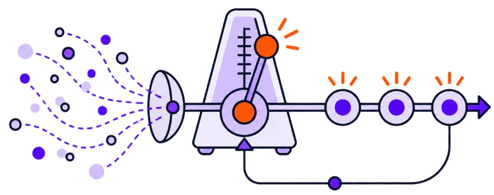

Inside:

The market first, then the signals · How to find market demand · Step 2: Define the signals, four classes · Define the signals with sales, not for them

17 PAGES TO PRINT AND FILL IN · 4 DIAGRAMS

AN OPERATING PARTNER FOR REVENUE SYSTEMS · HUBSPOT DIAMOND PARTNER

Free 60-minute Pipeline Generation Launchpad

pages.thesalesplaybook.com/meetings/erik-steffen/-pipeline-generation-assessment-call

Contents

Every page stands on its own. If you are short on time, jump straight to the step where it is stuck at your end.

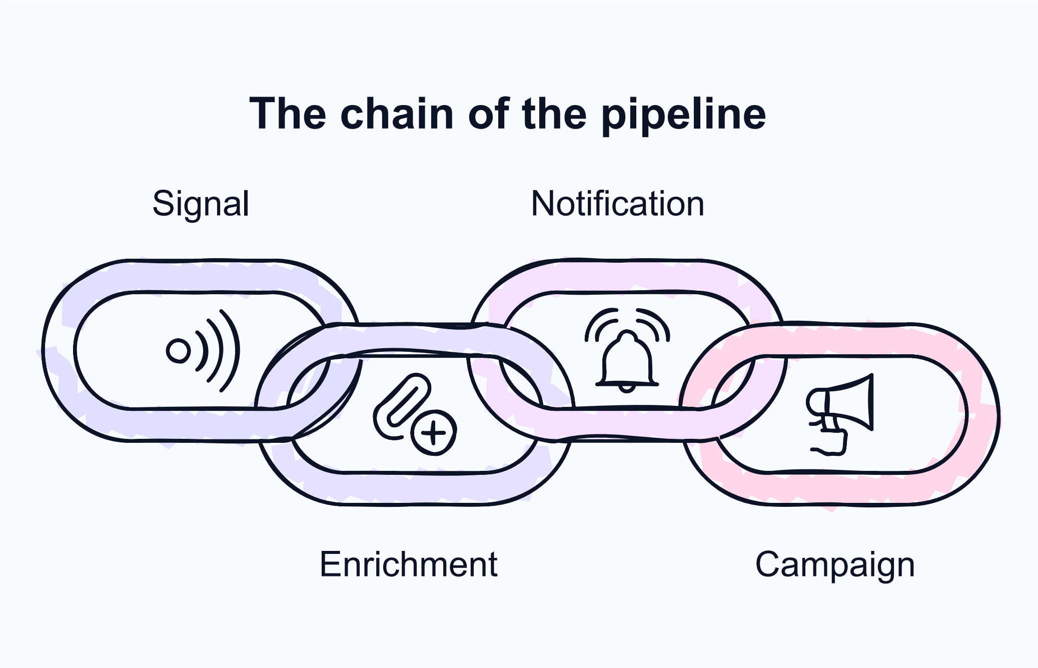

SKETCH · THE CHAIN

What this is for, and what it is not

Detecting buying signals is the easy part. There are tools for it, they are affordable, and they work. The hard part starts afterwards: turning a detected signal into a message that lands, and messages into qualified opportunities that arrive predictably.

What you end up with

A defensible number for how large your reachable market actually is. Three to five signals that are observable in that market. A chain from signal to first message, with its breaking points named. A data layer somebody else can read. And a list that is send-ready rather than merely large.

What this is not: a tool comparison. Specific tools appear because part of the mechanics attaches to them and cannot be described concretely otherwise. That is not a recommendation. If you work with different tooling, you transfer the mechanics.

This is for you if you run or are building outbound and have noticed that volume no longer helps. The lists get bigger, the reply rates do not, and nobody can say precisely why one campaign worked and the next one did not.

Not for you if your target segment is undefined. Signal-based work assumes it is clear who a signal counts for. Without that foundation it produces more noise, not less.

Time required: steps 1 and 2 are one two-hour session with sales and marketing together. Steps 3 to 5 are implementation and take days to weeks depending on the state of your data.

DOCUMENTED · CALIBO

Calibo turned 700k impressions into C-level meetings and a 400k deal, in four months. [Read the case study]

The market first, then the signals

Most outbound builds start with tools and lists. The step before that decides whether any of it can work: **how many companies realistically qualify for you?**

That number answers three questions at once, and all three are expensive if you ask them later.

QUESTION	WHAT THE NUMBER DECIDES
Does outbound carry at all?	If the market is too thin for volume, outbound is the wrong answer
What can you promise?	Market size bounds what volumes and outcomes you can honestly commit to
How much infrastructure do you need?	Contact volume determines the number of senders, inboxes and domains

GRAPHIC G1

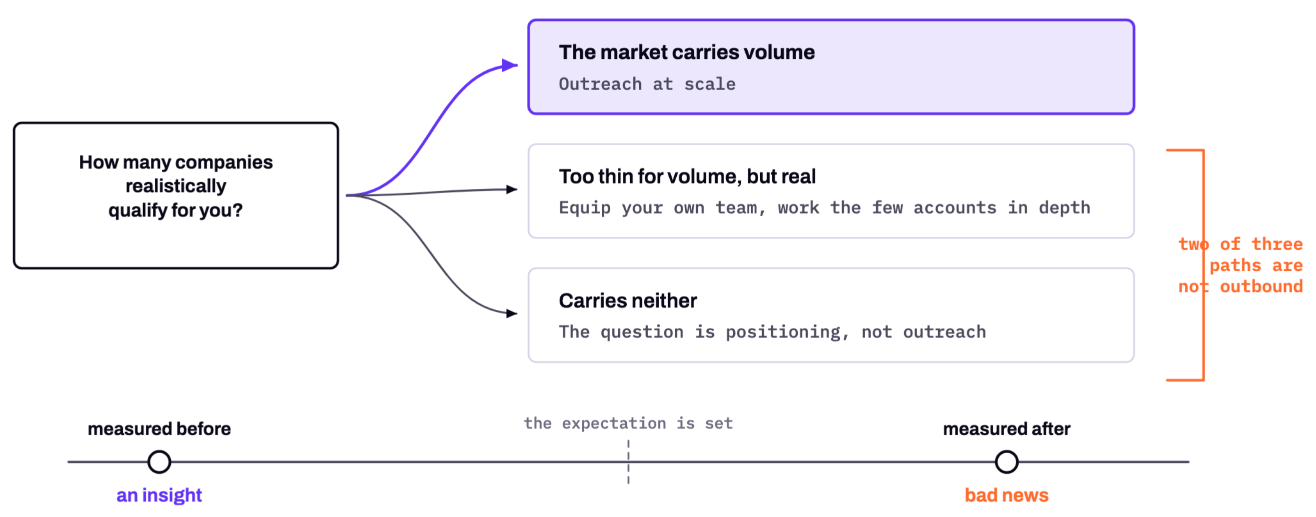

Three possible outcomes, and two of them are not outbound:

- **The market carries volume.** Outreach at scale makes sense.

- **The market is too thin for volume, but real.** Then the right move is to equip your own sales team and work the few accounts in depth, rather than running volume campaigns. Scaling anyway burns the market inside a quarter.
- **The market carries neither.** Then the question is not the outreach, it is the positioning.

Why the order matters. Measuring market size afterwards means measuring it once the expectation is already set. A number that comes in low is then no longer an insight, it is bad news.

The rule of thumb for depth, and it depends only on market size:

REACHABLE COMPANIES	APPROACH
50 to 1,000	Map the market precisely, work account by account, depth over volume. Never burn an account twice
2,000 to 10,000	A mix: the strongest segments in depth, the rest with clean automation
over 10,000	Filter efficiently, use signals to pre-sort, reserve depth for those who responded

The question a market stands or falls on is four words long. Manuel Hartmann puts it this way in his LinkedIn newsletter *Diary Of A CRO*: is there still room here? And behind it: can the market carry it? ([Read the issue, in German]())

Both belong before the method choice on the next page. A cleanly built list in a saturated market is clean and still worthless.

> GO DEEPER
>
> [starting the plan at the number]()

How to find market demand

The method depends on where your target segment is findable at all. Four cases, and the choice decides effort and accuracy.

IF THE SEGMENT	THEN	WHY
sits in a **regulated or institutional** market	public registers and directories	the record there is complete and verified, which beats any database
is **well represented on LinkedIn** (typical for SaaS)	databases plus targeted searches	coverage is high, currency is workable
has a **physical presence**	query mapping services programmatically	locations are more complete there than in company databases
is **very narrow**, under roughly 500 to 700 firms	research agents	most expensive method, but finding and enriching happen in one pass

The fourth case gets decided wrongly most often. With a small market the expensive method feels disproportionate. It is not: at a few hundred companies, the usual combination of sourcing plus a separate enrichment run is slower and less accurate than one pass that does both. The output still has to be checked.

Market is not list

The distinction gets blurred constantly and costs weeks:

- **The market** is the set of all companies that could qualify. You build it once, then maintain it, and it changes slowly.
- **The list** is the slice that goes into one campaign. It is drawn from the market, enriched and consumed.

Mixing the two means rebuilding the market for every campaign and never knowing which part has already been worked. The market lives in one place, lists get drawn from it.

And a check that belongs before scaling: build a small, **fully** enriched sample from each segment, ten to twenty companies. Not as a test of the tools, but as a test of the definition. If the sample does not convince you, the filter is wrong, and twenty rows is a cheaper place to learn that than five thousand.

GO DEEPER

cut by job, not by size

Step 2: Define the signals, four classes

A signal is worth having when it is **observable** before the competition sees it. Anything else is catching up.

EVIDENCE [what a signal really tells you]

GRAPHIC G2

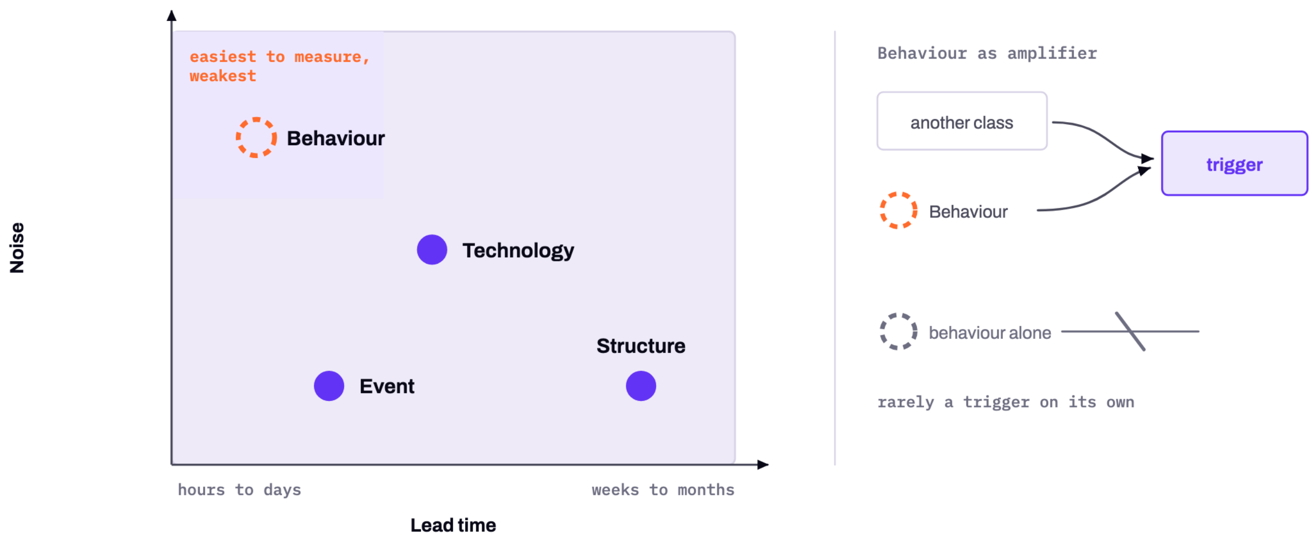

CLASS	EXAMPLES	LEAD TIME	NOISE
Behaviour	site visit, content consumed, interaction	hours to days	**high**
Event	funding, merger, regulation, market event	days to weeks	low
Technology	tool adopted or dropped, integration visible	weeks	medium
Structure	new role filled, team growing, site opened	weeks to months	low

Sorted by lead time, short to long. Noise is the share of a class's hits that are not buying interest. Low means almost every hit is one. High means most are not.

The behaviour class is the most tempting and the weakest. It is easiest to measure, has the shortest lead time and the highest noise. A website visit can be buying interest, a competitor, a job applicant or an accident. It works as an **amplifier** of another signal, rarely as a trigger on its own.

The selection questions, per candidate

1. **Is it observed or inferred?** A posted job is observed. "That company is probably growing" is inferred.
2. **How long is the window?** A signal with a three-day window needs a process that reacts in three days. Without one, that signal is not usable by you.
3. **What follows from it for the outreach?** If no opening line can be derived, it is a filter rather than a trigger.
4. **How many hits per month?** Under ten does not justify the build. Over several hundred it is not a signal, it is a segment.

Three to five signals is the target. More means none of them gets processed properly.

Signal 1: class __________ · observed ☐ · window __________ · hits/month __________

The opening line derived from it:

____

GO DEEPER

from signal to conversation on LinkedIn

Define the signals with sales, not for them

This step routinely happens in marketing alone, and that is why the resulting notifications get ignored later.

What a conversation with sales gives you that data does not:

Which signal actually changes a conversation. An AE knows from experience which trigger produces a reply and which gets politely ignored. That information appears in no data source.

How the trigger can be raised without sounding intrusive. Some signals are public but mentioning them feels invasive. The person having the conversations knows where that line is.

Which signals preceded the deals you won. This is the most valuable source and the least used. Walk backwards through five won deals and ask what happened before first contact. That produces better candidates than any signal library.

The session structure, two hours:

TIME	WHAT
30 min	Five won deals, backwards: what happened before first contact
30 min	Collect candidates without evaluating them
40 min	The four selection questions from page 6, per candidate
20 min	Pick three to five, discard the rest with reasons recorded

Discarded candidates get documented, not deleted. In six months somebody proposes the same thing, and then the reasoning is available rather than needing to be rebuilt.

Where the machine stops

The *Diary Of A CRO* issue on house rules for selling with judgment in the age of AI reduces to one sentence: AI builds the target account list, a human writes the letter. (Read the issue, in German)

For this step that means market and signals run on data, while the wording of the trigger runs through the person having the conversations. Automate both and you get messages that visibly came out of a field, and those are exactly what train a recipient to stop reading.

Step 3: Build the chain, link by link

A detected signal with no chain behind it is a notification nobody works. That is the most common state in teams that have just introduced a signal tool: detection runs, notifications arrive, and the pipeline does not change.

GRAPHIC G3

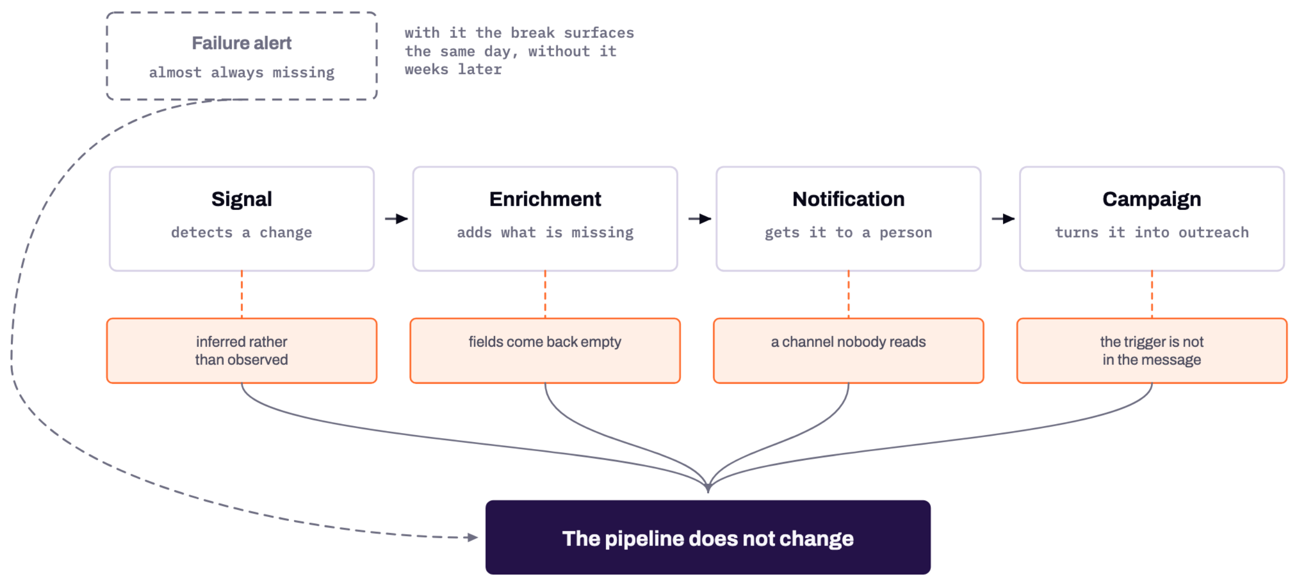

LINK	WHAT IT DOES	WHERE IT BREAKS
Signal	detects that something changed	the signal is inferred rather than observed
Enrichment	adds what the outreach needs	enrichment runs, but the fields come back empty
Notification	gets it to a person	it lands in a channel nobody reads
Campaign	turns it into outreach	the trigger is not mentioned in the message

Link 1 · Detect

The requirement is not completeness, it is repeatability. A source that finds some of the cases and runs reliably beats one that finds nearly all of them and fails three times a quarter.

What gets recorded: source, query frequency, and what happens when the source fails. The third is almost always missing: a signal chain without a failure alert goes quietly empty, and that only surfaces weeks later.

Link 2 · Enrich

Enrich **only what the outreach needs.** Every additional field costs credits, runtime and error surface.

THE OUTREACH NEEDS	IT DOES NOT NEED
Role and remit of the target person	a full org chart
The specific trigger, in one sentence	the company's full history
One contact route, verified	three contact routes, unverified
One detail showing somebody looked	twenty data points nobody reads

The last row matters most. One researched detail beats ten enriched fields, because it becomes visible in the text. The recipient never sees the other nine.

Source · Clay University, the basics of the chain · university.clay.com

GO DEEPER

the chain with Clay and HubSpot

Step 3, continued: notify and hand over

Link 3 · Notify

Three questions, each with an answer that lives in operations rather than in the concept:

QUESTION	WHAT COUNTS AS AN ANSWER	WHAT DOES NOT
Who receives it?	one named person, with a stand-in	a channel where several people might read
What is in it?	trigger, person, recommended first step	the raw record
What if nothing happens?	a fallback rule with a deadline	nothing. The signal then disappears without trace

The third row is the one that is missing. A signal nobody works produces no error. In every report it looks exactly like one that was worked: the chain ran, the notification went out. So put a deadline and a second address behind it. If a signal sits for two working days it moves on, and whoever sees it then knows it is the second time.

Still link 3 · The handover

The question to settle before launch: who handles the replies? A campaign whose responses belong to nobody produces conversations nobody has. It is the most common place where a technically clean chain ends in nothing, and it costs nothing but an agreement.

Three decisions are enough:

1. **One named person takes the replies, or they get assigned.** Either works. "It will sort itself out" does not.
2. **At what signal strength an AE steps in** instead of an SDR. Without that line, either the AE always steps in or never does.
3. **Where the notification lands**, and who sees it when that person is away.

A handover without a trigger is not a handover. If the message does not say why this contact now, the first sentence of the call is a question rather than an observation. That is the same mistake as an unsorted list, one level later.

The fourth link: the campaign

The trigger belongs in the message, or the whole chain was pointless. That is the sentence separating signal-based outreach from cold outreach.

What belongs in a signal-based first message:

ELEMENT	WHY
The trigger, in one sentence, without preamble	shows this is not a mail merge
The link to your offer, in one sentence	answers "and what does that have to do with you"
A question answerable without preparation	lowers the cost of replying

What does not belong: a list of your capabilities, a calendar link in the first message, and any phrasing implying you studied the company intensively. The last one reads as dishonest on a public signal, because the recipient knows it is public.

On personalisation, because it is usually misread: a generic compliment is not personalisation, it is a recognisable template block. What works are details the recipient genuinely cares about, and a link to role and situation that is real. If the sentence reads just as well with a different company name in it, it is not personalised.

The check before sending: can the recipient tell in ten seconds why this message is going to them, and why now? If not, the trigger is missing or phrased too vaguely.

And the case that damages the chain fastest: the same signal fires twice because two sources report it, and the recipient gets two messages. A per-contact cooldown belongs in the chain, not in the team's discipline.

Which triggers signal a change in strategy

Not every signal triggers a campaign. Some trigger a review of the strategy, and confusing the two is expensive: running a campaign on a strategy trigger is activity in the wrong place.

TRIGGER	FREQUENTLY MEANS	RESPONSE
Reply rate falls across several campaigns	the audience or the approach no longer holds	review the audience, do not test subject lines
A signal suddenly delivers ten times the hits	the source changed, not the market	check the source before working the list
Meetings happen but deals do not close	the outreach reaches the wrong role	review the buying roles
A competitor newly appears in conversations	the comparison set has shifted	review positioning
One segment suddenly moves faster	something happened there that may transfer	ask, do not just scale

The second is missed most often. When a signal source suddenly delivers far more, the natural reaction is to be pleased about the volume. The appropriate reaction is suspicion: frequently a definition changed on the provider side, and the new hits are weaker.

A cadence for this, because otherwise the checks do not happen: 30 minutes once a month walking the five rows above. No more. The purpose is to ask the question at all.

Unsure whether your market carries volume?

That is the question from page 4, and it decides everything after it. It can be answered roughly within an hour, long before any tools get bought. In a free 60-minute Pipeline Generation Launchpad we size it with you and tell you whether outreach at scale is your lever at all.

pages.thesalesplaybook.com/meetings/erik-steffen/-pipeline-generation-assessment-call · no pitch, an honest assessment

Step 4: Keep the data layer readable

An enrichment table gets messy fast, and then nobody except the person who built it can say what is an input, an output, a formula or an agent result. That is not cosmetic: it decides whether the work is transferable or attached to one person.

Four measures, all cheap, all effective:

1 · Group columns and mark the output of each group. Inputs, lookups, formulas, results. Anyone opening the table for the first time sees the flow rather than forty columns.

2 · Save recurring groups as templates. A suppression check is identical in every campaign. Build once, drop in everywhere.

3 · Make every filtering step its own view. Do not filter the table, stack named views: "over 200 employees," "best segment," "send-ready." That produces a funnel where it is visible how much is lost where.

4 · Saved searches and reusable agents. A search configuration with a dozen inclusion and exclusion criteria is half an hour of work. Building it twice is half an hour too much.

The mistake that quietly burns credits

> Lookups run against the default view, not the filtered one.

The sequence: somebody carefully builds a view funnel, filters 5,000 companies down to 300 qualified ones, starts an enrichment inside that view and assumes 300 rows get processed. 5,000 get processed.

How you notice: consumption does not match the row count in the view. It goes unnoticed on small tables and surfaces on the invoice for large ones.

MEASURE	EFFECT
Check the row count an enrichment refers to before starting it	catches the case before launch
Copy qualified records into their own table rather than filtering	the default view then holds only the 300
Always start with a small test batch, then scale	limits the damage to ten rows

The third row is the base rule of any enrichment anyway: dry run, then a limited batch, look at the output, then scale.

> GO DEEPER
>
> exhaust the free tiers first

Step 5: Send-readiness, the six questions

A large list is not the same as a send-ready list. Send-ready means every row can go into a message without rework, and nobody is on it who must not be.

GRAPHIC G4

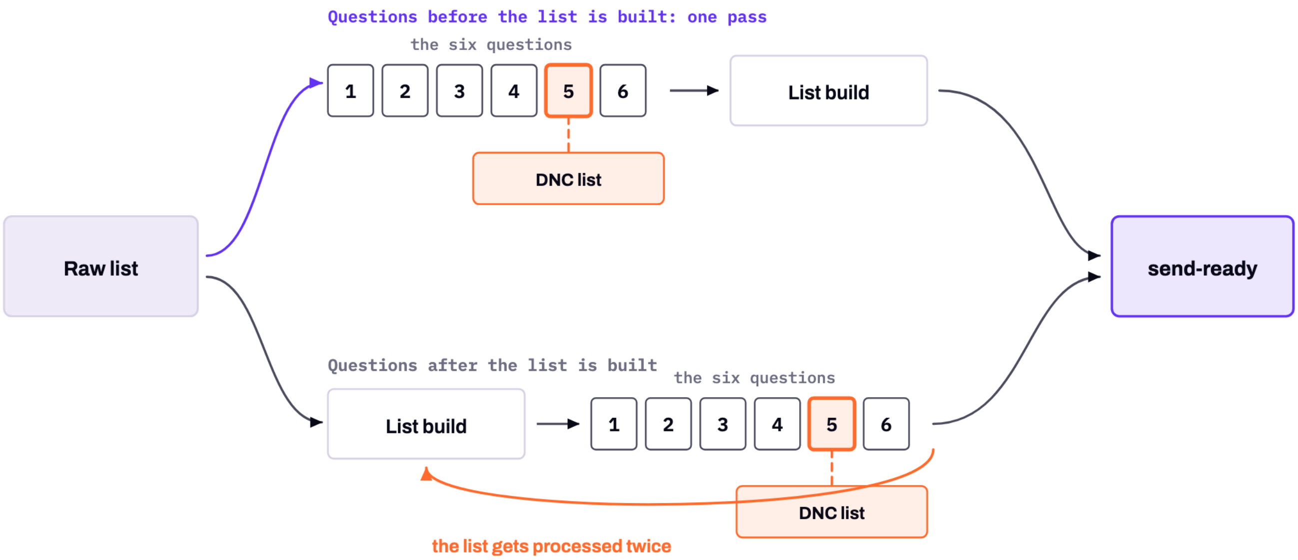

Five questions shape the row. Question 5 removes it, and it is the only one with legal consequences.

You answer the six questions before the list is built, not after. That is the whole point: applying them retroactively means processing the list twice.

#	QUESTION	WHAT HAPPENS WHEN IT IS MISSING
1	What form does the company name take in a sentence?	the full legal form mid-sentence, or a registered name nobody uses
2	First name, last name, or both? What if fields are empty?	"Hi ," in the merge
3	Does a title belong in the salutation?	in markets where the title counts, omitting it is a mistake
4	Formal or informal address?	a decision per market and audience, not per writer
5	Which suppression lists apply?	existing customers, live deals, opt-outs, competitors, your own staff
6	Cold or warm, and who works it?	the same approach to both groups costs you the warm one

Question 5 is the only one with legal and commercial consequences. A message to a live deal is embarrassing. A message to someone who has objected is a breach.

Question 1 is the most underestimated. The company name appears in almost every message, and a badly formatted one is the most visible tell of a mail merge there is.

Source · Google's sender requirements · support.google.com
Source · DMARC overview · dmarc.org
Source · connecting your inbox to HubSpot · knowledge.hubspot.com

GO DEEPER

settle tracking and logging first

The checks that run before the first send

Six answered questions are the policy. These six checks are the implementation, and they run on every list.

EVIDENCE why one check is not enough

#	CHECK	RULE
1	Apply the suppression list	always ask where it lives and what is in it. The rest runs even without one
2	Names complete and in the right field	swapped first and last names surface when you run the forename check on **both** fields
3	Normalise the writing	strip titles and credentials from the name field. Fix all-upper or all-lower casing. Do not force-capitalise name particles
4	Complete abbreviated names	derive from the local part of the address first, only then research
5	Check the match	if the first name fits but the surname starts with a different letter than the source, the address may belong to a different person
6	Verify addresses twice	once at enrichment, once in the sending tool immediately before launch

The rule that holds it together

> **Flag and hold, never silently delete and never guess.**

Any row where something stays unclear goes on a hold pile: out of the active send, not deleted, released once somebody has resolved it. Deleting destroys work. Guessing destroys trust with the recipient.

A threshold that has proven itself: an automatically derived attribute is only accepted when it is **highly** confident. Anything below stays blank and goes on the hold pile. A blank field is visible and gets filled in. A wrongly guessed salutation ships and cannot be recalled.

And one figure to know before you start: the share of addresses that come back only as "cannot be verified conclusively." They are sendable, but their share belongs reviewed and deliberately decided before launch, rather than quietly eroding your deliverability.

Source · setting up subscription types · knowledge.hubspot.com

GO DEEPER

the last checks before sending

The hygiene profile as a document

The answers to the six questions belong in a document that applies to every campaign, not in the head of whoever is building the list.

SECTION	CONTENT
Name form	Rule for the short form of the company name, with three examples
Salutation	First/last name, title rule, fallback for empty fields
Address form	Formal or informal, per market
Do Not Contact (DNC)	Which sources, how often reconciled, who maintains them
Handling	Cold and warm separated, who works which
Language and format	Date, number and currency format per market

Why a document rather than a convention: a convention holds while the same person builds the lists. At the first handover or the first outsourcing it fails, and the error ships to recipients.

The last row is the case nobody thinks about. A date written 09/01/2026 means two different days in two markets. In a message proposing a meeting time, that is not cosmetic.

Check before every send: pull three rows at random and look at the message as the recipient gets it. Not the template, the rendered version. Most formatting errors are invisible in the template.

What gets automated and what a person decides

AUTOMATE	A PERSON DECIDES
Detect signal and create the record	whether this signal counts in this market
Enrich against fixed fields	which fields are needed
Reconcile suppression lists	what belongs on the suppression list
Format and populate templates	the template itself
Notify and assign	whether the trigger supports outreach
Follow up on schedule	whether to follow up at all

The compromise that holds: the system builds the draft completely, a person looks at it for ten seconds and approves or discards. **That is not proofreading.** Anyone rewriting every draft has manual work with an extra step, and in that case the template is wrong.

Two evidenced results, and the checklist

Two published cases from our own work, quoted in the wording of the case study itself.

node.energy Generates 5x More Demos in Three Months Through Targeted Outbound With Clay Integrated Into HubSpot. Source: thesalesplaybook.com/clients/node-energy

From 700K Impressions to C-Level Meetings to closing a 400k Deal in 4 Months with a LinkedIn GTM Motion. Source: thesalesplaybook.com/clients/calibo

Both show that the mechanics described led to those results in one specific engagement. They are **not a commitment:** the starting position decides, and it differs by company.

Checklist

- [ ] Reachable market quantified before tools were chosen
- [ ] Decided whether the market carries volume or demands depth
- [ ] Sourcing method matched to where the segment is findable
- [ ] Market and campaign list kept separate
- [ ] Sample of ten to twenty companies per segment fully enriched
- [ ] Three to five signals selected, each observed rather than inferred
- [ ] Window and monthly hit count known per signal
- [ ] Signals defined with sales, five won deals walked backwards
- [ ] Discarded candidates documented with reasons
- [ ] All four chain links named, with ownership each
- [ ] Failure alert configured for the signal source
- [ ] Notification goes to a person, not a channel
- [ ] Settled who handles replies and when an AE steps in
- [ ] Trigger appears in the first message, in one sentence
- [ ] Per-contact cooldown in the chain, not in team discipline
- [ ] Column groups and views built, output marked per group
- [ ] Row count checked before every enrichment
- [ ] Test batch first, then scale, every time
- [ ] Hygiene profile written as a document, six questions answered
- [ ] Six list checks running, hold pile instead of deletion
- [ ] Addresses verified twice, share of unverifiable ones known
- [ ] Automation boundary drawn, approval step defined
- [ ] Monthly 30-minute look at the strategy triggers

NEXT DECISION

whether the pipeline covers the quota

Next steps

The place this most often stalls is not the technology. It is two questions before it: how large your reachable market actually is, and which three signals inside it genuinely mean something. The second cannot be answered from data, it comes from the deals you have already won.

Where this document reaches its limit: if the market does not carry volume, outreach at scale is the wrong answer, and no tool changes that. That is an insight rather than a setback, and it is cheaper before the build than after.

In a **free 60-minute Pipeline Generation Launchpad** we look at your situation: which signals are candidates, whether the chain is buildable on your current stack, and where it would break. Including when the answer is that outbound is not your lever.

Free, 60 minutes, no pitch, an honest assessment of whether there is a fit.

WHO YOU WILL MEET IN THE PIPELINE GENERATION LAUNCHPAD

Florian Lussi

Pipeline Generation Lead & Partner · SalesPlaybook

Book your free 60-minute Pipeline Generation Launchpad → pages.thesalesplaybook.com/meetings/erik-steffen/-pipeline-generation-assessment-call

Fill in the self-assessment first → form.jotform.com/240112515833345

Sources quoted

The statements quoted verbatim come from the public LinkedIn newsletter *Diary Of A CRO* by Manuel Hartmann, founder and CEO of SalesPlaybook (https://www.linkedin.com/newsletters/diary-of-a-cro-7463115808336986112/). The issues are in German:

- "„Show me your customers and your future."" · https://de.linkedin.com/pulse/show-me-your-customers-future-manuel-hartmann-8jfze
- "10 Hausregeln für Vertrieb mit Verstand im Zeitalter von KI" · https://de.linkedin.com/pulse/10-hausregeln-f%C3%BCr-vertrieb-mit-verstand-im-zeitalter-von-hartmann-wzlee

Objections quoted verbatim come from our own project conversations and are anonymised: no client name, no industry, no figures.

DOCUMENTED OUTCOMES · 75 SUCCESS STORIES

PIPELINE GENERATION · CALIBO · 4 MONTHS

From 700k impressions to C-level meetings to a 400k deal in four months

PIPELINE GENERATION · KNK · LESS THAN 6 MONTHS

From manual outreach to an AI-automated sales pipeline in less than six months

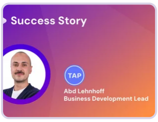

PIPELINE GENERATION · TAP · LESS THAN 5 MONTHS

How TAP generated 2x qualified pipeline in less than five months

All success stories: thesalesplaybook.com/clients

SalesPlaybook designs, builds and operates Pipeline Generation, HubSpot CRM and Revenue Enablement as one connected revenue system. An operating partner for revenue systems, HubSpot Diamond Partner, working with B2B technology companies with 25 to 500 full-time employees in the DACH region.

A workbook by SalesPlaybook AG

Free 60-minute Pipeline Generation Launchpad: pages.thesalesplaybook.com/meetings/erik-steffen/-pipeline-generation-assessment-call