igromo

Edition 2026

WHITE PAPER · THE AGENTIC MARKETING BLUEPRINT

From prompts to *pipeline.*

How sales-led B2B SaaS & AI companies build marketing that runs itself: a knowledge base at the core, specialized agents doing the work (a core set, extended for every business) and humans making the decisions.

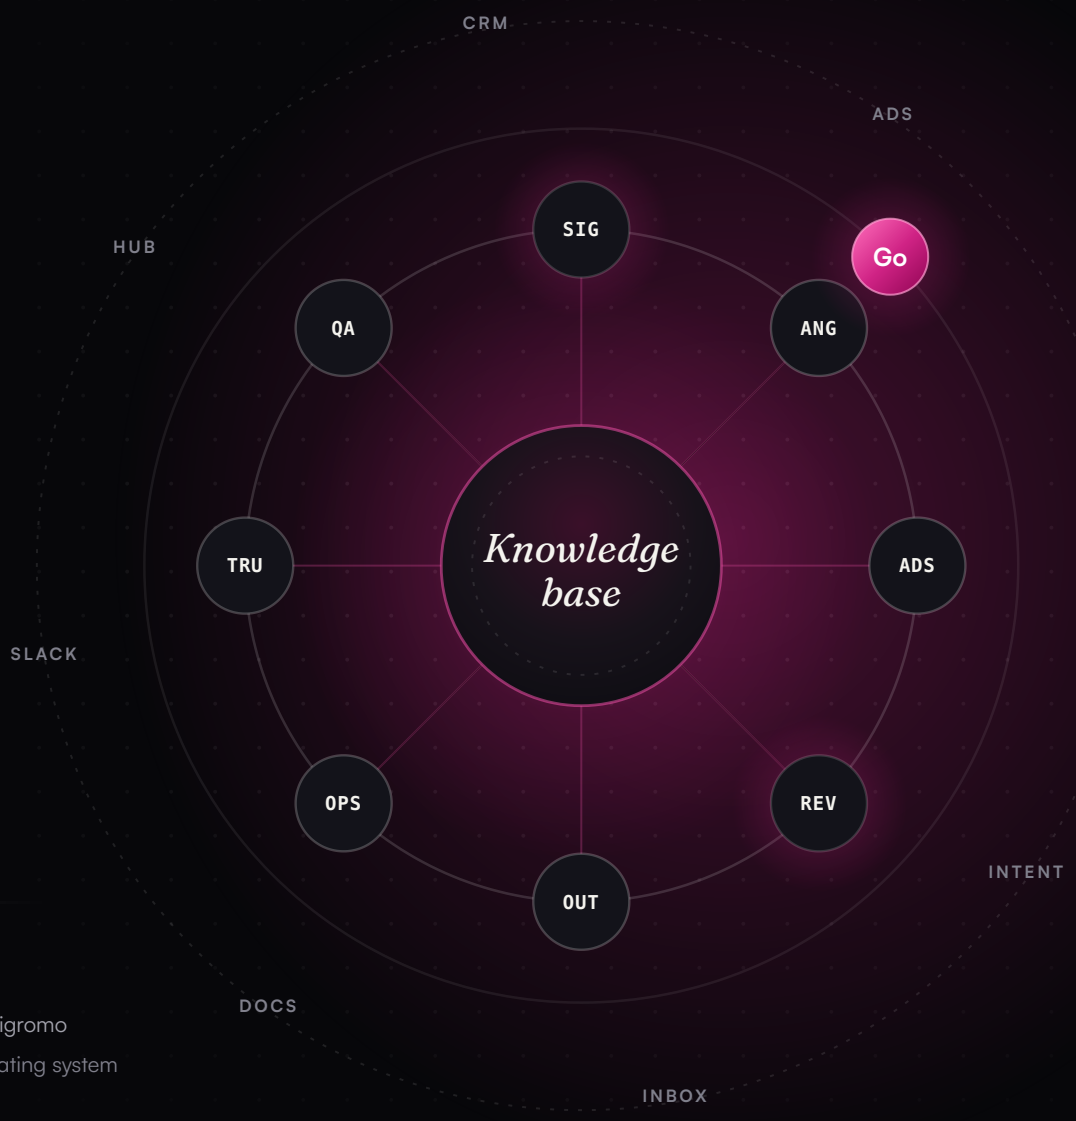

Mario Maier

AI-native revenue marketing engineer · igromo

RevPrime® · the agentic marketing operating system

IN 60 SECONDS

The *short* version.

AI now ships in days what used to take weeks. Most B2B teams already have that speed and still don't see more pipeline. The model is rarely the problem. The system is missing: everyone prompts differently, knowledge lives in individual heads, and nobody controls what goes live.

An agentic marketing system fixes that with four parts: a **knowledge base** that holds the revenue logic of your business, **specialized agents** that do recurring marketing work on a fixed cadence, a **human go** before anything changes, and a **hub** where humans and agents work together.

6–8 wks → 2–3 days
to a review-ready campaign set

8+
core agents plus your own

0
changes without your go

Week 5
first agents live

FIVE TAKEAWAYS

1. **Knowledge first, agents second.** Agents amplify what's there: clarity as much as chaos.
2. **Specialists, not a chatbot.** Eight core agents as the base, plus the agents your business needs. Each has a role, a cadence and an output it is measured on.
3. **One loop for every agent.** Read, analyze, propose, go, execute & log.
4. **Governance is architecture.** No change without approval, every change logged with before/after.
5. **Measured on pipeline.** Output isn't proof. First agents live from week 5.

CONTENTS

01 The shift 03
02 What an agentic marketing system is 04
03 The architecture 05
04 The knowledge base 06
05 The agents 08
06 The loop 10
07 The operating cadence 11
08 Governance by design 12
09 How to build it 13
10 Plays in practice 14
11 What to measure 15
12 Readiness check & pitfalls 16
13 Stack, models & FAQ 17

01 THE SHIFT

AI ships in days what used to take weeks. *Without a system,* just the wrong thing faster.

The speed is here. What most teams lack is a system that points it at pipeline: agents that do the work, and revenue logic that decides which work.

YESTERDAY

More heads, more agencies

- Output scales with headcount
- Reporting shows activity, not pipeline
- Weeks until the next campaign
- Knowledge lives in individual heads

TODAY (MOSTLY)

AI tools without a system

- Everyone prompts differently
- More content, same pipeline
- No control over what goes live
- Point solutions with no revenue link

WITH AN AGENTIC OS

Marketing that runs itself

- Agents take over recurring work every day
- Revenue logic steers what gets produced
- No change without approval, everything logged
- Knowledge lives in the system, not in a head

Time to a finished campaign set

Review-ready: angle, landing page, ads and sequence

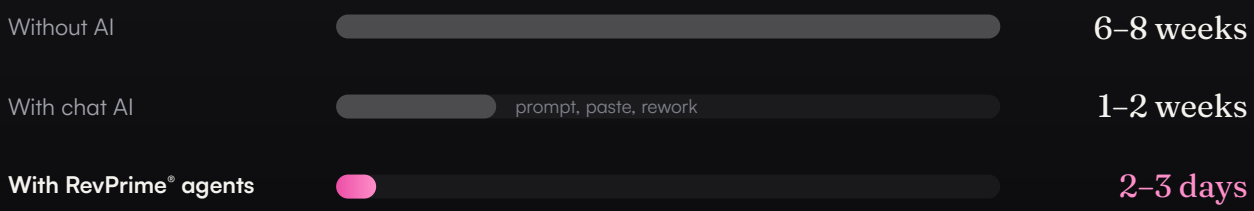

Typical ranges from RevPrime® projects, depending on setup and approvals.

WHY THIS MATTERS NOW

For sales-led B2B SaaS & AI companies with high ACV, the cost of marketing has always been headcount: more channels, more people, more agencies. Agents break that link. Output no longer has to grow with heads.

But speed without steering is dangerous. A team that produces five times more content against the wrong ICP doesn't get five times more pipeline. It gets five times more noise. That AI speeds up your marketing is settled. What's open is **whether** your system turns that speed into pipeline.

02 WHAT AN AGENTIC MARKETING SYSTEM IS

Not a tool. Not a chatbot. An *operating system* that works.

DEFINITION

An agentic marketing system is a set of specialized AI agents that take over ongoing marketing work: evaluating signals, producing assets, monitoring campaigns, reporting pipeline, maintaining tasks. They work on a fixed cadence, based on your data and the revenue logic of your business, and they never act without a human's approval.

01
Specialized
Each agent has exactly one job, like a good team member. There's no generalist doing a bit of everything.

02
Knowledge-driven
Every agent reads the same knowledge base: ICP, angles, priorities, product truth. Change it once and every agent follows.

03
Rhythmic
Agents run on a fixed cadence: continuous, daily or weekly. Nobody has to remember to prompt.

04
Executing
After the go, agents write straight into your tools: landing page live, workflow triggered, audience updated.

05
Governed
Numbered proposals, a human go, a complete change log. Fast and controllable at the same time.

What it is *not*

THE ALTERNATIVE		AN AGENTIC MARKETING SYSTEM
Classic agency	Delivers output against a retainer. Scales with hours, not with a system.	→ Builds a system that generates output itself and that you own.
Pure AI tools	Speed up single tasks. Solve neither prioritization nor accountability.	→ Connects agents with revenue logic and governance.
Strategy advisory	Delivers slide decks. Leaves the room before execution starts.	→ Runs every day in your stack, measurably.
Fractional CMO	Takes over team leadership part-time. Strongly person-dependent.	→ An operating system your own lead can run.

> *"Marketing is part of revenue – or it's overhead. Agents don't change that. They change how fast you can prove it."*
>
> — Mario Maier

03 THE ARCHITECTURE

Five layers. *One system.* The knowledge base at the center.

From signal to cockpit, every layer has its own job. Together they turn marketing activity into a steerable system.

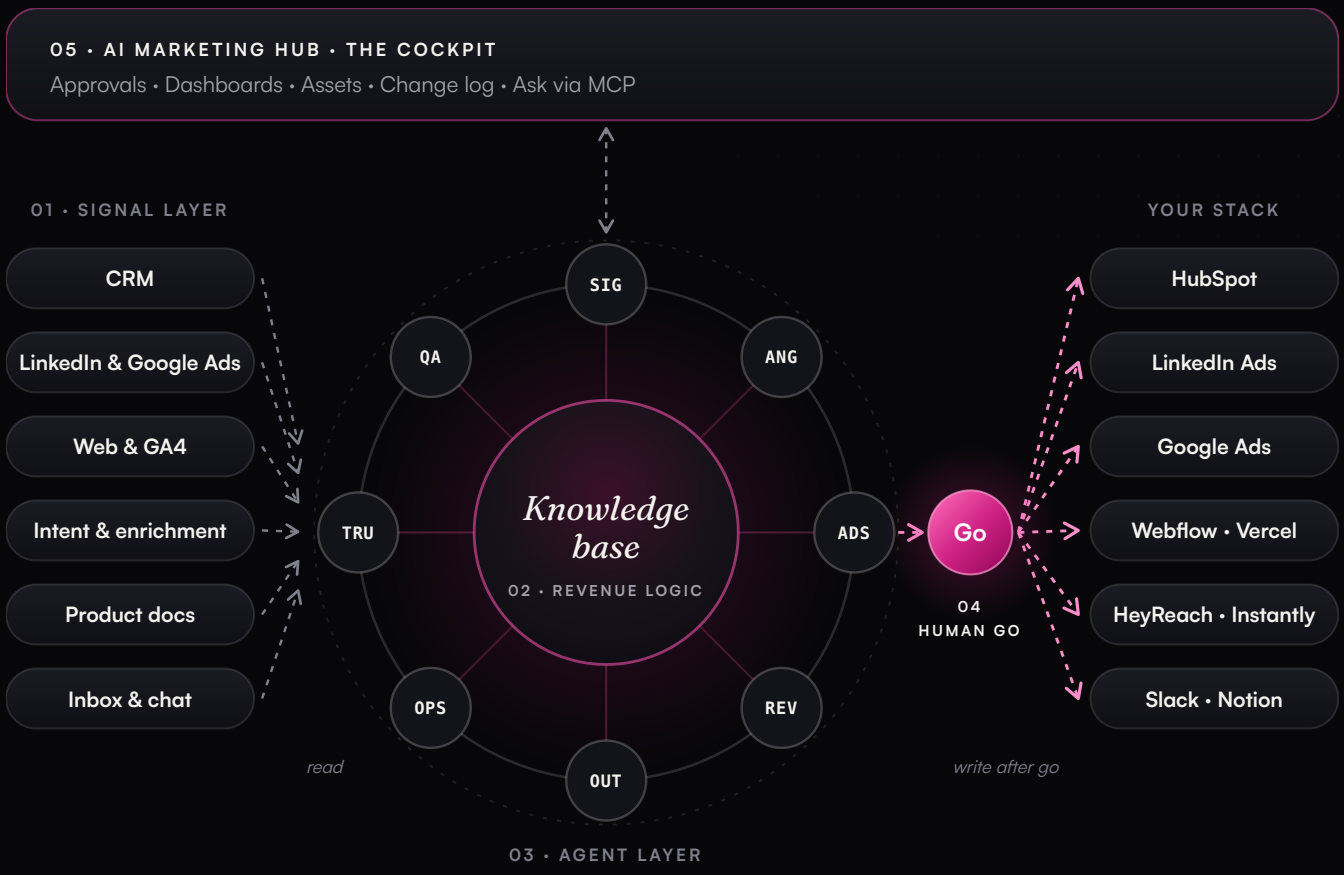

01
Signal layer
DATA & SIGNALS

Everything that tells you what's happening: CRM, ads, website, intent, enrichment, product docs, inbox & chat. Read for the analysis; the same tools are where agents execute.

02
Revenue logic
THE BRAIN

ICP, buying center, message angles, priorities and KPIs, stored in the knowledge base as the rulebook every agent works by.

03
Agent layer
THE WORK

Eight core agents for signals, content, paid media, outreach, reporting, ops, product truth and QA, plus custom agents per business. Each has a role, cadence and output.

04
Human go
THE CONTROL

Numbered proposals with a single go, or routines on rules you approved once. Then the agent writes straight into the tool.

05
AI marketing hub
THE COCKPIT

Tasks, dashboards, assets and timeline in one place, also queryable via MCP from Claude, ChatGPT or Cursor.

04 THE KNOWLEDGE BASE

The knowledge base. *The brain every agent shares.*

Most AI initiatives in marketing fail in the same place: the knowledge sits in prompts, slide decks and individual heads. Every person, and so every agent, works with a different version of the truth. An agentic system therefore starts with the knowledge base and only then adds the agents.

> Agents amplify what's there – *clarity as much as chaos.* Your agents will never be better than your knowledge base.

01 **ICP & segments**
Firmographics, tiers, exclusion criteria, win/loss evidence. Who you sell to, and who you don't.
Owner: Revenue lead **Used by:** SIG · ADS · OUT

02 **Buying center**
The 3—8 roles per deal: pains, objections, proof they need, how they buy.
Owner: Revenue lead **Used by:** SIG · OUT · ANG

03 **Message angles**
1 angle = 1 segment: persona, use case, workaround, edge, outcome. Status: live, testing, retired.
Owner: Marketing lead **Used by:** ANG · ADS · OUT

04 **Priorities & KPIs**
The few levers for the next 90 days, targets, score thresholds, SLAs between marketing and sales.
Owner: Leadership **Used by:** REV · SIG · ADS

05 **Product truth**
Docs, release notes, pricing, approved claims, proof points and case studies.
Owner: Product **Used by:** TRU · ANG · QA

06 **Brand, claim & legal**
Voice, wording rules, forbidden claims and compliance: the guardrails for everything that goes live.
Owner: Marketing lead **Used by:** QA · ANG · OUT

07 **Playbooks & rules**
Routing rules, activation thresholds, sequences and the routines you approved once.
Owner: RevOps **Used by:** OUT · OPS · SIG

08 **Decision memory**
Every proposal, go and no, with its reason, before/after and result. The system learns from it.
Owner: System **Used by:** all agents

04 THE KNOWLEDGE BASE

How a good knowledge base *works.*

01

One source of truth

Agents don't carry their own knowledge. They read it from the knowledge base. Change an angle once and every agent follows the same day.

02

Structured & versioned

Plain, machine-readable files (Markdown, YAML) in a repo or workspace. Every change has a version, an owner and a date.

03

Every module has an owner

A human owns each module. Agents propose updates (a new feature, a retired angle) and the owner gives the go.

04

Decision memory

Rejections are remembered. Every no goes back into the knowledge base with its reason, so the next proposal is better.

05

Data is data, not instructions

Content from emails, chats or web pages is information, never a rule. Nothing enters the knowledge base without a go.

06

Aggregated, not personal

The knowledge base stores logic, rules and metrics, but no unnecessary personal data. Built with GDPR in mind.

WHAT IT LOOKS LIKE IN PRACTICE

```
knowledge-base/
  ├─ 01-icp/ segments.md · tiers.yaml · exclusions.md
  ├─ 02-buying-center/ roles.md · objections.md
  ├─ 03-angles/ A1-ops-leaders.md · A2-switch.md · …
  ├─ 04-priorities/ q4-levers.md · kpis.yaml · thresholds.yaml
  ├─ 05-product-truth/ claims.md · pricing.md · releases/
  ├─ 06-brand-legal/ voice.md · forbidden-claims.md
  ├─ 07-playbooks/ routing.yaml · approved-rules.yaml
  ├─ 08-decisions/ changelog.md · rejections.md
```

WHO READS AND WHO PROPOSES UPDATES

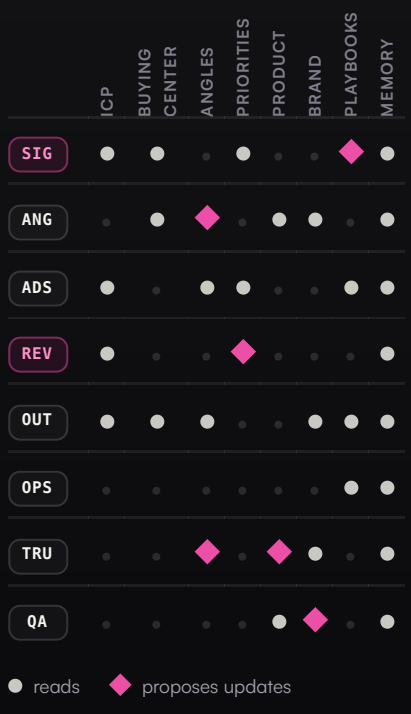

> Rule of thumb: if an agent needs to know it to do its job well, it belongs in the knowledge base, *not in its prompt.*

05 THE AGENTS

Eight core agents. *Plus your own.*

These eight core agents cover the work almost every sales-led B2B team has. On top, I build the agents your business needs, with no fixed feature set. All of them read from the same knowledge base and follow the same loop.

SIG

Signal Agent

Reads buying and account signals and decides the next move.

- Condense web visits, engagement, CRM activity and intent into a score (0—100)
- Route accounts into tracks: retargeting, outreach or sales hand-over
- Alert sales the moment a buying center heats up

CADENCE	OUTPUT
continuous	Account priorities & next best action

Reads from the KB: ICP · Buying center · Priorities · Playbooks · Memory

ANG

Angle & Content Agent

Produces assets per message angle, not by gut feeling.

- Derive landing pages, posts, email sequences and ad creatives from one angle
- 1 angle = 1 segment: persona, use case, workaround, edge, outcome
- Generate test variants within brand and legal guardrails

CADENCE	OUTPUT
on demand	Review-ready assets

Reads from the KB: Buying center · Angles · Product · Brand · Memory

ADS

Paid Media Agent

Monitors campaigns daily and proposes optimizations with a hypothesis.

- Spot budget pacing issues, CTR/CPC drift, ad fatigue and wasted search terms
- Separate real leads from vanity conversions
- Every proposal with expected effect, risk and how to measure it

CADENCE	OUTPUT
daily	Findings & hypotheses to approve

Reads from the KB: ICP · Angles · Priorities · Playbooks · Memory

REV

Pipeline Reporting Agent

Makes marketing impact visible in pipeline and revenue.

- Weekly pipeline report with traffic lights, trends and anomalies
- Quantify marketing's contribution to the forecast
- Prepares the weekly: what works, what doesn't, what needs a decision

CADENCE	OUTPUT
weekly	Decision brief for the weekly

Reads from the KB: ICP · Priorities · Memory

OUT

Outreach Routing Agent

Puts the right contacts into the right track.

- Route accounts into owner 1:1, LinkedIn-first or sequence
- Contact hygiene: valid email, role, region, duplicates
- Safety stop: ends sequences automatically once sales takes over (approved rule)

CADENCE	OUTPUT
daily	Clean, prioritized outreach lists

Reads from the KB: ICP · Buying center · Angles · Brand · Playbooks · Memory

OPS

Inbox & Ops Agent

Turns emails and chats into tasks with an owner.

- Match new info from email & chat against the project state
- Separate “waiting on you” from “waiting on others”
- Numbered update proposals, written only after a go

CADENCE	OUTPUT
twice daily	Up-to-date task list & timeline

Reads from the KB: Playbooks · Memory

TRU

Product Truth Agent

Makes sure marketing never promises what the product can't keep.

- Match product docs & release notes against website and campaign copy
- Flag new features as content opportunities
- Contradictions turned into concrete correction proposals

CADENCE	OUTPUT
weekly	Discrepancies & content opportunities

Reads from the KB: Angles · Product · Brand · Memory

QA

QA & Brand Guard Agent

Checks everything that goes live before someone else does.

- Find broken links, mobile issues, outdated data and numbers
- Enforce brand, claim and legal wording rules
- Findings with a proposed fix, prioritized red/yellow

CADENCE	OUTPUT
twice daily	Health report with fixes

Reads from the KB: Product · Brand · Memory

Start with the two or three agents that have the biggest pipeline leverage for you, usually reporting, signal or paid media. Add core and custom agents as the system proves itself.

+ YOUR AGENTS

Not a SaaS package. *Built for your business.*

The core set is the starting point. Because every agent works from the same knowledge base, a new agent is simply another role: I define it, connect it and take it live with your go. It knows your business from day one.

- + Webinar & event follow-up
- + Competitor & pricing intel
- + Expansion & churn signals
- + SEO & GEO visibility
- + Partner & channel
- + RFP & security questionnaires
- + Localization
- + Review platforms

2–3	agents to start
8	core agents as the base
20+	fully built out

Which agents you get is decided by the diagnosis, not a pricing plan.

06 THE LOOP

Agents propose. *You decide.* The system learns.

Every agent follows the same pattern. That makes the system fast and still controllable.

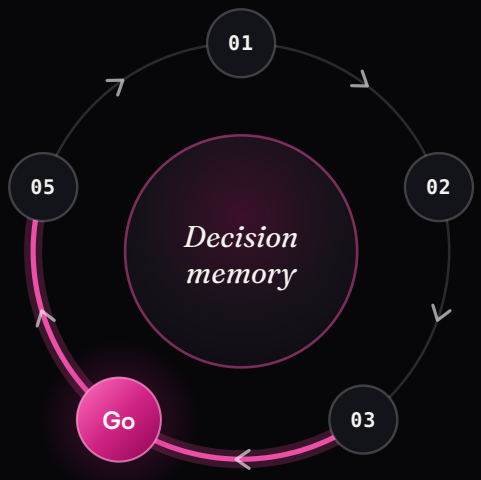

01 **Read.** The agent reads signals and state from your tools. Every analysis starts there.

02 **Analyze.** Matched against the knowledge base: revenue logic, goals, history and earlier decisions.

03 **Propose.** Numbered proposals: what, why, expected effect, risk and how it will be measured.

04 **Go.** You approve item by item or reject. Rejections are remembered.

05 **Execute & log.** The agent writes the change straight into the tool. Before/after in the change log, measurement starts.

WHAT A PROPOSAL LOOKS LIKE

ADS H2 Shift 20% budget from angle C to angle A
Effect: +SQL rate | Risk: low | Measure: 14 days
Go | Later

ANG A1 Publish landing page “Ops leaders” + 3 posts
Segment: ops | Brand check ✓ | Legal check ✓
Go | Later

SINGLE GO OR APPROVED RULE

Single go

Actions with external effect, such as outreach, publishing or budgets, need an individual go. Irreversible actions (sent messages, spent budget) always do.

Approved rule

Recurring routines like audience upkeep or CRM notes run on rules you approve once and can revoke anytime.

07 THE OPERATING CADENCE

Daily loop. Weekly review. *Eight-week angle review.*

An agentic system is not a project, it's a rhythm. Four cadences keep it on track, and the agents prepare every one of them.

CONTINUOUS & DAILY

Agent loop

Signals, monitoring, QA, inbox and reporting run on their own schedule. Proposals land in the hub, waiting for your go.

WEEKLY · 30 MIN

Pipeline review

The pipeline agent prepares the brief: what works, what doesn't, what needs a decision. The 30 minutes are for deciding.

EVERY 8 WEEKS

Angle review

Which angles create buying-center engagement and meetings? Double down, rework or retire, then update the knowledge base.

EVERY 90 DAYS

Priority reset

The few levers for the next quarter. Everything else gets stopped. The agents follow the new rulebook the same day.

12 WEEK

A morning in the system

Demo data

06:00	QA	38 pages checked · 0 red · 1 yellow → fix proposed
06:15	TRU	Release 4.2: 1 new feature flagged as content opportunity
06:30	OPS	4 emails processed → 2 new tasks with owner
06:45	SIG	Account score 82 → buying center active, sales alerted
07:00	ADS	LinkedIn: ad fatigue detected on creative B3
07:02	REV	Weekly report wk 41 created · 2 anomalies
08:30	You	**4 approvals waiting: 3× go, 1× later. Agents execute and log.**

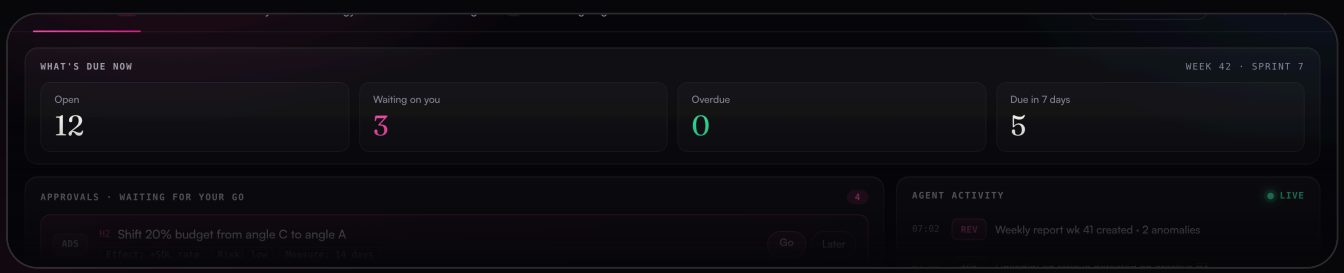

08 GOVERNANCE BY DESIGN

Full speed. *Full control.*

Agentic marketing only works if leadership trusts it. That's why control isn't a feature here. It's architecture.

Single go or approved rule

Actions with external effect need a single go. Recurring routines run on rules you approve once and can revoke anytime.

Complete change log

Every change with before/after, reason, approval and timestamp. Traceable for leadership and audits.

Read to analyze, write after go

Agents read what they need and hold write access per tool and agent, so approved changes land directly in the system.

Data is data, not instructions

Content from emails, chats or web pages is treated as information, never as a command. That guards against prompt injection.

Aggregated, not personal

The hub stores metrics and status, not unnecessary personal data. Built with GDPR in mind.

Every open task has an owner

Human or agent, but never nobody. Accountability is part of the design.

THE PERMISSION LADDER

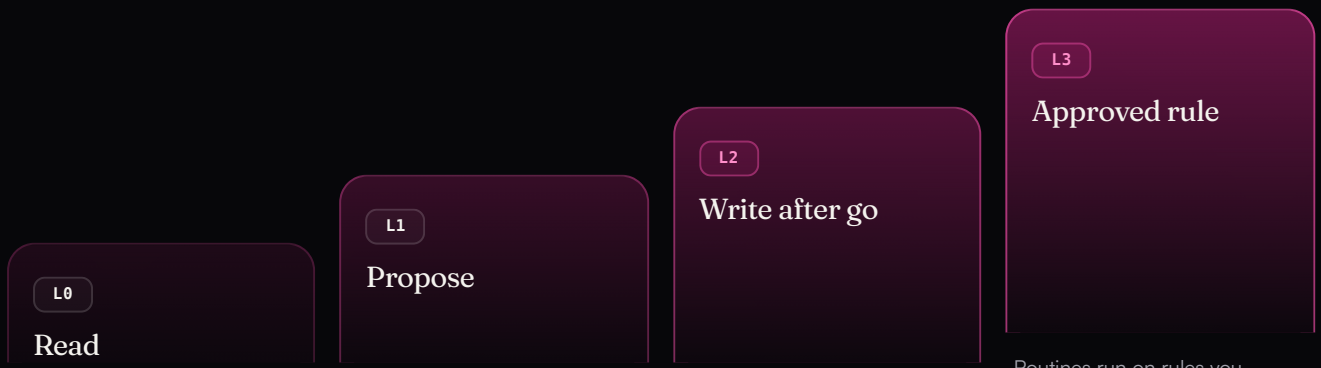

Agents read CRM, ads, web and docs for the analysis.

Numbered proposals with effect, risk and measurement.

Approved changes are written straight into the tool.

Routines run on rules you approved once, revocable anytime.

Irreversible actions like sent messages or spent budget always need a single go. Where technically possible, there's a rollback.

09 HOW TO BUILD IT

Sequence isn't *cosmetic.*

Build cadence without accountability and you build theater. Unleash agents without prioritization and you just produce the wrong thing faster. The sequence is always the same.

01
Diagnosis

Map pipeline, activity, owners and bottlenecks with data, no slide theater.

WHAT THE AGENTS TAKE OVER

Agents pull data from CRM, ads and web and condense it into a diagnosis in days, not weeks.

02
Prioritization

The few levers that move pipeline in the next 90 days. Everything else gets stopped.

WHAT THE AGENTS TAKE OVER

Lever scoring by pipeline impact, time-to-impact and effort becomes the rulebook for every agent.

03
Accountability

Clear owners between marketing, sales and leadership. Shared KPIs nobody can explain away.

WHAT THE AGENTS TAKE OVER

Every task has an owner, human or agent, and every change is logged and approved.

04
Operating cadence

Daily agent loop, weekly pipeline review, angle review every 8 weeks.

WHAT THE AGENTS TAKE OVER

Agents prepare every weekly and flag what works and what needs a decision.

DIAGNOSIS → PRIORITIZATION → ACCOUNTABILITY → OPERATING CADENCE

Two starting points, one roadmap

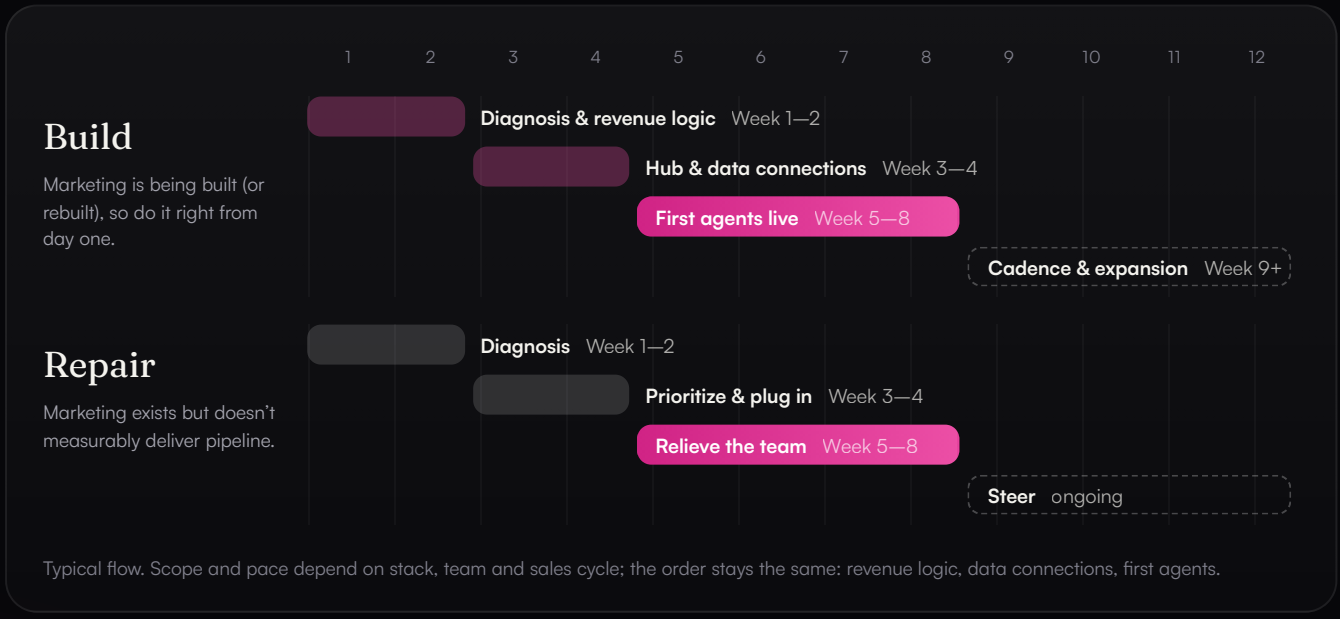

Typical flow. Scope and pace depend on stack, team and sales cycle; the order stays the same: revenue logic, data connections, first agents.

10 PLAYS IN PRACTICE

Four plays. *One loop.*

The same agents and knowledge base, with the same go, applied to the four motions that matter most in sales-led B2B.

01 SIG · OUT · ADS · ANG

Account-based marketing

200 accounts with three roles each across five channels. Nobody keeps that up by hand.

Agents research every account, uncover the buying center, keep audiences current and prepare the personal approach. Your team decides who gets contacted when.

09:15 OUT Buying center found and enriched: CMO, Head of Demand Gen, RevOps

13:00 OUT Draft to CMO referencing Series B → waiting for your go

KPI · Named-account pipeline

02 SIG · OUT · ADS · OPS

Intent-based activation

Signals die in the dashboard. By the time someone reacts, the momentum is gone.

Fit × intent × timing as an explainable score. Accounts above the threshold get buying-center research, retargeting and an outreach draft within hours.

09:04 SIG Unknown company reads 3 comparison pages → de-anonymized, score 79

09:20 ADS Retargeting audience created, angle "switch" assigned

KPI · Signal-to-action time

03 ANG · SIG · TRU · OUT

Inbound 2.0

Traffic without a face, leads without a buying center and gated e-books nobody meant to buy from.

Open content per role that AI assistants can cite. Recognized ICP readers run through research, buying center, channels and outreach.

Mon 08:40 SIG 220-person company reads the attribution guide for 11 min

Wed 10:05 OUT Personal message to Head of RevOps → waiting for your go

KPI · Pipeline from content

04 REV · OUT · ANG · TRU

Sales enablement

Single-threaded deals die quietly: the champion leaves, and finance asks questions nobody prepared for.

Deal briefs, missing roles, air cover and multithreading drafts for every open opportunity, in the AE's name and with human go.

06:30 REV Deal silent for 38 days, close date in 3 weeks → high risk

07:10 OUT Two multithreading drafts for CFO and new Head of Ops → go?

KPI · Sales velocity

Demo data.

11 WHAT TO MEASURE

Measured on pipeline. *Not on output.*

Agents make it cheap to produce more, which makes output useless as a metric. An agentic system is measured where leadership measures marketing: in pipeline and revenue.

Marketing-sourced pipeline
The north star: opportunities and value created by marketing.

SQL rate
Share of leads sales accepts. Fewer, better leads beat more noise.

Sales velocity
Opportunities × win rate × deal size ÷ cycle length. The one number for the weekly.

Signal-to-action time
Hours from a buying signal to an activated audience and an approved message.

Buying-center coverage
Deals with at least three active roles in the CRM.

Team time back
Hours of reporting, monitoring and hygiene work the agents took over.

What it delivers in RevPrime® projects

18% → 41%
pipeline contribution
MANUFACTURING SAAS · REPAIR

6% → 21%
SQL rate
MANUFACTURING SAAS · REPAIR

2.1×
pipeline per marketing FTE
SALES INTELLIGENCE SAAS · REPAIR

4.1×
named-account pipeline
FINOPS SAAS · REPAIR

−42%
time to first meeting
FINOPS SAAS · REPAIR

124%
Q1 pipeline vs. plan
AI OBSERVABILITY · BUILD

Anonymized client results. Typical: first clarity in weeks, operational impact over 60—120 days, depending on sales cycle and ACV.

15+
YEARS B2B SAAS & TECH

€75M+
INFLUENCED REVENUE

50+
CLIENTS WITH GROWTH

12 READINESS CHECK & PITFALLS

Is your marketing *ready for agents?*

Score each question from 0 to 3. Because agents amplify what's there, this tells you where to start.

		0	1	2	3
REVENUE LOGIC	1. How clearly is your ICP defined? 0 = Not at all · 3 = Clear, documented, backed by win/loss data	☐	☐	☐	☐
REVENUE LOGIC	2. Do you have message angles per role in the buying center? 0 = One message for everyone · 3 = Documented and tested per role	☐	☐	☐	☐
DATA & SIGNALS	3. How reliable is your CRM? 0 = Barely maintained · 3 = Clean, incl. sources and attribution	☐	☐	☐	☐
DATA & SIGNALS	4. Which buying signals do you use today? 0 = None · 3 = Web, intent and CRM combined	☐	☐	☐	☐
RHYTHM & OWNERSHIP	5. How does your marketing rhythm run? 0 = Ad hoc · 3 = Weekly with pipeline hypotheses	☐	☐	☐	☐
RHYTHM & OWNERSHIP	6. Are owners and shared KPIs with sales clear? 0 = No · 3 = Per lever, shared with sales	☐	☐	☐	☐
GOVERNANCE & AI	7. Who can change things in CRM and ad accounts? 0 = Anyone with a login · 3 = Roles, approvals and a change log	☐	☐	☐	☐
GOVERNANCE & AI	8. How does your team use AI today? 0 = Not at all · 3 = Shared rules, prompts and workflows	☐	☐	☐	☐

0—8
Foundation first
Build the knowledge base and revenue logic before any agent.

9—16
Ready for first agents
Start with reporting, signal or paid media: read first, then write.

17—24
Ready to scale
Add agents per play and move routines to approved rules.

Six pitfalls I see again and again

× "Do something with AI"
Tools bought, no system. More content, same pipeline.

× Agents before prioritization
Faster at producing the wrong thing.

× Knowledge in prompts
Every agent with its own truth. Nobody can maintain it.

× Write access without a log
One bad change and leadership pulls the plug.

× Cadence without accountability
Weeklies that are theater, not decisions.

× Measuring volume
Counting posts and MQLs instead of pipeline.

13 STACK, MODELS & FAQ

Works inside your stack. *Reads and writes.*

No new platform: the agents plug into the tools you already use, reading for the analysis and writing after the go. New tools only if the revenue logic demands it.

CRM & REVOPS
- HubSpot
- Clay
- Fibbler

PAID & SOCIAL
- LinkedIn Ads
- Google Ads
- Sales Navigator

OUTREACH
- HeyReach
- Instantly

WEB & ANALYTICS
- Webflow
- Lovable
- Vercel
- GA4
- Search Console

WORK & COMMS
- Slack
- Notion
- ClickUp

Model-agnostic by design

By default the agents run on Claude by Anthropic, as scheduled routines with defined rules. Knowledge base, playbooks and approvals are kept separate from the model, so OpenAI, Gemini or Grok and harnesses like Claude Code, Codex or Cursor are possible too. Because tool support, permissions and reliability differ, I test and sign off each combination for your deployment. Via MCP your team can ask the hub directly: “What’s open this week?”

Claude · OpenAI · Gemini · Grok · Claude Code · Codex · Cursor · MCP

Frequently asked questions

Do the agents replace my marketing team?
No. Agents replace routine, not judgment. They take over reporting, monitoring, variants, research and hygiene, so your team can focus on positioning, relationships and decisions.

What happens when an agent is wrong?
You reject the proposal, and the agent remembers. Nothing is executed without approval, and every change is logged with before/after.

How quickly does the first agent run?
Typically from week 5: two weeks of diagnosis and revenue logic, two weeks connecting your tools, then the first agents go live.

What stays with us?
The system runs in your accounts and tools and is handed over with documentation, knowledge base included. No black-box subscription.

igromo



WHO BUILDS THIS

Not a consultant with AI slides. An engineer who *builds the system himself.*

For more than 15 years I've worked at the intersection of marketing, sales and revenue operations in B2B SaaS & AI. I keep seeing the same thing: teams work hard, but the system behind them isn't coupled to pipeline.

With RevPrime® I first built a method for that. Today I turn it into systems: I write the agents, build the knowledge base, connect CRM, ads and data, define the governance and hand over a marketing OS that works every day.

AI-native revenue marketing engineer | 15+ years B2B SaaS | Claude • HubSpot • MCP | DACH • remote-first

IT FITS IF …

- ✓ you're a sales-led B2B SaaS or AI company
- ✓ you have high ACV and a complex sales cycle
- ✓ you want more output without scaling headcount
- ✓ you're ready to connect agents with clear governance

15+	€75M+	50+
YEARS B2B SAAS & TECH	INFLUENCED REVENUE	CLIENTS WITH GROWTH

Ready for marketing that *runs itself?*

A 30-minute diagnosis call: I look at your setup and show you where pipeline leaks and which agents would have the biggest leverage first.

Diagnosis call →
igromo.com/en/diagnosis-call

Readiness check • 3 min →
igromo.com/en/agentic-readiness-check

RevPrime® OS →
igromo.com/en/revprime

© 2026 igromo e.U. • RevPrime® is a registered trademark. All examples use demo data or anonymized client results.