THE $0 → $20M ARR PLAYBOOK

Sales & GTM for AI Scale-ups.

What 26 years of enterprise sales still teaches, and what had to be rebuilt AI-native after November 2022.

Andrew Wesbecher

linkedin.com/in/wesbecher · awesbecher@gmail.com

CRO / VP SALES · AI & AI SECURITY

TRACK RECORD

26
YEARS ENTERPRISE GTM

6
GTM ORGS BUILT & SCALED

10
GTM PLANS FOR AI COMPANIES

MEDDPICC to AEO.
Pipeline math to agent workflows.

00 · WHY THIS EXISTS

ChatGPT changed everything. The math held.

I have carried a bag and built GTM engines for 26 years: TIBCO, Meraki, ThousandEyes, Lacework, Contrast, Traceable. The last three years, I have run AI and AI security scale-ups as a fractional CRO and advisor, with an AI-first GTM stack doing the ops work.

This playbook is what survived the collision: the enterprise foundations that still close seven-figure deals, and the modern layer that AI-era buyers now expect.

Nothing in it is aspirational. It is what I run now.

26	10	4
YEARS ENTERPRISE GTM	GTM PLANS, AI COS	AI SECURITY CLIENTS

THE PLAYBOOK

WHO THIS IS FOR

Founders and boards scaling $0 → $20M ARR. Every framework here has shipped inside a real plan.

Your buyer moved. Three years ago.

01 **Research happens in the answer box.** Buyers ask ChatGPT, Claude and Perplexity who the leaders are before your SDR ever calls.

02 **Adoption starts at the practitioner.** Engineers install the OSS and form opinions first. The exec meeting only ratifies it.

03 **Proof means production.** A slideware demo earns nothing. Buyers expect a POV in a prod-like environment inside days, not a quarter.

04 **Procurement grew an AI checklist.** Model risk, data boundaries, agent governance. Meet it in week 2, or lose to it in week 11.

05 **Categories form in months.** A model release can commoditize your moat mid-quarter. Speed of iteration is the durable advantage.

WHAT IT COSTS YOU

IF YOU ARE NOT IN THE ANSWER

You are not on the shortlist, and no one tells you the deal ever existed.

IF THE PRACTITIONER NEVER TRIED IT

The exec meeting is a courtesy. Someone else already ran the bake-off.

IF THE POV TAKES A QUARTER

You lose to a vendor who proved it in nine days with worse technology.

THE JOB

Get into the model's answer, the practitioner's terminal, and production, in that order. The rest is execution.

The GTM delta, in one table.

Same disciplines. Different mechanics. The right-hand column is where AI-native companies actually operate.

DISCIPLINE	THE OLD PLAYBOOK	THE 2026 PLAYBOOK
Demand capture	SEO, SEM, gated PDFs	AEO: be the answer the model gives. Ungated technical depth
Entry point	Top-down exec outreach	Bottoms-up practitioner adoption; exec alignment engineered later
Proof	Custom demo, 90-day pilot	Self-serve trial + scoped POV with signed success criteria
Technical selling	SE demos over Zoom	FDE builds alongside the customer, in their environment
GTM operations	SDR armies and ops headcount	Headless agent workflows; humans on judgment and calls
Content	Quarterly gated whitepaper	Continuous technical content practitioners actually share

THE TELL

The left column still works in legacy categories. In AI, it reads as a vendor who has not used their own product.

What 26 years says still wins.

Six disciplines that predate the model era and survive it intact.

01
MEDDPICC
Metrics, economic buyer, decision and paper process, pain, champion, competition. Deals get inspected, not narrated.

02
COMMAND OF THE MESSAGE
Value framing the economic buyer signs, and discovery that earns the right to deliver it. AI buyers still buy outcomes.

03
MANUFACTURED PIPELINE
PG quotas next to ARR quotas, activity math, coverage inspected weekly. Hope is not a pipeline source. It never was.

04
POV DISCIPLINE
Pre-agreed success criteria, frozen scope, readout to the economic buyer. The POV got faster, never optional.

05
COACHING ON CADENCE
Call review, scorecards, 60-day ramp gates on every new hire. AI does not fix a rep who cannot run discovery.

06
FORECAST AS ARITHMETIC
Stage exit criteria in CRM. Commit, upside, closest-to-pin. Judgment sits on top of math, never instead of it.

THE POINT

AI changed the how. It did not repeal the math. Every module after this one sits on top of these six.

Two engines, one capacity model.

Enterprise and velocity run on different math. They land in the same forecast.

FIELD ENGINE · ENTERPRISE + GROWTH	
ACV $80-400K	by tier; 90-120 day cycles
MEDDPICC + POV	run through nine stages, slide 7
Anchors	meeting → POV 4-5% · POV → win 75-85%
Staffing	AE + BDR pods, SE at 1:3, FDE on strategic

VELOCITY ENGINE · PLG + COMMUNITY	
$30-60K entry	AI-native and mid-market; 30-45 days
PQL-gated	seats + prod-adjacent usage + findings
Anchors	PQL → win 15-25% · reverse trials, 14 days
Staffing	one velocity pod works the whole flow

RULE

One capacity model, one forecast. Separate staffing, separate math, one number. Slides 21-24 build the model.

The nine stages have not changed.

Exit criteria, not activities. A stage advances when the criterion is met, not when the rep feels good.

#	STAGE	CLOSE %	EXIT CRITERIA	ADVANCE →
0	Qualify	0%	ICP fit, BANT, a real initiative	50-55%
1	Discovery	10%	Pain quantified, champion identified	~30%
2	Scoping	20%	Technical fit, budget confirmed, POV criteria drafted	~30%
3	EB Alignment	35%	Economic buyer met, reverse timeline agreed	~90%
4	POV	50%	Success criteria signed, scope frozen	~80%
5	Win the Decision	75%	POV readout to the EB, verbal to proceed	~90%
6	Negotiate & Close	90%	Proposal, mutual action plan, paper process mapped	~95%
7	Closing	99%	Signatures in motion	100%
8	Closed Won	100%	Order form executed	–

3-4%
MEETING → WIN

Nothing structural changed. The POV now runs in days, not weeks, so the whole table spins faster. Stage 4 is shaded because that is where the deal gets decided.

The FDE motion: selling by building.

The forward-deployed engineer is a sales motion with an engineering cost structure. Scope it like one.

WHAT IT IS, WHEN TO RUN IT

- **Forward-deployed engineers** embed with the customer, build the integration in their environment, and map expansion. The build is the proof, and the customer keeps it.
- **Deploy on** agentic or complex products, design-partner phases, and lighthouse logos where the demo cannot carry the sale. Never on a deal an SE can close alone.
- **Palantir normalized it.** OpenAI and Anthropic industrialized it. AI-native buyers now expect a builder in the room, and they can tell inside an hour if you sent one.

THE RULES THAT KEEP IT A SALES MOTION

- Scope FDE time like a POV: success criteria, exit date, EB readout
- FDE cost is CAC, not services revenue. Watch the margin line
- 1 FDE per 2-3 strategic deals in flight; classic SE 1:3 for growth
- FDE builds twice, product ships once. Weekly roadmap loop
- Exit criteria: repeatability. Without it you built a consultancy

SCOPE
Like a POV
criteria, exit date, EB readout

→

COST
CAC, not services
watch the margin line

→

EXIT
Repeatability
without it you built a consultancy

Bottoms-up and top-down, same account.

The practitioner installs it. The economic buyer contracts it. Four steps, in order.

01 **Practitioner installs.**
OSS or free tier, no gate, generous limits. Treat the giveaway as marketing spend, not lost revenue.

02 **Champion forms.**
Usage becomes evidence: findings, saved hours, incidents avoided. The AI stack spots the signal.

03 **Platform team scopes a POV.**
Success criteria signed in writing, not implied on a call. Economic buyer aligned on a reverse timeline from go-live.

04 **The EB contracts the platform.**
Land $50-120K. Security review done once at the platform, then inherited by every team after.

EXPAND VECTORS

Workloads → environments → modules → teams.
Expansion runs on triggers, not on an account manager's mood. NRR 105-110% landing year, 115-120% at scale.

PRICING ARCHITECTURE

- Value metric: per governed workload, not per seat
- Floors by tier; exceptions are CEO-only
- Annual prepay default; multi-year needs cash up front
- Free tier generous; enterprise buys governance

THE POINT

Free adoption is the demand gen line item. The contract is signed one floor up, by someone who never touched the product.

Qualify the company, then the committee.

Five company prerequisites. Three buying rings. Nobody prospects outside them.

ICP PREREQUISITES. ALL FIVE, OR IT IS NOT ICP.

01 AI agents or LLM apps in production, or within two quarters

02 A named owner of the AI platform on the org chart

03 Security budget with an AI line item, this fiscal year

04 A compliance driver: EU AI Act, SOC 2, model risk management

05 Meaningful model and inference spend, growing

Scored weekly by the AI stack. Nobody prospects off-list.

ECONOMIC

CISO, with the CFO co-signing

Buys risk reduction and audit readiness

1-2 PEOPLE

PRODUCT DECISION

VP AI Platform · Dir AI Infra

Owns the platform your product lands on

3-6 PEOPLE

EVALUATORS

AI + security engineers · MLOps

Daily users. Opinions form in the free tier.

8-15 PEOPLE

PRIMARY AXIS

VP AI Platform → Dir AI Infra → AI Engineers. Sell where the pain lives. Contract where the budget lives. The POV bridges the two.

The stack I pay for but never log into.

Every tool below is run by Claude, through official MCPs or Claude in Chrome. I prompt; the stack executes.

KNOW · DATA, RESEARCH, RECORD

Attio
CRM + call recording. System of record

Clay + Apollo
Enrichment, ICP + persona sourcing

Parallel
Account research + web scraping

AthenaHQ
AEO: visibility in AI answers

ONE INTERFACE: A PROMPT

CLAUDE
The control plane

Official MCPs · Claude in Chrome
Cowork agents · skills · projects

Slack + CRM
Where insights and coaching land

ACT · EXECUTION CHANNELS

Origami + Instantly
Outbound email sequencing

HeyReach
LinkedIn outbound campaigns

Nooks
AI autodialer + call sessions

AgentMail
Agent-owned inbox: inbound SDR

THE POINT

Ten paid tools. Zero logins. A human approves every external send. Everything else runs headless.

A prompt is the new login.

Four prompts I actually run. Each one ends in a system I never opened.

“Move Acme to Stage 3 in Attio, log the EB meeting, set next steps.”

→ CRM updated, tasks created. Attio never opened.

“Find 50 Series B AI companies hiring security. Enrich in Clay, draft N=1 copy.”

→ List, enrichment and copy live in Origami within the hour.

“Launch the HeyReach campaign for the conference follow-up list.”

→ LinkedIn touches running by lunch, replies triaged back to me.

“Parse today’s call recordings. Coach the team in #sales, log notes to CRM.”

→ Every recording lands in a Claude project; insights hit Slack.

ALWAYS ON, NO PROMPT NEEDED

- **Meeting intelligence.** Every call parsed. Coaching to Slack, notes and next steps to CRM.
- **Inbound agent.** Claude plus AgentMail: replies in minutes, qualifies, books the meeting.
- **Copy engine.** Cadence copy drafted by Claude skills, tuned on what converts.

THE OPERATING MODEL

Claude is the interface. The stack is the execution layer. One operator runs what used to take a pod.

THE REQUIREMENT

An official MCP, or Claude in Chrome. That covers every tool above, today. No custom integration work, no engineering ticket.

AEO is the new demand capture.

Buyers ask an assistant before they run a search. Five moves, four measures.

01 **Entity-clean docs and llms.txt.** Make the product legible to models: crisp category language, structured docs, consistent naming.

02 **Comparison and alternatives pages.** Answer the questions buyers actually ask assistants: X vs Y, best tools for Z.

03 **Community surface.** Reddit, HN, Stack Overflow, GitHub. Models source opinions where practitioners argue.

04 **Review corpora.** G2 and peer reviews are retrieval fodder. A 30-review gap is a demand-capture gap.

05 **Ungated technical depth.** Models cite what they can read. Every gated PDF is invisible to the answer box.

MEASURE IT

- Monthly share-of-answer audits across ChatGPT, Claude and Perplexity
- AI-referral sessions tracked as their own channel
- Branded-query lift measured after content ships
- Percent of inbound that says an assistant recommended you

THE SHIFT

Rank #1 on Google and still lose the deal. The model handed your competitor a buyer who never searched.

The headless workflow catalog.

Eight jobs that used to be headcount. Every external send still passes a human.

WORKFLOW	STACK	CADENCE	HUMAN GATE
Signal sensing: funding, hiring, tech adoption	Claude Cowork	Daily	None: read-only
Enrichment waterfall + account scoring	Clay + Apollo	On signal	None
Outbound cadences, drafted N=1	Claude → Instantly	Daily	Every send approved
LinkedIn campaigns + reply triage	Claude → HeyReach	Daily	Every send approved
CRM hygiene + stage enforcement	Attio / Salesforce	Nightly	Exceptions flagged
AEO, SEM and content pipeline	Claude Cowork	Weekly	Publish approved
Event ops: lists, pre-books, follow-ups	Claude + CRM	Per event	Sends approved
Win/loss tags + forecast prep	Claude + CRM	Weekly	CRO judgment

THE HIRE

One GTM engineer runs this entire table. That role replaces the ops admin as the first RevOps seat, by roughly $3M ARR.

What the engine buys you.

Not fewer people. The same people, spending their hours where judgment actually pays.

100%

ACCOUNT COVERAGE

Every tiered account researched continuously, not the top 20%.

Under 5 min

SPEED TO LEAD, 24/7

PQLs and inbound routed, researched and answered around the clock.

~3 FTE

SDR CAPACITY ADDED

Without the headcount. Pods spend their hours in live conversations.

15 → 2 min

RESEARCH TIME

Brief, angle and opening line waiting in the CRM before the rep dials.

N=1

PERSONALIZATION

Copy drafted per account from live signals, never from merge tokens.

Nightly

CRM HYGIENE RUNS

Stages, next steps and close dates enforced before the forecast call.

NOT A CONCEPT

My daily driver.

This is the production stack I run across portfolio companies today. Claude Cowork orchestrates Attio, Clay, Apollo, HeyReach and Instantly on headless daily jobs, with human-approved sends.

This playbook was researched, written and designed by that same stack.

ATTIO · CLAY · APOLLO · HEYREACH · INSTANTLY · NOOKS · ATHENAHQ

THE TRADE

Headcount stays. The hours move to judgment, relationships and closing.

Four stages, one question each.

You graduate on evidence, or you do not get to spend like the next stage.

$0-1M

PROVE IT

EVIDENCE, NOT BOOKINGS

Founder-led selling, an FDE in the account, design partners signed at real prices.

THE QUESTION

Do buyers with no tie to you buy, deploy and renew?

$1-5M

REPEAT IT

PROCESS IS THE PRODUCT

First AEs, process of record formalized, PG quotas set, community funnel running.

THE QUESTION

Can two non-founder sellers hit 70%+ of quota?

$5-10M

ENGINE IT

THE DECIDING STAGE

Dual engine live, monetization fixed, capacity modeled with a haircut on ramping reps.

THE QUESTION

Does the math clear field economics at your real ACV?

$10-20M

COMPOUND IT

GROWTH GETS CHEAPER

Channel partners, first-line managers, an international expansion machine.

THE QUESTION

Does growth get cheaper as the number gets bigger?

HOW TO READ IT

Stages are gates, not vibes. The next four slides are the playbook for each one.

Stage 1 · $0-1M: sell truth, not scale.

Nothing here scales. That is the point. You are buying evidence, not bookings.

01 **Founder sells. An FDE builds.**
Nobody hires a VP Sales yet. The founder cannot outsource learning what makes buyers move.

02 **10-20 design partners at real prices.**
Discounts trade for references and roadmap input, never for silence. Free pilots teach nothing about willingness to pay.

03 **Run pricing experiments on purpose.**
Three packaging tests before you claim repeatability. The metric that scales is tied to value delivered.

04 **Instrument everything from deal one.**
Win/loss on every deal, CRM discipline from the first opportunity, signal capture running even now.

THE STAGE IN ONE NUMBER

10-20

DESIGN PARTNERS, REAL PRICES

GRADUATE WHEN

Two net-new logos, sourced outside the network, live through a full renewal cycle.

FAILURE MODE

Renting revenue from the founder's network and calling it repeatability.

THE ONLY PROOF

Two logos with no tie to the founder. Bought, deployed, renewed. Everything before that is anecdote.

Stage 2 · $1-5M: the first repeatable pod.

Two sellers who are not the founder have to hit a number. Everything here serves that test.

01 **Hire 2-3 AE athletes, not a big-company VP.** Sellers who prospect, run their own POVs and thrive in ambiguity. The professionalize-it VP comes later.

02 **Process of record in CRM, week one.** Stages, exit criteria, MEDDPICC fields. Forecast becomes arithmetic the day the fields exist.

03 **Pipeline-generation quotas from day one.** AE sources one net-new ICP meeting a week, BDR two. Coverage inspected Fridays. Hope is not a source.

04 **Formalize the community funnel.** PQL gates, reverse trials, DevRel that ships technical content. Bottoms-up stops being luck.

05 **Hire the GTM engineer by ~$3M.** The AI stack scales before headcount does. One person runs the entire workflow catalog.

THE STAGE IN ONE NUMBER

~$3M

WHEN THE GTM ENGINEER LANDS

GRADUATE WHEN

Two non-founder reps at 70%+ attainment, two consecutive quarters.

FAILURE MODE

Scaling outbound spend before message-market fit is proven outside the community.

THE TEST

Two sellers, 70% of quota, twice running. Anything less and you are still at Stage 1 with a bigger payroll.

Stage 3 · $5-10M: the engine, math shown.

Illustrative: $4.2M today, $10M planned. New logo carries $5.2M of that gap, net expansion carries the rest.

FIELD ENGINE · $3.6M NEW LOGO

18 / week
AE 30% · BDR 35% · marketing 25% · partners 10%

864 meetings
ICP only, held across the year

× 48 WEEKS

37 POVs
scoped, success criteria signed

× 4.3% TO POV

VELOCITY ENGINE · $1.6M NEW LOGO

420 free-tier orgs
fed by the OSS and community funnel

160 PQLs
3+ seats · prod-adjacent · findings in 14 days

× 38% TO PQL

64 velocity opps
worked by the velocity pod, not the field

× 40% TO OPP

62 wins · $5.2M new logo
Field: 37 POVs × 81% win = 30 at $120K blended ACV. Velocity: 64 opps × 50% win = 32 at $50K ACV.

THE HAIRCUT

Staff the seat grid against a 30% team failure rate. If capacity net of the haircut does not clear the number, change the plan, not the story.

Stage 4 · $10-20M: make growth compound.

The question stops being can you sell it and becomes does each new dollar cost less than the last.

15% → 35%

CHANNEL SHARE OF NEW

Marketplaces first, since EDP and MACC budgets unlock stuck procurement. Then two security VARs.

6-8

REPS PER FIRST-LINE MANAGER

Players stop coaching at this span. Managers get hired on coaching evidence, not on their old logo slide.

115%+

NRR, FROM PLAYBOOKS

Whitespace maps, trigger-based plays, QBRs run with customers, not at them. Expansion stops being a surprise.

INTERNATIONAL

Rides evidence, never ahead of it. An EMEA pod once three or more organic logos already exist in region. Hiring into a region with no pull is the most expensive way to learn it has no pull. Open with one AE and one SE on local paper, local support hours and a reseller already selling into those same accounts.

THE BOARD FORECAST

Built bottom-up, seat by seat. Capacity per seat, coverage entering the quarter, a tripwire attached to every assumption. Ramp curves, start dates and the 30% failure haircut all sit in one model, so moving a hire date reprices the year on the spot. The Series B gets raised on evidence, not narrative.

GRADUATE WHEN
Channel clears 25% of new business, two quarters running.

FAILURE MODE
Scaling spend on an engine whose Stage 3 unit math never closed.

Sales capacity planning, defined.

The model that turns a growth target into a hiring calendar, and a hiring calendar into a number you can commit.

A capacity model forecasts ARR from rep production, seat by seat. Ambition is not one of the inputs.

Each seat carries a ramp schedule and an expected first-year yield. The sum of scheduled seats, net of a failure haircut, is the number you can commit.

The fundamental advantage: the revenue plan and the headcount plan become the same document. You cannot change one without repricing the other.

HOW IT WORKS

01 **Define the unit.** One seat's ramp schedule and first-year yield, at the real median deal size.

02 **Schedule the seats.** Hire dates decide when capacity exists. The grid is the calendar.

03 **Haircut the grid.** Some reps will miss. The model absorbs 30% before the board has to.

04 **Commit the output.** What survives the haircut is the number, and it is already a hiring plan.

WHY IT PREDICTS

Misses become seat-level variances you diagnose on Friday. Not annual surprises you explain in Q4. Executed well, the growth number is an arithmetic consequence of hiring dates.

A rep is a $600K asset in year one.

Example scenario: a company exits this year at $1.2M ARR, targeting $7.0M in 12 months. Start with one seat, one ramp.

BOOKINGS RAMP, ONE NEW SEAT

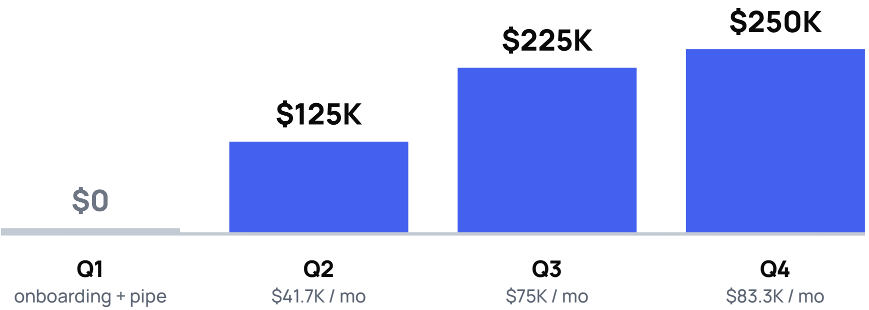

THE UNIT, SUMMARIZED

$600K	**$1.0M/yr**	**$120K**	**$300-330K**
YEAR-1 YIELD	STEADY, M10+	MEDIAN DEAL	OTE, 50/50

THE REALITY CHECK · 2026 MARKET DATA

Industry data: 252 cloud sales orgs, ~58,500 quota-carrying sellers.

43.6% — MEDIAN QUOTA ATTAINMENT
Cybersecurity runs 40.3%. Plan at production, then haircut the team 30%.

178 days — MEDIAN CYCLE, $100-200K DEALS
Add a 90-day ramp: a month-7 hire delivers a stub, not a year.

2.65 — LEADFLOW SENTIMENT, RECORD LOW
Capacity without pipeline is idle payroll. Every seat carries a PG quota.

Boards that plan at 100% attainment are planning fiction.
2026 data says 44%.

THE RULE

Plan capacity at production, not quota. Haircut the team 30%. Front-load hiring: late seats pay next year.

The activity math behind the $600K.

One ramping seat, first 12 months, locked at the $120K median deal.

INPUTS · LOCKED

$120K
MEDIAN NEW-LOGO DEAL

$1.0M/yr
STEADY-STATE CAPACITY

0-50-90-100%
QUARTERLY RAMP

$600K
EXPECTED YEAR-1 YIELD

= 5 new logos required in year one

THE FUNNEL, COMPOUNDED FROM THE NINE-STAGE SYSTEM

147 MEETINGS HELD → 16% → **23** QUALIFIED OPPS → 27% → **6** POVS STARTED → 81% → **5 WINS** = $600K BOOKED

ACTIVITY, PER RAMPED SEAT	ANNUAL	QTR	MO	WK
New discovery meetings (S0)	147	37	12	3
Qualified opportunities (S2)	23	6	2	0.5
POVs started	6	1.5	0.5	0.1
Closed-won new logos	5	1.3	0.4	0.1

Compounded from the nine-stage table: meeting → qualified 16% · qualified → POV 27% · POV → win 81%.

THE LINK

Three net-new meetings a week per seat: exactly the pod's dual PG quota (AE 1 + BDR 2). The capacity model and the activity system are the same machine.

From $1.2M to $7.0M: a hiring schedule.

Example scenario: a company exits this year at $1.2M ARR and wants to reach $7.0M total ARR in the next 12 months.

SEAT	START	Q1	Q2	Q3	Q4	YEAR 1
Rep A	carried in, ramped	$250K	$250K	$250K	$250K	$1.00M
Rep B	carried in, ramped	$250K	$250K	$250K	$250K	$1.00M
Rep C	carried, prior Oct	$125K	$225K	$250K	$250K	$850K
Rep D	carried, prior Nov	$83K	$192K	$242K	$250K	$767K
Reps E-F	2 hires, month 1	–	$250K	$450K	$500K	$1.20M
Reps G-H	2 hires, month 2	–	$167K	$383K	$483K	$1.03M
Rep I	hire, month 3	–	$42K	$158K	$233K	$433K
Rep J	hire, month 4	–	–	$125K	$225K	$350K
Rep K	hire, month 6	–	–	$42K	$158K	$200K
Rep L	hire, month 7	–	–	–	$125K	$125K

Cell shading steps with the dollar figure: the capacity staircase, made visible. 2026 cloud sales benchmarks.

THE HIRING RULES

- Front-load: all 8 hires land by July. Later seats pay next year.
- Pair the first four hires: shared onboarding, an honest A/B.
- The number costs money first: about $3.3M of in-year AE comp.
- Q4 annualizes to $10.9M gross: the schedule pre-builds next year.

Working model ships as a spreadsheet.

GROSS CAPACITY		FAILURE HAIRCUT		NET NEW LOGO		STARTING ARR		EXPANSION, NET		EXIT ARR
$6.96M	–	**$2.09M**	=	**$4.87M**	+	**$1.20M**	+	**$0.93M**	=	**$7.00M**
12 seats		30% of the grid		~41 at $120K		the ARR carried in		base, less churn		total company ARR

The functions you are actually hiring.

Sixteen functions. Which exist by $5M, $10M and $20M is the hiring argument.

	SALES FUNCTION				MARKETING FUNCTION			
NEW LOGOS & EXPANSION	RSMs 5M **Enterprise Sales**	AEs 5M **Commercial Sales**	BDRs 5M **Business Development**	SEs 5M **Sales Engineering**	PIPELINE 5M **Demand Generation**	DEVREL 5M **Dev Community**	PMM 10M **Product Marketing**	CONTENT 10M **Content Marketing**
CHANNEL & ALLIANCES	CAMs 20M **Channel Sales**		PARTNERSHIPS 20M **GTM Alliances**		FIELD 20M **Field & Channel Marketing**		CO-MARKETING 20M **Partner Co-Marketing**	
RENEWALS & SUCCESS	AMs 10M **Account Managers**		CSEs 10M **Customer Success Engineers**		ADVOCACY 20M **Customer Marketing**		BRAND 20M **Brand & Corporate**	

THE 17TH FUNCTION

It spans every function above. Systems, forecast, comp, territory, data. One person: the GTM engineer.

Who to hire, and when.

Four stages. The seat count is the plan; the role each stage unlocks is the argument.

STAGE	SELLERS	SUPPORT	ROLE THE STAGE UNLOCKS
$0-1M	Founder + 1 FDE	Contract SE help	FDE: the builder in the room
$1-5M	2-3 AEs + 1-2 BDRs	1 SE · DevRel	GTM engineer by ~$3M
$5-10M	6-8 AEs, dual-track + velocity pod	2-3 SEs · 2 BDRs	CRO or VP Sales, player-coach
$10-20M	10-14 AEs · EMEA pod	SE team · channel lead	First-line managers · VP Marketing

PAIR EVERY SELLER HIRE

Two at a time, shared onboarding,
and an honest A/B on whether the bar held.

HIRE AHEAD OF THE RAMP

A month-7 hire contributes a stub this year.
The ramp is the constraint.

BACKFILL ON DAY 60

Two misses at the 60-day activity gate,
open the req that same day.

THE RULE

Every seat on this table is a line in the capacity model. Change a hire date and the ARR number reprices itself.

What the team costs.

2026 ranges for AI and AI-security scale-ups in US metros. Quota-to-OTE is the number that decides the model.

ROLE	OTE RANGE	SPLIT	MECHANICS
Enterprise AE	$300-330K	50 / 50	Quota-to-OTE 3-3.5x. Accelerators 1.25x above plan, no cliffs.
FDE	$220-280K	80 / 20	Bonus on POV wins and on repeats that get productized.
Sales Engineer	$230-260K	70 / 30	Variable on POV win rate and POV cycle time.
MM / Velocity AE	$200-250K	50 / 50 to 60 / 40	Paid on tier wins and on conversion SLAs, not on discovery volume.
GTM Engineer	$160-200K	Base + bonus	Owns the AI stack. Bonus on pipeline-per-pod efficiency.
BDR	$90-110K	75 / 25	Paid on held meetings that convert to Stage 2, never on dials.

55-60%
YEAR-1 YIELD VS QUOTA

Nobody carries a full number before month seven. Plan year-one production against this yield, and gate 10% of variable comp on pipeline generation.

The operating rhythm.

Leading indicators get inspected. Lagging indicators get judged. The calendar enforces both.

LEADING · INSPECTED WEEKLY

- Stage-0 meetings by pod and source, vs the weekly bar
- PQL flow vs gate, velocity pod SLA compliance
- Coverage entering quarter: **3.5x new · 1.2x expansion**
- Stale pipeline **under 15%** older than two quarters
- MEDDPICC score on every Stage-2+ opportunity

LAGGING · JUDGED MONTHLY

- Bookings vs plan, by segment and by source
- Blended new-logo ACV vs the field-economics floor
- Stage-2 to win **20%+** · POV-to-win **75-85%**
- NRR **115%+** · GRR **90%+**
- CAC payback **under 18 months** · burn multiple

THE CADENCE

MON	WED	FRI	MONTHLY	QUARTERLY
Commit and upside call	Pipeline scrub by stage	PG scorecard by pod	Win/loss and pricing	QBRs and the board pack

THE RHYTHM

Nothing on this page is a status meeting. Every session ends in a decision, an owner and a date.

Five ways AI scale-ups die.

Every risk gets a mitigation and a numeric tripwire. When the tripwire fires, the decision is already made.

	RISK	MITIGATION	TRIPWIRE → ACTION
01	Adoption runs slower than the hype	Anchor on what is in production today. PLG keeps CAC low while the category matures.	Under 40 PQLs in a quarter → shift the mix to field, revisit the segment bets
02	Priced as a tool, not a platform	Pricing floors, platform packaging, POV outcomes the economic buyer signed.	Field ACV under $120K for two quarters → repackage and requalify up-market
03	Rep failure runs past the haircut	Hold the hiring bar. Scorecards from day one, activity gates at day 60.	Two reps below gates at day 60 → reallocate pipeline, open the backfill req that day
04	Paid tiers burn community goodwill	Clean OSS and paid boundary: the community keeps the tool, enterprise buys governance.	Downloads down 20% for two months → packaging review with founders
05	Enterprise deals go single-threaded	MEDDPICC enforced: no Stage 3 without an economic buyer meeting on the calendar.	Over 30% of late-stage pipeline single-threaded → daily deal reviews on those accounts

WHY THIS EXISTS

A plan without failure conditions is a pitch. Wire each tripwire to a decision and review all five at every QBR.

The operator behind the playbook.

Enterprise revenue, rebuilt around an AI-native stack.

ANDREW WESBECHER

26 years building enterprise security and infrastructure revenue.

VP Sales and CRO roles scaling early-stage SaaS.

Today, fractional CRO across AI security: AI-SPM, identity, agent security, detection and response.

This playbook was researched, written and designed with the AI stack it describes.

WHERE THE WORK HAPPENED

TIBCO · MERAKI · THOUSANDEYES
LACEWORK · CONTRAST · TRACEABLE

$11B+

COMBINED EXIT & PEAK VALUE

Lacework $8.3B peak. Meraki $1.2B and ThousandEyes $925M to Cisco. Traceable $900M to Harness.

$1M → $25M

EARLY-STAGE ARR, SCALED

The path this playbook maps, run more than once as first or second sales leader.

4

AI SECURITY CLIENTS TODAY

Fractional CRO and advisory work. One client is an RSAC 2026 Innovation Sandbox finalist.

10

GTM PLANS FOR AI COMPANIES

Every framework in this playbook shipped in a real plan first, then survived a real pipeline.

NEXT STEP

If this maps to your next twelve months, I would rather talk than pitch. awesbecher@gmail.com

The foundations are not nostalgia.
The AI layer is not a gimmick.

Operators who hold both will define this decade of enterprise software.

If you are scaling an AI company through $5M, $10M or $20M and want to compare notes, my inbox is open.

Andrew Wesbecher
linkedin.com/in/wesbecher · awesbecher@gmail.com

WHAT TO TAKE AWAY

01 **The foundations still price the deal.**

02 **The AI layer changes who does the work.**

03 **The plan is a hiring calendar.**

 ·