

## HOW GTM RADAR IS BUILT

# Methodology & Decision Log

*The buckets, the filters, the scoring, and every pivotal scope call made during the build.*

GTM Radar is a research workflow, not an automated pipeline. Its quality depends on what gets onto the watchlist and what gets thrown out. This page documents both — the operating logic and the decisions that shaped it. Read it as the operator's manual.

## Six buckets, one filter

AI is the company filter, not the role filter. A VP Sales hired at a company racing to ship AI is the briefing's target. A 'Head of AI' or 'VP of AI' is not. The six buckets organize the watchlist around where Betts's book actually lives, with a bias toward AI-late incumbents (the largest opportunity) over AI-native cohorts (the smaller, more visible one).

- 01 AI-Native (15 cos)** — AI is core product architecture. Harvey, Writer, Glean, Cresta, Observe.AI, Electric, Gong, Notion, plus seven scaling cohorts in customer support / knowledge work / legal.
- 02 RevTech (12 cos)** — sales / marketing / RevOps tooling. Salesloft, Outreach, Clari, 6sense, Demandbase, Apollo.io, ZoomInfo, Common Room, UserGems, Clay, Sybill, Lavender.
- 03 Fintech + Cyber (10 cos)** — adjacent sectors where AI is reshaping the GTM motion. Ramp, Brex, Mercury, Rho, Abnormal, Vectra, Dropzone, Prophet, Sysdig, Red Points.
- 04 AI-Late Adopter (16 cos, biggest bucket)** — mid-cap SaaS being pushed to adopt AI. Zendesk, DocuSign, Box, Dropbox, Twilio, Intercom, Okta, Atlassian, Asana, Monday, Smartsheet, Zoom, Samsara, Klaviyo, PagerDuty, Procure.
- 05 Incumbent Context (2 cos, context-only)** — HubSpot, Gainsight. Tracked for market awareness; not Betts search targets.
- 06 Pre-Discovery (8 cos)** — vertical-AI companies that match the Betts ICP but sit off the conventional GTM-recruiting beat. Trunk Tools, Mendel.ai, Sixfold, Cleric, 11x, Nooks, Tessel, Speak. The strongest version of this product surfaces deals the firm doesn't know about yet.

## Three filters decide what's on the list

All three must be true. The filters are conservative on purpose: they keep the watchlist defensible and the briefing density high.

- 01 Active GTM hiring need.** Series B through pre-IPO, or mid-cap public / PE-backed. No pre-PMF.
- 02 No internal recruiting army or brand magnet.** Anthropic gets 50K inbound applications. The pain has to be real for Betts to win.
- 03 Matches Betts's documented pattern.** Tech / SaaS, GTM-heavy, volume hires or exec search. Anchored in public Betts case studies — Harvey, Zendesk, Grammarly, Sysdig, Tray.io, Electric, Attentive, Staffbase.

## Scoring rubric

Every signal is scored 1–5. Default cutoff for inclusion in the briefing is  $\geq 3$ . The scoring is a forward-looking hiring-likelihood read, not a press-release weight.

SCORE	LABEL	PREDICTS	EXAMPLE
5	Hot	Hires within 30 days	Late-stage raise + AI-native + active reqs
4	Warm	Hires within 60 days	Recent round or new exec at scaling AI co
3	Watch	Hires within 90 days	Product launch or early-stage raise
2	Background	No direct GTM read	Monitor — does not enter the briefing
1	Noise	Drop	Wrong sector, wrong stage, or unverified

## Pivotal scope decisions

Every entry is a real decision made during the build. Each changed what gets included or how the briefing reads.

- 01 **Cut hyperscalers and brand magnets.** Anthropic, OpenAI, Perplexity, Cursor, Cognition, Poolside, Magic excluded from v1. They don't need Betts. Healthcare AI scoped out — different recruiter network.
- 02 **Re-classified Gong and Notion to AI-Native.** Both have AI as core product architecture, not a late-added layer. Moving them out of incumbent-pivot sharpened the bucket distinction.
- 03 **Defined the role filter.** AI is the company filter, not the role filter. 'Head of AI' / 'VP of AI' / 'AI Product Manager' explicitly excluded — those aren't Betts's book. A VP Sales at the same company is.
- 04 **Added 'too big for Betts' exclusion.** Salesforce and ServiceNow removed. Internal TA armies at scale.
- 05 **Created 'context-only' bucket.** HubSpot and Gainsight tracked for market-pitch awareness but not as search targets. Their events appear in TL;DRs / sidebars, not as full signal cards.
- 06 **Expanded watchlist by 12 cos after a quiet W17.** First two scans returned zero on-target signals. Added seven AI-native scaling cohorts (Sierra, Decagon, Hebbia, Forethought, Maven AGI, Crescendo, EvenUp), three RevTech (Clay, Sybill, Lavender), two AI-late adopters (PagerDuty, Procure). The expanded scan surfaced four verified launches in the W17 window.

## The verification rule

Every signal traces to a primary source — company press release, blog, SEC filing — or a tier-1 secondary (TechCrunch, Reuters, WSJ, Information, Bloomberg) within the published week's window. The brand of this artifact is 'never hallucinates.' One fabricated signal kills the credibility of the briefing.

### ~~Rho Series B~~

Page-update on a 2024 announcement

### ~~Apollo \$110M~~

Conflated with the 2022 raise

### ~~Abnormal FedRAMP~~

2025 announcement — outside window

### ~~Vectra-Phillips CRO~~

Dec 2025 — outside window

### ~~Mercury \$5B~~

Rumor — no primary source

### ~~Zendesk AI rollout~~

Mar 30 — outside window

6 REJECTED · 5 PUBLISHED · THE REJECTED LIST SHIPS INSIDE EACH BRIEFING.

## What this could become

Three directions worth pressure-testing if there's appetite to extend the work. None are built yet. Each is a different shape of bet on what makes recruiter intel actually useful at the desk.

- 01 **Two more verticals.** Extend coverage to cohorts Betts already places into but the current sources don't reach — vertical SaaS (Toast, Procure, ServiceTitan-adjacent), developer infra (Vercel, Supabase, Resend), or whichever cohort sits closest to your active searches.
- 02 **An interactive query layer.** A recruiter pulls up a target company mid-call and asks 'who's likely available' or 'what's the hiring read here?' Not a dashboard — a conversational query against the

same research substrate these PDFs are built on.

- 
- 03 Backtest the rubric.** Take six months of confirmed Betts placements, replay the watchlist as if those placements hadn't happened yet, and check whether the 1-5 scoring would have flagged them in the right window. Turns the 30/60/90-day claim from a hypothesis into a measurement.
- 

**WHERE I'D LIKE FEEDBACK**

*Of the three directions above, which one would actually be useful to a Betts recruiter at their desk on a Monday morning? Or — what's missing? The PDFs were the easy part. Building the next version against a real read beats another guess.*

**Kyle Findlay**

kylefindlay828@gmail.com