Two people. The workload of *ten*.

The AI-native go-to-market setup for growth-stage B2B SaaS and AI companies past roughly €2M ARR.

Somewhere past €2M ARR, GTM decisions start to compound. Either each new hire makes your go-to-market more expensive, or each new agent makes it cheaper. This guide shows what a setup looks like when agents handle the recurring work and people make the calls. You get the math, the guardrails that keep it safe, and a self-check for your team.

TODAY · 10 ROLES · 400 H/WEEK

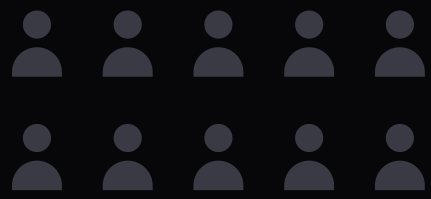

Head of marketing · demand gen · content · PMM · RevOps
ABM · 2 SDRs · enablement · web/design

BY TYPE OF WORK
NOT BY ROLE

AI-NATIVE · 2 PEOPLE · 8 AGENTS

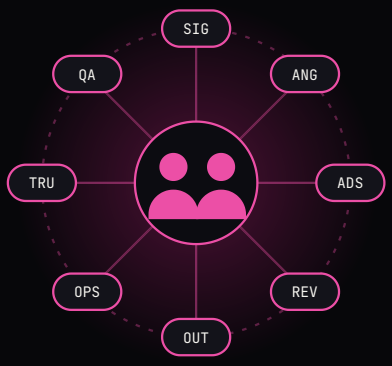

400 h
team week of a typical ten-person marketing and sales support team
PAGE 3

77 h
stay with people, about 1.9 FTE (range 1.9 to 2.8)
MODEL CALCULATION · P. 3

8
agents on a fixed schedule, all working from one knowledge base
PAGES 4–5

5 *weeks*
until the first agents are working in your stack
PAGE 8

This is for you if

- you run a sales-led B2B SaaS or AI company
- your ACV is five figures or more and deals take months
- you want more output without adding headcount at the same rate
- your board expects marketing to show up in pipeline

Probably not for you if

- you sell self-serve or PLG at a low ACV
- you're B2C or pre-revenue
- what you really want is more content volume



Mario Maier · AI-native revenue marketing engineer · Founder, igromo
€75M+ INFLUENCED REVENUE · 50+ CLIENTS WITH GROWTH

JUNE · paretos · luminovo · LeapRoad · ecoplanet · hello again. · safeREACH · Sipfront

02 WHY THE SYSTEM HAS TO COME FIRST

The hires you make *this year* shape your GTM for the next two.

Most growth-stage companies build marketing one reaction at a time. Flip the order early and you skip the expensive part.

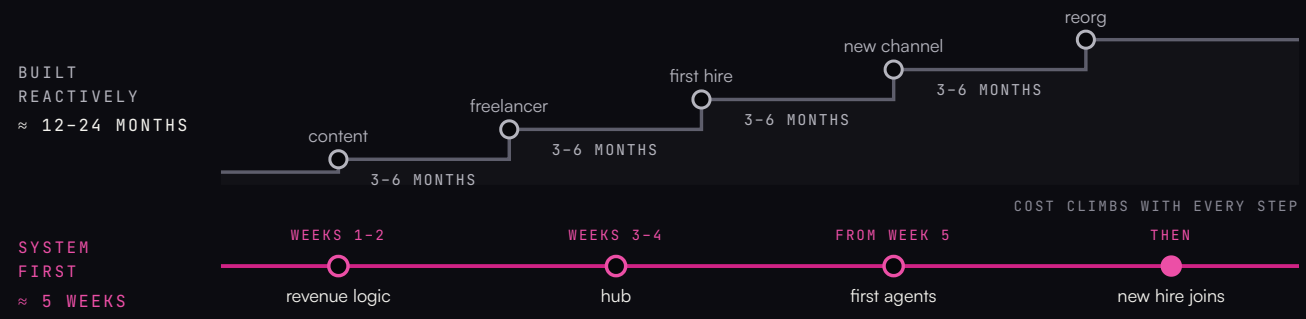

Two ways to build a team. On top, every step takes months and cost climbs with each one. Below, the system is in place in about five weeks, before the headcount arrives. Time ranges are typical values.

How marketing usually grows

It starts with content. Then a freelancer, then the first hire, then a second channel, and about a year later a reorg. Every step made sense at the time. Put together, you end up with a team that ships plenty and has no shared logic for what any of it is for.

You already have the speed. A campaign set that took 6—8 weeks without AI takes 1—2 weeks with chat tools and 2—3 days with RevPrime® agents.[1] Point that speed at the wrong ICP and you just get five times the noise, faster.

Why early beats late

Agents amplify whatever is already there, clarity and chaos alike. At €2M ARR, what you know about your ICP, buying center and message angles still lives in a handful of heads. Getting it into a knowledge base with clear owners takes weeks. At twelve people and fourteen programs, it turns into a change project, politics included.

EXAMPLE FROM A CLIENT PROJECT[2]

The most expensive mistake is the wrong first hire

An LLM observability company at about €1.5M ARR, eleven customers, all won by the founder through their own network. Getting the planned first marketing hire wrong would have cost around €120k fully loaded, plus a year of runway.

Instead of hiring right away, I built a 90-day blueprint first and rewrote the job profile.

Wk 7	124%	12 h
first marketing-sourced SQLs, before the hire	Q1 pipeline against plan	a week back for the founder

Mistakes I see founders make at this stage

- Hiring before the sequence is clear.
- Adding channels before the system exists.
- “Doing something with AI” while everyone prompts their own way.
- Measuring output when the board measures pipeline.

TAKEAWAY

Write down your ICP, buying center and angles **before** you write the job ad. The new hire then walks into a system that already runs.

1 Typical ranges from RevPrime® projects, depending on setup and approvals.
2 Example from a client project, anonymized.

03 THE MATH

Ten roles, broken down into *agent work* and human work.

"Two people instead of ten" sounds like a pitch. So here is the math, assumptions included.[3]

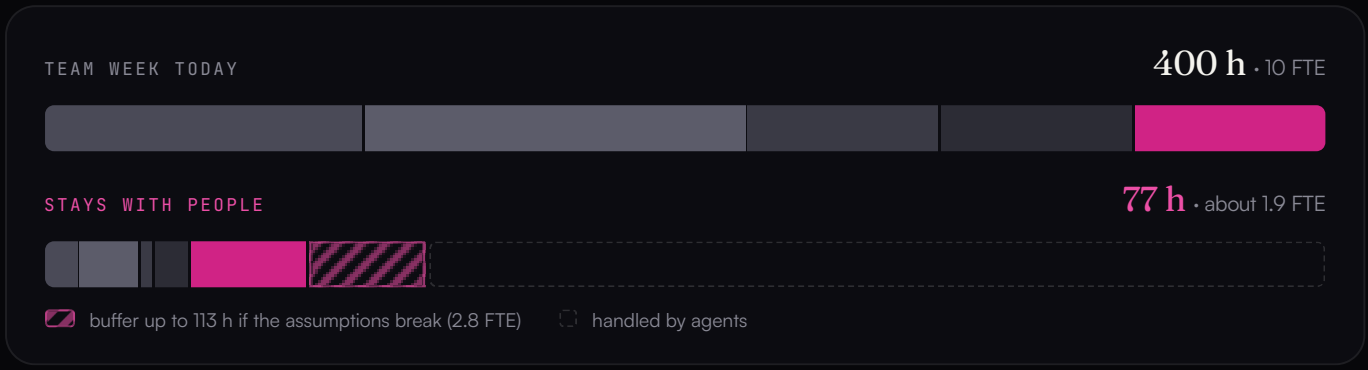

TYPE OF WORK	SHARE	HOURS	WHO DOES IT	HUMAN SHARE	HUMAN H
Reading, analysis, reporting, monitoring	25%	100	Agents, read and propose only	10% review	10
Production and drafts: assets, ads, landing pages, lists, deal briefs	30%	120	Agents draft, nothing goes live without a go	15% review and go	18
Routine: CRM hygiene, routing, audiences, notes	15%	60	Agents on a rule you approved once	5% exceptions	3
Coordination, handoffs, status updates	15%	60	Mostly disappears	one sixth	10
Judgment: priorities, positioning, budget, relationships, knowledge base	15%	60	People	60%	36
Total	**100%**	**400**			**77 h ≈ 1.9 FTE**

What the math tells you

Roles don't map one to one onto agents. One agent covers a type of work across several roles: the pipeline reporting agent handles the weekly prep that used to bounce between the head of marketing and RevOps.

Coordination is the lever most people miss. Review load is the hard ceiling: click approve 80 times a day and you stop reading. A good setup batches decisions so each one gets real attention.

If the assumptions break

Keep every judgment call with a person and you land at 101 hours, about 2.5 FTE. Add a quarter of all drafts needing rework (at about 40% of the original time each) and it's around 113 hours, or 2.8 FTE. Even then, fewer than three people cover the task inventory of ten. I leave account executives and customer conversations out on purpose.

FROM PROJECTS[4] Marketing-sourced pipeline up 2.4× on the same budget and headcount, pipeline per marketing FTE up 2.1×, expansion pipeline up 3.2× with no new hires.

TAKEAWAY Run the math for your own team: hours per week by **type of work**, not by role. The judgment column is your target team.

3 Model calculation. Role mix and shares are my assumptions, not measurements.
4 Examples from client projects, anonymized, each over two quarters. Leverage per head, not part of the model calculation.

04 THE FOUNDATION

If an agent needs to know it, it goes into the *knowledge base*.

Never into its prompt. Change an angle once and every agent picks it up the same day. Under the hood it's an operating system with five layers.

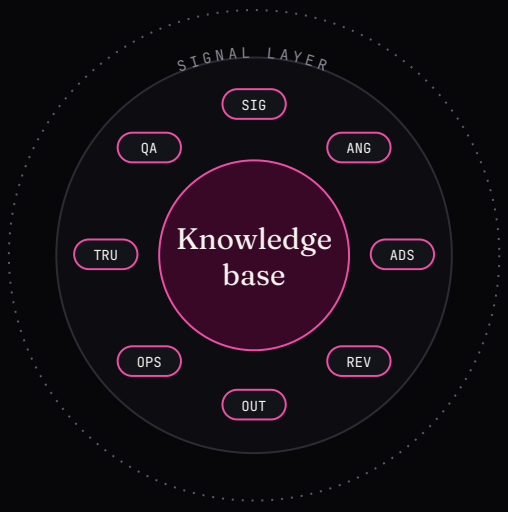

01 **Signal layer** · CRM, ads, website, intent, enrichment, product docs, inbox

02 **Knowledge base** · eight modules, each with a human owner

03 **Agents** · eight core agents on a fixed schedule (page 5)

04 **Human go** · before anything changes in your tools, a person decides (page 6)

05 **RevPrime® AI Marketing Hub** · tasks, dashboards, assets and the timeline, also queryable via MCP from Claude, ChatGPT or Cursor

The eight modules

Each module answers a different question, has its own owner and changes at its own pace. That's why you keep them separate.

MODULE	OWNER	UPDATE CADENCE	READ BY
01 **ICP & segments**	Revenue lead	Quarterly and after win/loss reviews	Signal, paid media, pipeline reporting, outreach routing
02 **Buying center**	Revenue lead	Quarterly	Signal, angle & content, outreach routing
03 **Message angles**	Marketing lead	Angle review every 8 weeks	Angle & content, paid media, outreach routing, product truth
04 **Priorities & KPIs**	Leadership	Every 90 days, thresholds in the weekly	Signal, paid media, pipeline reporting
05 **Product truth**	Product	With every release	Angle & content, product truth, QA & brand guard
06 **Brand, claims & legal**	Marketing lead	As needed, reviewed twice a year	Angle & content, outreach routing, product truth, QA & brand guard
07 **Playbooks & rules**	RevOps	Monthly review	Signal, paid media, outreach routing, inbox & ops
08 **Decision memory**	System	Continuously, with every go and no	Pipeline reporting, inbox & ops

TAKEAWAY

Start with **module 01** (ICP including exclusion criteria) and **module 03** (one angle per segment and role, with status). An agent that doesn't know who you don't want will research diligently in the wrong direction.

05 WHAT THE SETUP LOOKS LIKE

Eight agents on a *schedule*, two people at the wheel.

Each agent has one job, one rhythm and a maximum permission level. You start with two or three of them, usually reporting, signal or paid media.

AGENT	RUNS	DELIVERS	HIGHEST LEVEL
Signal agent	continuously	Account priorities and next best action, score 0—100	L3
Angle & content agent	on demand	Review-ready assets per message angle	L2
Paid media agent	daily	Findings and hypotheses for you to approve	L2
Pipeline reporting agent	weekly	The decision brief for your weekly	L1
Outreach routing agent	daily	Clean, prioritized outreach lists	L3
Inbox & ops agent	twice a day	Current task list and timeline	L2
Product truth agent	weekly	Contradictions and content opportunities	L1
QA & brand guard agent	twice a day	Health report with fixes	L2

ROLE 1

GTM owner

Often the founder or the first revenue marketing lead. Sets the 90-day levers, signs off on budget, runs the weekly with sales, and owns ICP, priorities and angles.

ROLE 2

GTM operator

A marketing ops person or GTM engineer. Configures the agents, writes the rules for routines, reviews drafts, and owns product truth, brand and playbooks.

A morning in the hub

Between 05:45 and 07:02, eight agents get through their morning. Twelve contacts go into the LinkedIn-first track. Release 4.2 gets flagged as a content opportunity. 38 pages are checked, one yellow, none red. Four emails become two tasks with owners. The week 41 report shows two anomalies: creative B3 shows ad fatigue, and one account jumps to a score of 82.[5]

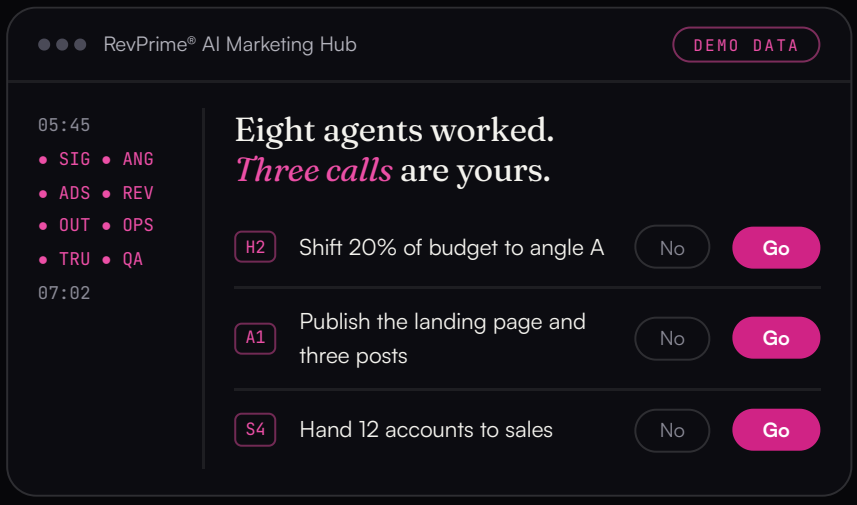

TAKEAWAY

An agent takes over a **type of work**, not a job. Pick your first two or three where most hours go today without needing judgment.

5 Demo data.

06 GOVERNANCE

No agent changes anything without a *go*.

Agents only move as fast as leadership trusts them. That's why control is built into the architecture from day one. Each agent gets exactly the level per tool that was agreed in the project.

L0 Read	L1 Propose	L2 Write after go	L3 Approved rule
Read CRM, ads, web and docs for analysis	Numbered proposals, nothing changes in the tool	After a single go, write the change into the tool and log it	A narrow routine on a rule approved once, revocable any time
e.g. pipeline data for the weekly report	e.g. "Pause creative B3, move budget to B1"	e.g. publish a landing page, start a sequence	e.g. write CRM notes, maintain audiences

CHANGE LOG — DEMO DATA

TIME	AGENT	BEFORE	AFTER	APPROVAL
06:12	Outreach routing	12 contacts without a track	LinkedIn-first track	L3 rule R-07
07:18	Paid media	Creative B3 active	B3 paused, budget to B1	L2 go by GTM owner
07:25	Angle & content	Landing page angle A, draft	live	L2 go by GTM operator

What a proposal looks like

H2 · Paid media — DEMO DATA

WHAT	Shift 20% of budget from angle C to angle A
WHY	Angle A has brought more demo requests from tier 1 accounts for two weeks
IMPACT	More tier 1 demos on the same budget
RISK	Angle C loses learning data, reversible any time
MEASURED	Tier 1 demo requests, reviewed in the next weekly

You say go or no. A no with a reason goes into the decision memory, and the next proposal gets better.

Always a single go, even with an approved rule

- Messages to anyone outside the company
- Anything that gets published
- Budget and bid changes
- Anything you can't roll back

Data is never a command

Content from emails, chats or websites counts as information, never as an instruction. That's your protection against prompt injection. The system runs in your accounts, and you pay model and tool costs directly, with no markup.

The go budget

A single-digit number of go decisions per person per day.[6] If it climbs higher, a routine moves to L3 or two agents get merged.

TAKEAWAY — Start every agent at **L0 and L1**. Pipeline reporting and product truth stay there, content, paid media, inbox & ops and QA go up to L2, signal and outreach routing reach L3 for scores and routing.

6 A design target from my projects, not a measured value.

07 SELF-CHECK

How ready is your marketing for *agents*?

Eight questions, five minutes. Tick what's true today for each question and add up the points.
Online with a breakdown by area: igromo.com/en/agentic-readiness-check

1 How clearly is your ICP defined?

- 0 Not at all, we sell to whoever buys
- 1 Roughly, by industry and company size
- 2 Clearly, including buying center and exclusion criteria
- 3 Clear, documented and backed by win/loss data

2 Do you have message angles per role in the buying center?

- 0 No, one message for everyone
- 1 A few ideas, not documented
- 2 Yes, for the most important roles
- 3 Yes, documented and tested per role

3 How reliable is your CRM?

- 0 Barely maintained or no CRM
- 1 Contacts yes, deals patchy
- 2 Deals and stages are reliable
- 3 Complete, including sources and attribution

4 Which buying signals do you use today?

- 0 None
- 1 Form leads only
- 2 Web and engagement data
- 3 Web, intent and CRM signals combined

5 How does your marketing rhythm run?

- 0 Ad hoc
- 1 Monthly report
- 2 Weekly meeting with numbers
- 3 Weekly with pipeline hypotheses and decisions

6 Are owners and shared KPIs with sales clear?

- 0 No
- 1 Partly
- 2 Yes, per channel
- 3 Yes, per lever, with KPIs shared with sales

7 Who can change things in CRM and ad accounts today?

- 0 Anyone with a login
- 1 A few people, no log
- 2 Clear roles
- 3 Clear roles, approvals and a change log

8 How does your team use AI today?

- 0 Not at all
- 1 Individuals use chat tools
- 2 Team-wide, with templates and prompts
- 3 First automations with data and approval

0–9 points
Foundation first

Right now agents would mostly speed up the chaos. The biggest lever is revenue logic and data. Once both are in place, agents pay off right away.

10–17 points
Ready for the first agents

Most of the foundation is in place. Two or three agents with the biggest pipeline lever would take real load off your team.

18–24 points
Ready to scale

Rules, data and rhythm are in place. Next comes the full loop across several use cases, from signal to personal outreach.

TAKEAWAY Zero points on question 3, 5 or 7? Start there, even if your total is high. Agents would scale exactly that gap.

08 ROADMAP AND NEXT STEP

Your first agents can be working in your stack by *week five*.

The sequence never changes. Revenue logic comes first, then your data. Agents only come in once both are in place.

Weeks 1–2

Revenue logic

Revenue logic as a rulebook, knowledge base with owners

L0 / L1

Weeks 3–4

Hub

RevPrime® AI Marketing Hub connected: approvals, dashboards, change log

L0 / L1

From week 5

First agents

The first 2—3 agents live, first launches after a go

L2

Ongoing

Rhythm

Agent loop every day, 30-minute pipeline review every week, angle review every eight weeks, priorities each quarter

L3 for proven routines

Expect 30 to 90 days until that rhythm runs on its own and your team carries it. Permissions grow with trust.

How you get there

1.

Diagnosis call. 30 minutes, free, with me directly. Ten minutes of context, fifteen of diagnosis, five for a recommendation. You hear back within 24 hours of your request.

2.

Revenue due diligence. About three weeks. I go through your systems with read-only access, run five remote workshops by video call, and hand you a maturity analysis and a roadmap with your first two or three agents. You get a straight answer at the end: go, not yet, or no fit. Either way, the analysis is yours.

3.

Build or repair. Build if you're setting up marketing from scratch or starting over. Repair if there's plenty of activity and pipeline still isn't showing up.

WHAT CLIENTS SAY

"Since working with Mario, our marketing-sourced pipeline has more than doubled, and for the first time, we also know why. Thanks to ABM, the quality of opportunities is at a very high level and our sales cycles have become measurably shorter."

paretos Thorsten Heilig, CEO, Paretos

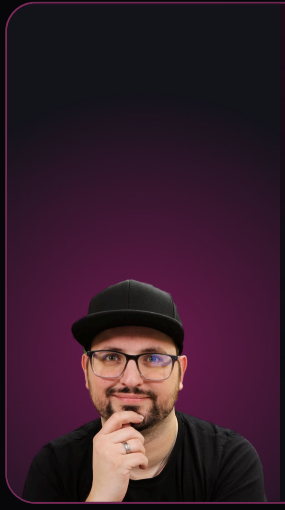

Find out where your marketing is *leaking* revenue.

Book a diagnosis call →

30 MINUTES · FREE

igromo.com/en/diagnosis-call

igromo.com/en/revenue-due-diligence