A PLAYBOOK FOR SCALING B2B FINTECHS

The GTM Operating System

How scaling B2B fintechs turn complexity into disciplined execution.
Strategy, domain expertise and AI, moving as one.

RedSnap
The fintech agency · redsnap.nl

Contents

The AI value gap

Two big enterprise AI studies came out in the back half of 2025. McKinsey surveyed nearly 2,000 companies. 88% of them are using AI somewhere in the business. Only 5.5%, one in twenty, can attribute more than 5% of EBIT to it. BCG's study of 1,250 executives found the same shape from the other side: 5% of companies qualify as "future-built" for AI, and those companies are seeing 1.7x the revenue growth of their peers, 3.6x the three-year shareholder return, 1.6x the EBIT margi.

Most of the early reaction to numbers like these blamed the technology. We don't buy it. In almost every B2B project we've seen go sideways, the AI wasn't the issue. It got bolted onto a go-to-market that wasn't coherent in the first place. A strategy that never made it into the CRM. Domain knowledge that lived in one person's head. Workflows firing into the dark.

Who this is for

You're a CCO, CMO, CRO or RevOps lead at a scaling B2B fintech. The board has started asking, politely or otherwise, why last year's AI spend hasn't shown up in pipeline.

Inside: a model of what a working AI-enabled GTM engine looks like, a diagnostic for finding where yours is lopsided, and a 90-day plan for getting the first loop live.

AI doesn't fix a broken go-to-market. It amplifies one.

The flip side of the story is just as quantifiable. Gong's analysis of 7.1 million sales opportunities, published December 2025, found that sellers using AI well generate 77% more revenue per rep than peers who don't. The companies on the right side of the gap aren't doing more AI. They've rebuilt the engine around it.

What GTM complexity actually is

Every scaling fintech hits a version of the same wall. Up to about fifty people, everything fits in the founder's head. Between fifty and two hundred, the commercial engine is a set of relationships and a spreadsheet. Past two hundred, in multiple markets, with a buying committee of six-plus stakeholders and a regulator per country, it stops being a set of relationships and starts needing to be a system.

That transition is what we mean by GTM complexity. It isn't a tooling problem. It's the point at which strategy, execution and infrastructure stop moving as one.

Three symptoms every operator recognises

Pipeline you can't trust. The forecast review becomes an argument about numbers instead of a conversation about accounts. Dashboards multiply. Confidence in any one of them shrinks.

Tools that multiply, value that doesn't. The typical scaling fintech now runs thirty-plus revenue tools. Six of them claim to be "AI-powered." Three overlap. Two contradict. The stack grew faster than the operating model.

Agents that add cognitive load. You deployed an AI SDR, a content agent, an enrichment layer. The reps now spend more time curating AI output than they spent writing emails before. Someone quietly says, "we were faster without it." And they're right. Not because the agent is broken, but because nothing told it what good looks like.

If any of those sound familiar, the instinct is to reach for another tool. It won't help. The problem isn't the tools. It's that the three things that should hold the engine together (a clear strategy, real domain expertise, and disciplined automation) aren't aligned.

From stack to system

The typical scaling fintech runs 30+ revenue tools. Six claim to be AI-powered. Three overlap. Two contradict.

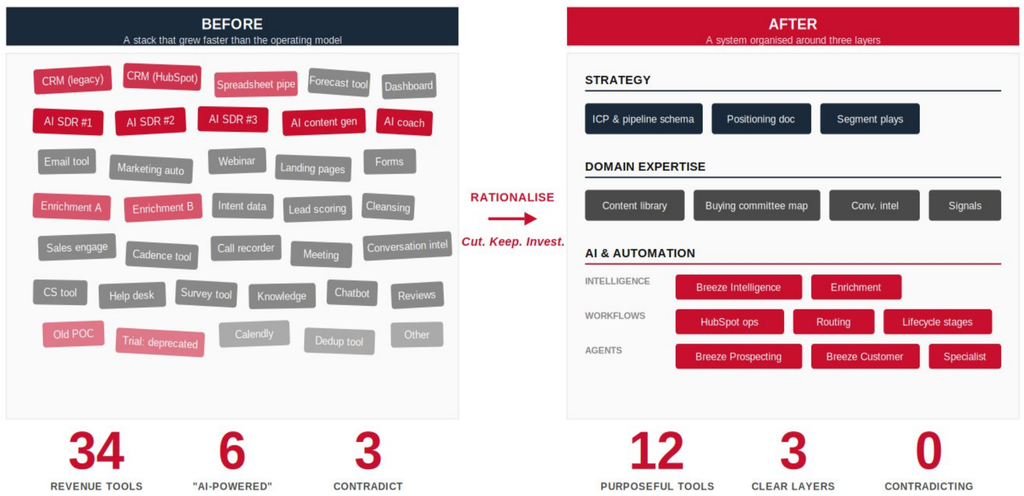

Rationalisation isn't about adding intelligence. It's about removing the noise that drowns it.

The GTM Operating System

A three-layer thesis

There's a place where strategy, fintech domain expertise and intelligent automation meet. It's not a buzzword and it's not a slide. It's an operating system: the set of decisions, assets and loops that let a scaling commercial engine run coordinated, repeatable and built to scale.

Each layer on its own is insufficient. Strategy without expertise is theory. Expertise without automation doesn't scale past a few heroic operators. Automation without strategy confidently amplifies the wrong thing at speed.

The GTM Operating System

Where strategy, fintech expertise and intelligent automation meet.

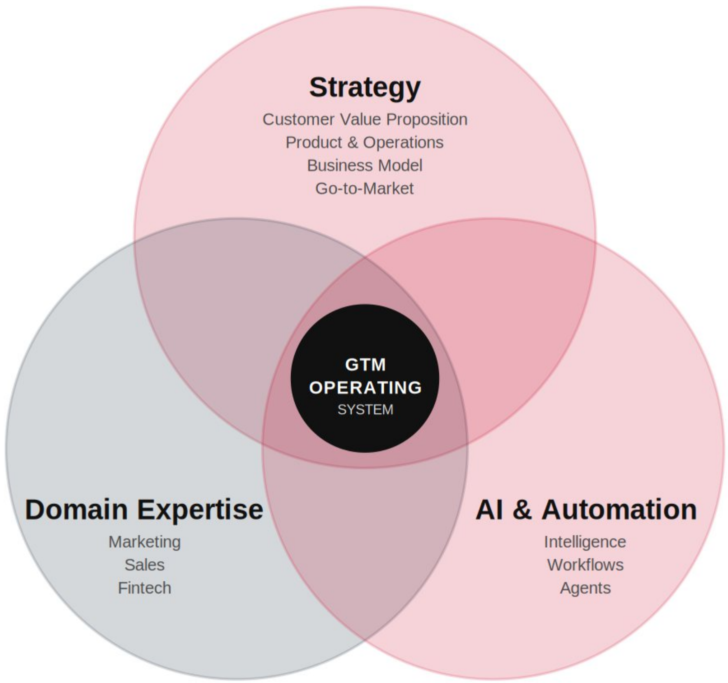

Each layer alone is insufficient. The intersection is the operating system.

What sits where

Layer	What it contains	What happens without it
Strategy	Customer value proposition Product & operations Business model Go-to-market	You automate the wrong motion faster
Domain expertise	Marketing Sales Fintech-specific patterns	Generic AI produces hyper-personalised noise
AI & Automation	Intelligence Workflows Agents	Expertise stays trapped in a few people's heads

The point

The GTM Operating System isn't the sum of these layers. It's the intersection. The teams winning in 2026 are the ones where a strategic decision on Tuesday shows up as a changed field in the CRM on Wednesday and a retuned agent on Thursday, and the team can tell you why.

The diagnostic

Where is your engine lopsided?

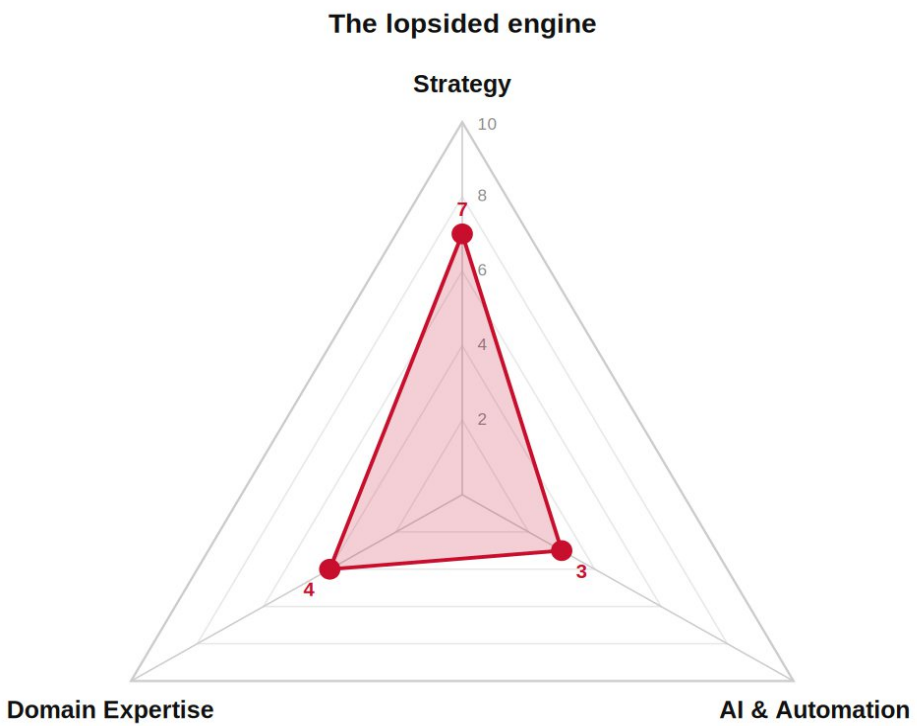

How to read it

Score 0–10 per layer.
Lowest layer is
where you start.
Widen the narrowest.

A typical scaling fintech: strong strategy, narrow expertise, weakest in automation.

Run through the fifteen statements below. Score each one honestly: 0 (no), 1 (partially), 2 (yes, with evidence). Tally by layer. The lowest-scoring layer is where you start. Not the most interesting one.

Strategy

1. Our ICP is written down and any new hire in sales or marketing can name the top three disqualifiers.
2. Our value proposition is segment-specific. We say something different to a Tier-1 bank than we do to a neobank.
3. Positioning is encoded in CRM fields, not just on a slide. Lifecycle stages and scoring reflect it.
4. We can draw our business model, pricing and unit economics on a whiteboard in under three minutes.
5. Our go-to-market motion is named: PLG, sales-led, partner-led or hybrid. The team agrees.

Domain expertise

1. We have documented the buying committee for our top two segments, including compliance, InfoSec and legal.
2. Our sellers know the three fintech-specific objections most buyers raise in the first thirty minutes, and have defensible answers.
3. Our marketing distinguishes regulated-buyer messaging from (fintech) builder-buyer messaging. They are not the same.
4. Compliance and data residency are positioned as GTM assets, not legal disclaimers at the bottom of the deck.
5. We have a referenceable point of view on our market: one a CEO would quote in a board meeting.

AI & Automation

1. Our CRM data is clean enough that an AI agent could make a trustworthy next-best-action recommendation today.
2. Our core workflows (lead routing, lifecycle, renewal, escalation) are documented, not tribal.
3. We can name the top five signals that predict a deal advancing, and they fire somewhere in our stack.
4. We know which Breeze agents we have live, which we've rejected, and why.
5. We have a governance answer for an AI agent that says something wrong to a regulated customer.

The pattern we see most often is a 6 or 7 on Strategy, a 4 on Domain Expertise (because it lives in two or three people), and a 3 on AI & Automation. That's the layer the company is trying to fix first. That's backwards. You don't climb the stack. You widen the narrowest layer.

How to read your score

0–3 per layer: foundational work required. Don't add AI here yet.

4–6: functional. One layer down, focused investment will unlock the others.

7–8: competitive. Rare. Protect it.

9–10: either world-class or over-scoring. Ask a skeptical outsider to re-score you.

Layer 1: Strategy, made operational

Strategy is the layer most companies think they've handled. They haven't. They have a slide. Strategy is operational only when it's been translated into the structure of the commercial engine: which fields exist in the CRM, which lifecycle stages a contact can move through, how leads are scored, what "qualified" actually means for each segment.

For a fintech selling into both Tier-1 regulated banks and fast-moving neobanks, this isn't a theoretical point. The two buyers have nothing in common. A Tier-1 bank moves through a 6 to 12 month cycle with compliance, InfoSec, legal and CISO sign-offs. A neobank can evaluate an API in a weekend. If they share a pipeline stage with the same conversion expectation, every forecast you build is wrong.

What translating strategy into the engine looks like

Segment the pipeline. At minimum, split regulated-buyer motions from builder-buyer motions. Different stages, different SLAs, different scoring.

Encode the ICP as data. Not a paragraph in Notion. A set of fields (regulator, geography, scheme membership, program type, team maturity) that lets any AI agent reason about a lead instead of guessing.

Name the disqualifiers. The three types of deals you should refuse. Put them in the CRM. Reward the team for the "no" as much as the "yes."

Make positioning a field, not a slogan. If your differentiation is "cloud-native issuer-processor for emerging markets," then "emerging market" has to be a property with a clear definition, not a vibe.

This is most of what we do in the first thirty days with a new client. We bring a fintech-specific ICP schema into HubSpot: regulator, scheme membership, BIN sponsor, program type, geography, team maturity. We split the pipeline into the regulated and builder tracks with their own stages and scoring. We name the three disqualifiers. None of this is exotic. It's the work most engines have skipped because the slide felt like the deliverable.

Field Note: NorthArc Pay, month 1

NorthArc Pay is a global issuer-processor, ~400 people across 60+ countries. A new CCO arrives. In the first weekly pipeline review he notices a Tier-1 bank RFP and a neobank API trial sitting in the same "Qualified Opportunity" stage with the same 30% weighting.

The engine is telling him they're the same kind of deal. They're not. The first move isn't a new tool. It's re-architecting the pipeline so the engine stops lying to him.

Layer 2: Domain expertise, encoded

This is the layer most at risk of being invisible. It tends to live in two or three people (the founder, the head of sales, one SME) and it leaves with them. Encoding it is what turns a heroic commercial engine into a repeatable one. It's also the single biggest lever for making AI useful instead of embarrassing.

Generic AI misses the intricacies related to fintech. It doesn't know that "BIN sponsor" isn't a synonym for "payment partner." It doesn't know that messaging a compliance officer about "move fast and break things" will end the deal. It will write a perfectly personalised email that's technically correct and commercially fatal.

Five fintech-specific patterns worth encoding

→ **The regulated-buyer committee.** Six-plus stakeholders, with legal, compliance, InfoSec and procurement often outweighing the economic buyer. Messaging and content have to be built for each of them.

→ **Compliance as a GTM asset.** Certifications, audits, data residency and scheme membership aren't footnotes. They're top-of-funnel proof. Put them in the lead gen, not the MSA.

→ **Trust before evaluation.** In fintech, a buyer decides whether to evaluate you before the demo. The first three touches carry more weight than the next thirty.

→ **Partner-led motion physics.** Integrators, schemes, BIN sponsors, banking partners etc ... are distribution channels. They need enablement, not just a contract.

→ **Signal decay in long cycles.** A "hot" signal on a 6-month cycle is almost worthless at month 4. Most prospecting tools don't know this.

Most of our middle months with a client are spent here. Encoding the five patterns above into the engine.

Buying-committee maps as HubSpot records. Compliance and scheme membership as searchable properties, not PDFs in a shared drive.

Signal decay rules wired into automations so a stale lead stops firing alerts the team has learned to ignore.

Two or three people's tacit knowledge becomes something the next ten hires can use on day one.

Field Note: NorthArc Pay, months 2–3

The marketing team had one content library. It spoke to everyone, which meant it spoke to no-one. In 60 days, the team splits it: regulated-buyer content (compliance depth, scheme partnerships, references from Tier-1 banks) on one track; fintech builder-buyer content (API docs, time-to-first-card, developer testimonials) on another.

The enrichment layer is reconfigured so a lead flagged "regulated" never receives the "ship your card in 72 hours" nurture. What looked like an AI problem was a segmentation problem.

Layer 3: AI & Automation, disciplined

Only now does this layer become useful. With Strategy operational and Domain Expertise encoded, AI and automation have something real to work with. Before that, they're a loud way to amplify confusion.

There are three sub-layers, and they need to be built in order. Skip a step and the one above collapses.

1: Intelligence

Data an agent can reason over. In a HubSpot-centred stack, that usually means Breeze Intelligence (intent and enrichment), a cleaned and deduplicated contact database, and a defined set of signals: hiring moves, funding events, product launches, partnership announcements, scheme membership changes. Intelligence is the input. If it's dirty, everything downstream lies.

2: Workflows

The ten workflows every scaling fintech should automate first are unglamorous: data hygiene, lead routing, lifecycle progression, MQL-to-SQL handoff, opportunity stage transitions, forecast roll-up, renewal 120-day flag, NPS follow-up, churn-risk escalation, executive-level win/loss capture. None involve a chatbot. All make the chatbot useful when it arrives.

3: Agents

Only now do agents become relevant. In a HubSpot-native stack, Breeze Agents (Customer, Prospecting, Content, Data, etc..) or MCP, cover almost everything most scaling fintechs need. The mark of a mature operator isn't which agents are deployed. It's which ones the team has consciously decided not to use yet.

Breeze does the in-CRM work well. The harder questions sit outside the CRM: which fintechs are quietly in-market for a new issuer-processor, which are about to change BIN sponsor, which compliance officers are speaking at next month's industry event. That's what we built our own agents for. IMA, our in-market agent, watches the external fintech signal landscape (scheme announcements, regulator filings, partnership news, hiring patterns at known buyer accounts) and feeds the signals back into HubSpot with decay rules attached. Hubspot's Prospecting Agent acts on what IMA surfaces. The two are designed to run together.

The context gap, and how to close it

Breeze is native to the CRM, and that's its superpower and its limit. Agents that can only see CRM data can't answer questions that need product telemetry, support history, billing data or the regulatory calendar. The discipline is deciding, case by case, which questions are CRM-shaped and which require closing the context gap through integration.

In fintech, the context gap is usually outward-facing: signals that live in the regulator's feed, the scheme network, the partner ecosystem. IMA fills exactly that gap for our clients, and the resulting in-market signal lands in HubSpot as a property Breeze can reason over. The principle generalises. Every mature AI-enabled GTM engine has at least one purpose-built agent watching the part of the world the CRM can't see.

Field Note: NorthArc Pay, months 4–6

The team deploys Hubspot's Prospecting Agent first, against the neobank segment, where the data is cleanest and the cycle is shortest. Early results are good.

They try to extend it to the Tier-1 bank segment and it falls over. The agent can't see scheme relationships, doesn't know who the BIN sponsor is, and produces a message that confuses "cloud-native" with "unregulated." The team doesn't blame the agent. They close the context gap with IMA: scheme, sponsor and other relevant signals from the external fintech world land in HubSpot as properties, decay rules attached. Hubspot acts on what IMA surfaces.

From month 6 it works.

The 90-day build sequence

A working GTM Operating System isn't built in a sprint, but the first loop should be live in 90 days. If it isn't, the project drifts. Here's the sequence that works.

The 90-day build sequence

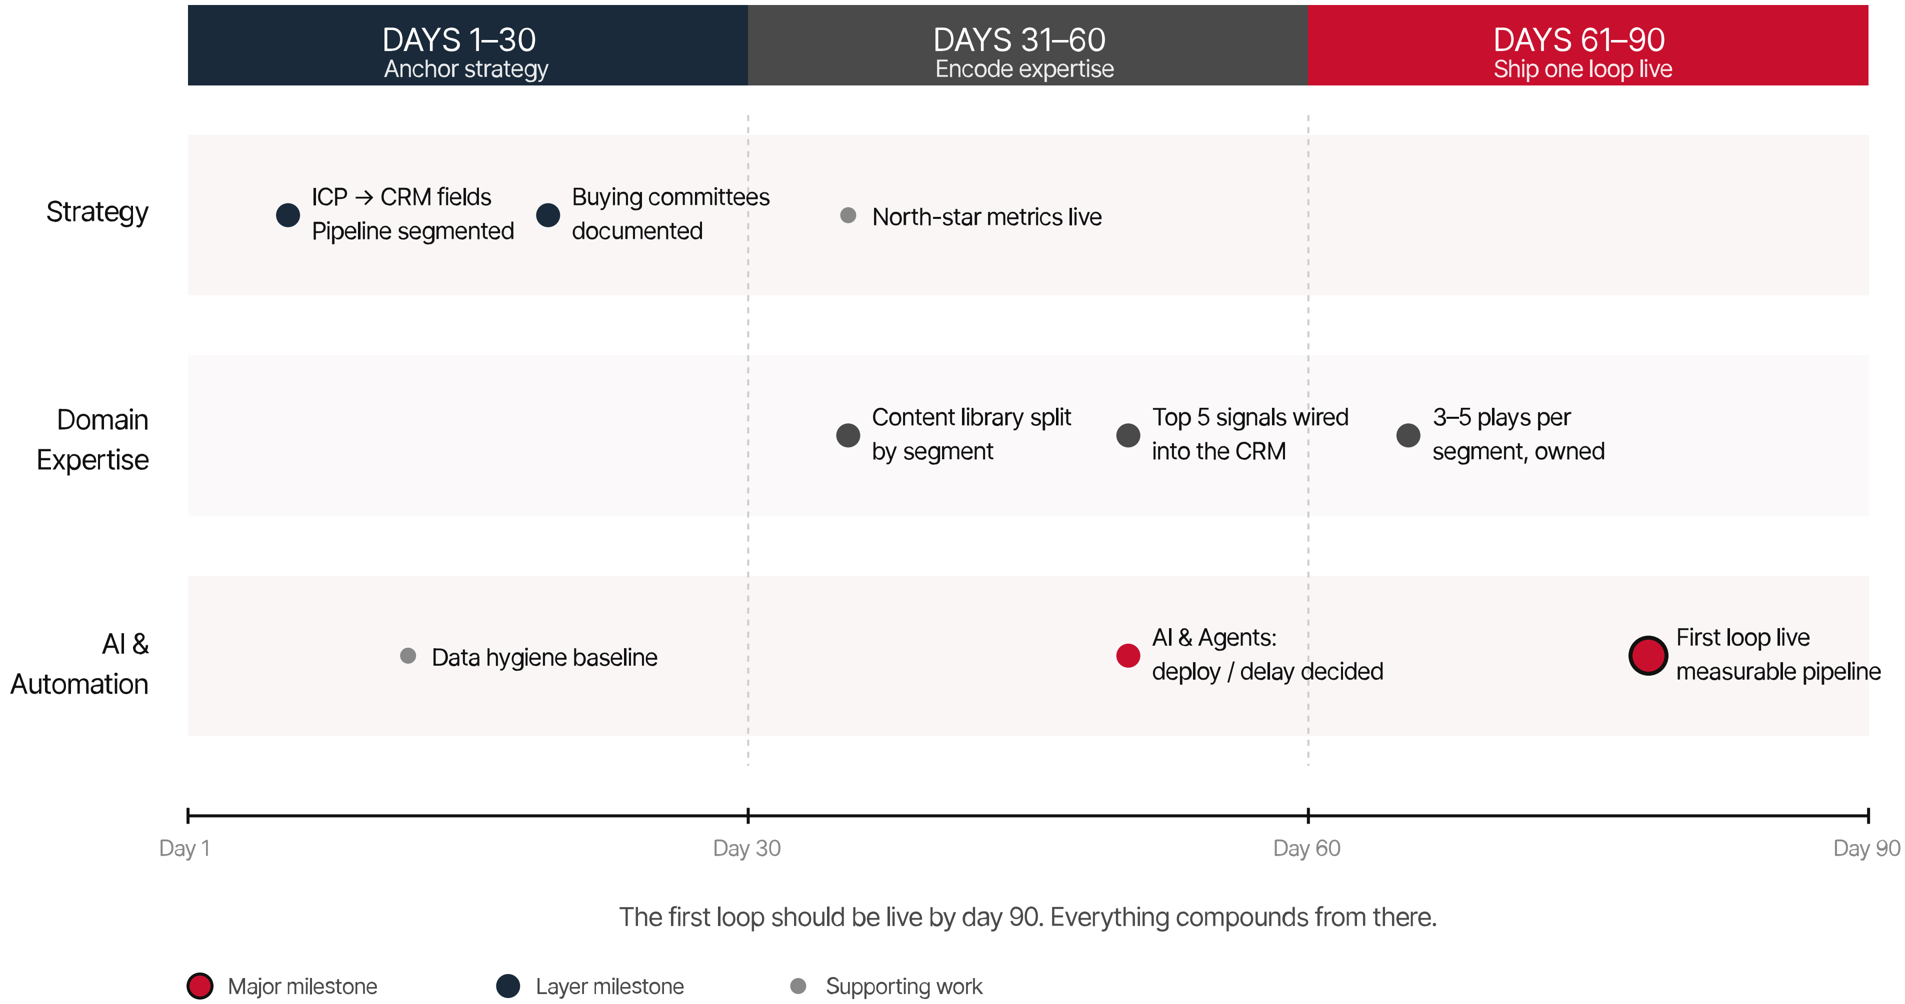

Days 1–30: Anchor strategy

- → Re-ground positioning and ICP. Segment the pipeline: regulated vs. (fintech) builder, at minimum.
- → Translate the ICP into CRM fields. Retire every field nobody has written to in 90 days.
- → Document the top two buying committees end to end, including the people who can kill the deal.
- → Agree one north-star metric per layer. Pipeline quality for Strategy. Segment message fit for Domain Expertise. Loop latency for AI & Automation.

Days 31–60: Encode expertise

- → Split the content library along segment lines. Kill anything without an owner.
- → Name the five most valuable signals and wire them into the CRM.
- → Document three to five “plays” per segment: what triggers them, who runs them, what good looks like.
- → Decide which agents you’re deploying, which you’re delaying, and write down why for each.

Days 61–90: Ship one loop live

- → Pick the narrowest possible scope: one segment, one motion, one signal, one play.
- → Deploy one agent against it. Measure the loop end to end, not the agent’s output.
- → Run a weekly review focused on whether the loop is learning, not whether the agent is fast.
- → At day 90, the artefact is one referenceable play with measurable pipeline contribution. That’s the proof point. Everything else compounds from it.

Five anti-patterns to refuse

Five patterns we see kill AI-enabled GTM projects at scaling fintechs. None are about the technology. All are about the operating model.

1 **Automating before positioning.**
If you can't articulate who you're not for, automation will find more of the wrong prospects faster.

2 **Buying agents before cleaning data.**
IMA, the Prospecting Agent, the Data Agent, the Content Agent all assume the underlying data is clean. It almost never is. Clean first.

3 **Measuring throughput instead of decision quality.**
"The AI sent 3,000 emails this week" isn't a number that matters. "Reply rate on the regulated segment went from 1.2% to 3.8%" is.

4 **Forecasting without acknowledging signal decay.**
Long fintech cycles need signals with shelf-lives. Treat them like milk, not canned goods.

5 **Treating Customer Success as a hand-off.**
Net revenue retention is decided in the first 90 days of a customer's life. If CS gets the customer after the close, the compounding's already gone.

Where this leaves you

The value gap at the start of this playbook isn't a technology story. It's an operating model story. The teams in the 5% aren't the ones with the biggest AI budget. They're the ones where strategy, domain expertise and automation move as one.

That intersection is where the next two years of compounding commercial advantage are built. The playbook above is the shape of it. The diagnostic in section 04 tells you where your engine is lopsided. The 90-day sequence in section 08 is how you ship the first loop.

If you work through the diagnostic and the lowest-scoring layer is AI & Automation, the instinct is to buy an agent. Don't. Widen the narrowest layer first. The compounding only works when all three move together.

Where RedSnap fits

We operate at the intersection: strategy, fintech domain expertise, and AI & automation inside HubSpot. Not as three separate capabilities. As one operating model, for teams that want growth done properly.

If strategy, expertise and scale need to move as one, we should talk.

www.redsnap.nl · The fintech agency

A note on NorthArc Pay. The story woven through this playbook is a composite drawn from scaling global issuer-processors we've worked alongside. The two-segment pipeline architecture, the content library split, the IMA & Prospecting Agent deployment sequence: those are the real ones. The name is the only fiction.

Sources cited: McKinsey, State of AI 2025 (November 2025). BCG, Build for the Future 2025 (September 2025). Gong, State of Revenue AI 2026 (December 2025).