

Nagent

PLAYBOOK · EDITION 2 · OCTOBER 2026

FOR HEADS OF GROWTH, MARKETING AND SALES

The AI-native *GTM playbook*

How a growth team runs the whole funnel, from being found to being renewed, with people and AI teammates in one room, on one number, with every action on one record.

COVERS	INCLUDES	PLUS	READING TIME
10 funnel stages	90-day adoption plan	Scorecard and readiness check	About 35 minutes

CONTENTS

00 Foreword 3
GTM was never a headcount problem

PART ONE · THE OPERATING MODEL

01 The multiplayer growth org 5
Who does the work, who holds the authority

02 Six principles of AI-native GTM 6

03 The autonomy ladder 7
How AI teammates earn the right to act

04 The funnel on one page 8

PART TWO · THE FUNNEL, STAGE BY STAGE

05 Ten stages, from context to expansion 10
Foundations, search, content, paid, social, capture, outbound, meetings, deals, customers

PART THREE · RUNNING IT

06 The control plane 21
Budgets, tools, memory and the record

07 How the team changes 22

08 The scorecard 23

09 The 90-day adoption plan 24

10 Readiness check and failure modes 25

11 Proof from the field 26

WHO THIS IS FOR

Anyone who owns a growth number: a head of growth, CMO, CSO or founder, and the team leads who run marketing, sales and content under them.

It assumes you already have a funnel and a team. It does not assume you have any AI in production yet.

HOW TO READ IT

Part one is the model: how a team of people and AI teammates is organised, and how trust is earned rather than switched on.

Part two walks the funnel. Each stage has the same shape: what breaks in the old playbook, the AI-native play, who does what, the one approval gate, and the numbers to set.

Part three is how to run it: budgets, roles, the scorecard, a 90-day plan and a readiness check you can score this week.

A NOTE ON HONESTY

Every capability here is described as it ships in Nagent in October 2026. Anything in beta or coming soon is marked as such. Customer figures are taken from published customer stories. Third-party research is cited on the last page.

00 · FOREWORD

GTM was never a *headcount* problem

For twenty years the answer to "we need more pipeline" has been "hire": another SDR, another content writer, another agency. Each hire adds capacity and, quietly, adds coordination. That second cost is the one that kills growth teams.

Look at where the hours actually go in a typical B2B growth team. A lead fills a form and waits for someone to notice. A rep spends the morning researching accounts before writing a single email. A demo is booked, prepped by hand, held, and then the recap is written two days later. Content is drafted, waits in a review queue, is published, and never reaches the people who asked the question it answers.

Most of these jobs are **coordination, not judgement**: moving context from one system and one person to the next. Coordination is exactly what AI teammates are good at, and judgement is exactly what your people are for.

The first wave of AI in GTM got this backwards. It handed every person a private assistant that writes faster. Output went up; the number did not. Ten people with ten chatbots are still ten people passing context to each other by hand, and now they have more drafts to review.

The AI-native playbook starts somewhere else. It treats the growth team as **one room**: people and AI teammates in the same shared context, working the same number, with every action, approval and result on one record. The AI teammates do the work end to end, all the way to pipeline and customers, not just content. People set direction, approve what goes out, and handle the moments that need a human.

Two rules keep that safe. AI teammates **earn autonomy** one level at a time, on a record of approved work, and lose it automatically when quality slips. And every agent runs inside a **budget**, so cost is a control you set, not a surprise you discover.

> "We stopped buying deliverables. We started buying the number."
>
> Ashutosh Modi, Emami

7×

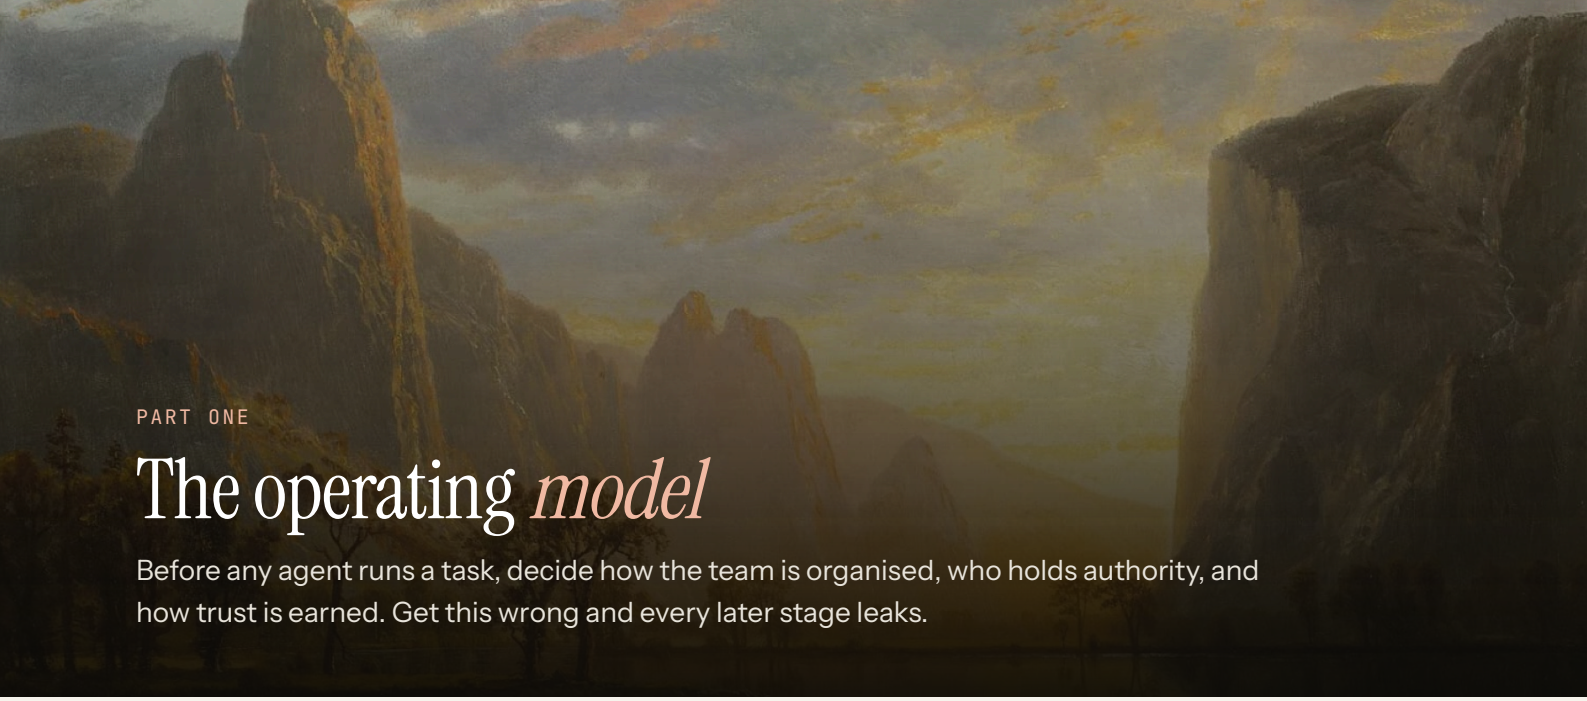
more likely to qualify a lead when the first contact happens within an hour. Most firms take far longer. (HBR)

68%

of US Google searches ended without a click in early 2026, up from 60% in 2024. Being the answer now matters as much as ranking. (SparkToro)

+41%

visibility in AI answers from adding quotations; keyword stuffing scored −8%. Evidence beats packaging. (GEO, KDD 2024)



PART ONE

The operating *model*

Before any agent runs a task, decide how the team is organised, who holds authority, and how trust is earned. Get this wrong and every later stage leaks.

WHAT CHANGES

From tools people use to teammates people manage

In the old model, software is a tool: a person opens it, does a task, closes it. In the AI-native model, an AI teammate holds a **standing job** (keep us visible in AI answers, keep the outbound pipe full, answer every inbound in minutes) and does it continuously, inside limits you set.

That changes the manager's job. You stop asking "who is going to do this?" and start asking three better questions:

- **What is the outcome?** Qualified meetings, cited pages, resolved conversations.
- **What must a person approve?** One gate per workflow, named for who approves it.
- **What can it spend?** A daily and monthly budget, checked before every call.

WHAT DOES NOT CHANGE

People hold the authority

Your people still own the number, the brand line and every relationship. A phone call or a chat message never approves anything; only a person pressing approve does, and the reason goes on the record.

Destructive and governance actions (deleting, redirects, changing an agent's limits, spending past a cap, creating agents) always wait for a person, at every autonomy level, forever.

Multiplayer AI, defined

People and AI teammates working in the same room, on the same number, with every action on one record. The record of who did what, and who approved it, is also the audit trail.

01 · THE MULTIPLAYER GROWTH ORG

An organisation, *not an assistant*

The AI-native growth team is shaped like a real org chart, because that is how accountability works. Each person talks to one teammate. Under it, an AI chief of staff runs each department with team leads who do the work, and your people work beside them.

Head of growth
CGO, CMO, CSO or CCO · owns the number

N · head of the growth org
The one teammate each person talks to, by voice or text. Hands work to the right chief of staff. BETA



MIRA
AI Marketing Chief of Staff
Owns reach and pipeline cost

Team leads
DRIS search and AI visibility · NIA paid · HOOK paid creative BETA · MOXA social EARLY ACCESS · Echo reputation BETA



SERA
AI Sales Chief of Staff
Owns qualified meetings

Team leads
NORA outbound · CORA CRM · VANI and LIPI email writers BETA · DEXA deals and RIVA revenue ops SOON



CREA
AI Content Chief of Staff
Owns output and the brand line

Team leads
Blog Studio · RUPA images · FOLIO long-form · Alpha video · LANDON landing pages



Convexa
Customer experience
Owns resolution and conversion

Channels
Chat, voice and video on your site; books meetings; pushes qualified leads to the CRM

Your people
Marketers, reps, editors and account owners work in the same rooms as the agents, approve what goes out, and take the conversations that need a human.

Why a chief of staff per department

A chief of staff plans an outcome brief and routes each piece to the teammate that owns it. Handoffs stay inside the room's roster, and the work runs under the receiving agent's own autonomy level and budget. Nobody gets a blank "do marketing" agent.

Why one teammate to talk to

Context dies in handoffs. When every person asks one teammate, by voice, email, Slack or the app, the question lands with the right department and the answer lands on the shared record, not in a private chat.

02 · PRINCIPLES

Six principles of *AI-native GTM*

These are the rules that separate teams whose AI moves the number from teams whose AI moves the word count. Every stage in part two applies them.

PRINCIPLE 1

Context before agents

An agent is only as good as what it knows: your ICP, products, proof, brand voice and what was decided last week. Load the shared memory first. An agent with no context writes generic copy at machine speed, which is worse than no copy at all.

PRINCIPLE 2

Outcomes, not outputs

Give each team a number it owns: qualified meetings, cited pages, resolved conversations, pipeline cost. "Posts published" and "emails sent" are inputs. A team measured on inputs becomes a cost centre that looks busy.

PRINCIPLE 3

One gate per workflow

Every workflow has exactly one human step, named for the role that approves it. Not five reviewers, not zero. The approval sits on the same record as the work, so the record is the audit trail.

PRINCIPLE 4

Trust is earned, and lost

AI teammates start supervised and move up one autonomy level at a time on a record of approved work. When dismissals, errors or complaints rise, they drop automatically. Autonomy is a ledger, not a setting.

PRINCIPLE 5

Budgets are the control

Every agent has daily and monthly caps, checked before each call. At high autonomy you stop approving every cost and let the budget do that job. If a cap is hit, the agent queues, pauses or notifies, as you choose.

PRINCIPLE 6

Carry it to the customer

The handoff is where pipeline leaks: from content to subscriber, subscriber to lead, lead to meeting, meeting to deal. AI-native teams design the workflow across the handoff, so the same record follows the buyer from first touch to renewal.

The test. For any AI initiative, ask: which number does it own, who approves its one gate, what can it spend, and what record does it leave? If you cannot answer all four, it is a tool someone will stop using, not a teammate.

03 · THE AUTONOMY LADDER

How AI teammates *earn the right to act*

The single most important design decision in AI-native GTM is not which model to use. It is how much each agent may do without asking. Nagent uses five levels, and every action an agent attempts gets one of three decisions: fire, queue for a person, or refuse.

L1
Observe
Reads and reports. Even drafts wait for a person. Right for a brand-new agent on sensitive work.

L2
Assist
Drafts on its own. Nothing changes and nothing leaves until a person approves.

L3
Operate
Changes records inside the workspace on its own: updates the CRM, moves a card, files a task.

L4
Represent
Acts outside the workspace, such as emailing a contact, within the rules you set for sends.

L5
Lead
Publishes on its own. Reserved for agents with a long, clean record.

Earning a level

- Nominated nightly after **30 days** at its current level, with at least **25 decided cards** and fewer than **5% dismissed**.
- A person decides. Up one step at a time, never jumping.
- Raising to L4 or L5 needs an org admin, chief of staff or team lead. On Enterprise, a raise to L5 needs two people, and nobody approves their own request.

Losing a level

- Over a 14-day window: more than 30% of decided cards dismissed or edited, more than 15% failures of its own making, thumbs-down from three or more people, or two poor outcomes.
- A budget breach drops one level immediately.

Action classes

CLASS	EXAMPLE	FROM
C1 Read	Pull Search Console, read the CRM	L1
C2 Prepare	Draft an email, a post, a brief	L2
C3 Internal act	Update a lead, create a task	L3
C4 External act	Send to a contact, book a meeting	L4
C5 Publish	Post, publish a page, go live	L5
Locked	Delete, redirect, change governance, spend past a cap, create agents	Always a person

Playbook advice. Start every new agent supervised, even if your plan allows more. A month of approved cards is the cheapest training data you will ever produce, and it gives your team a reason to trust the agent instead of a reason to fear it.

04 · THE FUNNEL ON ONE PAGE

Ten stages, *one record*

The whole funnel, with the AI teammates that carry each stage, the person who holds its one gate, and the single number that decides whether it works.

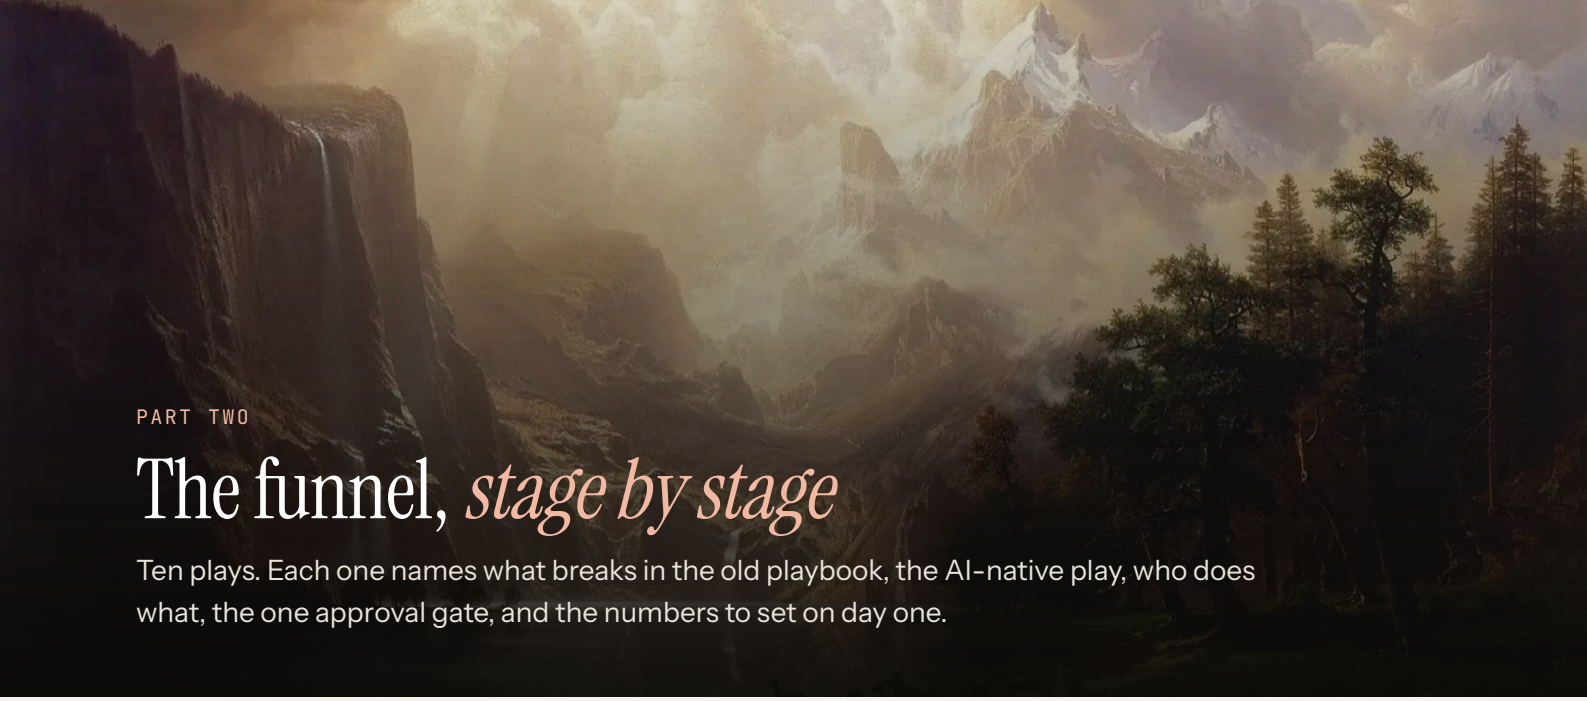
#	STAGE	AI TEAMMATES	PERSON AT THE GATE	THE NUMBER THAT MATTERS
01	**Foundations** ICP, proof, voice, memory	Knowledge Hub, Brand Lock, every chief of staff	Head of growth	Share of agent drafts approved without edits
02	**Be found** Search and AI answers	DRIS	SEO or content lead	Share of tracked questions where you are cited
03	**Content engine**	CREA, Blog Studio, RUPA, FOLIO, Alpha, LANDON	Editor or brand lead	Subscribers and leads generated per piece
04	**Paid demand**	NIA, HOOK BETA	Performance marketing lead	Cost per qualified pipeline, not per click
05	**Social and reputation**	MOXA EARLY ACCESS, Echo BETA, Alpha	Brand or community lead	Engaged ICP followers; review response time
06	**Capture and speed-to-lead**	Convexa, CORA, lead magnets, LANDON forms	Sales lead	Time to first touch on hot leads
07	**Outbound**	NORA, VANI BETA	Sales lead or SDR manager	Qualified outbound meetings booked
08	**Meetings and demos**	SERA, CRM sub-agents, room meetings	Account executive	Demos with an agreed next step in 24 hours
09	**Deals and pipeline**	SERA, CORA, DEXA and RIVA SOON	AE and sales head	Stalled deals with a named next step
10	**Customers and expansion**	Convexa, Echo, CORA, LIPI BETA	Customer success lead	Resolution rate; expansion conversations opened

Read it as a loop, not a line

What customers ask (stage 10) becomes the questions DRIS tracks (stage 02) and the content CREA writes (stage 03). Objections NORA hears (stage 07) become proof on the landing page. The shared record is what lets each stage feed the next.

Do not start everywhere

Almost no team should switch on ten stages at once. Section 09 gives a 90-day sequence. The short version: foundations first, then the one stage where your funnel leaks most, then the stage that feeds it.



PART TWO

The funnel, *stage by stage*

Ten plays. Each one names what breaks in the old playbook, the AI-native play, who does what, the one approval gate, and the numbers to set on day one.

HOW EACH STAGE IS LAID OUT

AI TEAMMATES Who carries the work	PERSON AT THE GATE Who approves the one human step
STARTS AT A sensible first autonomy level	NORTH STAR The one number that decides it

Then: the play as numbered steps, the metrics to set, and the mistakes we see most often.

A WORD ON STARTING LEVELS

The levels suggested in each stage are where we would start a team that is new to agents. They are deliberately conservative. Your plan sets the ceiling (L3 on Solo and Team, L4 on Business, L5 on Enterprise); your agents' records decide how fast they climb toward it.

On a new workspace, every outbound send waits for a person regardless of level. That is a feature: the first few hundred approvals are how your team learns what good looks like, and how the agent learns your taste.

STAGE 01 · FOUNDATIONS

Context before *agents*

The old playbook kept GTM knowledge in people's heads, a stale positioning deck and forty Google Docs. Every new hire, agency and tool re-learned it badly. AI teammates make this worse, not better, unless the context comes first.

AI TEAMMATES	PERSON AT THE GATE	STARTS AT	NORTH STAR
Knowledge Hub, Brand Lock, every chief of staff	Head of growth confirms ICP and voice	Setup, before any agent acts	Drafts approved without edits

The play

01 **Load the Knowledge Hub.** Six shelves: Documents, ICP, Products and services, Website, Brand and Assets. Upload what you have; the site is read for you. Every agent grounds its work in these passages.

02 **Write the ICP as a definition, not a list.** Industry, size, trigger events, the job the buyer is hiring you for, and who is out of scope. Scoring, targeting and copy all read from this one definition.

03 **Lock the brand.** Brand Lock extracts your voice from what you have published. It stays a draft until a person confirms it. Add banned words, claims you can prove and claims you cannot make.

04 **Name your proof.** Customer stories, numbers and quotes, each with its source. Agents may only use claims a person has approved, which is what stops confident invention.

05 **Set up rooms.** One room per team, with its people and agents on the roster. Decisions, files and pins live in the room, and team memory carries them forward so nobody re-explains last week.

Insight

The highest-leverage hour in the whole playbook is the one spent writing down what you will not say. A banned-claims list does more for quality than any prompt.

Metrics to set

Drafts approved without edits	>60% by week 6
Knowledge gaps flagged by agents	Falling weekly
Approved proof points on file	20+

Mistakes we see

- Switching on agents first and "fixing the voice later".
- An ICP so broad that scoring cannot separate good from bad.
- Proof with no source, which the agents then repeat.

STAGE 02 · BE FOUND

Be the answer, *not just a result*

Buyers now ask ChatGPT, Gemini, Perplexity and Google's AI answers before they ever visit your site. The old playbook tracked rankings and traffic. The AI-native one tracks whether you are named, cited and recommended for the questions your buyers actually ask.

 **DRIS** organic search and AI visibility

AI TEAMMATES	PERSON AT THE GATE	STARTS AT	NORTH STAR
DRIS, with gaps handed to CREA	SEO or content lead approves fixes	L2: audits and drafts fixes	Share of tracked questions cited

The play

01 **Track questions, not keywords.** Write the 50 to 100 questions your buyers ask before they buy. DRIS samples ChatGPT, Claude, Gemini, Perplexity and Grok for each, repeatedly, because a single run is noise: answers change day to day.

02 **Measure six dimensions.** Coverage, share of voice, entity strength, recommendation position, sentiment and agent legibility. Plot yourself against competitors.

03 **Fix retrievability first.** Site audit, metadata, structure, crawl access. Approved metadata fixes go live through the site connection; any one can be withdrawn. Index rules and canonicals never earn autonomy.

04 **Then add evidence.** Pages win citations by carrying quotations, statistics and sources. DRIS sends content gaps to CREA as tasks, capped at ten a day.

05 **Refresh what is cited.** Cited pages are overwhelmingly recent. Put your top pages on a refresh cycle.

Proof

Sophoz launched a new site on 10 September 2026. Its first Perplexity citation came six days later. Within weeks, 19 pages had been cited 185 times across Perplexity and ChatGPT.

Metrics to set

Questions where you are cited	`Up month on month`
Recommendation position	`Top three`
Pages refreshed in last 12 months	`>75% of top pages`
Fixes approved within a week	`>80%`

Mistakes we see

- Reporting a single "AI rank" from one run.
- Rewriting body copy for "GEO" before fixing structure, which can cost retrieval.
- Expecting AI answers to replace search traffic by volume. They are an influence channel first.

STAGE 03 · CONTENT ENGINE

Content that *earns a subscriber*

The old content team was measured on publishing consistency, so it published consistently and generated little pipeline. The AI-native content engine is measured on what each piece produces: subscribers, leads and conversations.

 **CREA** chief of staff

 **RUPA** images

 **FOLIO** long-form

 **Alpha** video

LANDON landing pages

AI TEAMMATES	PERSON AT THE GATE	STARTS AT	NORTH STAR
CREA and its team leads	Editor approves before publish	L2: drafts, waits to publish	Subscribers per piece

The play

01 **Run a topic queue with performance signals.** Topics come from DRIS gaps, sales objections and customer questions, ranked by what past pieces generated.

02 **Let Blog Studio do the pipeline.** Research, draft, edit, review, illustrate, internal links and SEO audit, grounded in your Knowledge Hub and brand book. A fail-closed editorial check runs at publish; a person approves.

03 **Make one idea travel.** RUPA renders on-brand images (Brand Lock on by default). Alpha turns the idea into short films with a storyboard you approve and a quote in credits before rendering. FOLIO turns a theme into a PDF playbook or report.

04 **Gate the best of it.** FOLIO can file a book as a gated lead magnet once a person presses Promote. LANDON builds the landing page and form; nothing goes live until someone presses Publish.

05 **Broadcast and bridge.** Every publish goes to subscribers. Opens and clicks raise intent, and a subscriber whose intent crosses your threshold becomes a lead with the content source attached.

Proof

Emami's content studio used HOOK for angles and Alpha for films: 50 persona ideas, 15 shortlisted, 45 videos, and more than 25 million YouTube views.

Metrics to set

Subscribers generated per piece	`Rank topics by it`
Engagement per subscriber	`Not emails sent`
Content-sourced conversations	`The only revenue number`
Cost per piece	`In credits, budgeted daily`

Mistakes we see

- Publishing more because it is now cheap. Volume without distribution is noise.
- Building the demand loop before 50 leads a month; the ratios are not yet measurable.

STAGE 04 · PAID DEMAND

Every channel *accountable*

Paid teams lose most of their week to building campaigns, checking structure and refreshing creative, and too little to deciding where the money should go. AI-native paid flips that: agents build and check, people decide and approve spend.

 NIA paid performance

 HOOK paid creative strategy BETA

AI TEAMMATES	PERSON AT THE GATE	STARTS AT	NORTH STAR
NIA, one sub-agent per channel; HOOK	Performance lead approves every live action	Below live-action level	Cost per qualified pipeline

The play

01 **Brief, then build.** NIA drafts Google Search, Shopping, Display, Video and Performance Max campaigns from a brief, with fixed checks on budget, structure, copy safety, conversion tracking and duplicates. Meta and ChatGPT ads run as beta channels.

02 **Everything starts paused.** Every campaign is created paused and every live action queues for a person. Meta runs under a hard spend cap.

03 **Creative from evidence.** HOOK mines what your buyers say, finds angles, writes scripts and packs them per channel (Meta 9:16 and 1:1, TikTok Spark, Google RSA). It needs a brand pack first, is read-only on ad accounts, and overriding its quality gate takes a second person.

04 **Reallocate weekly.** NIA recommends moving budget across channels. It only recommends; the performance lead decides.

05 **Close the loop to pipeline.** Judge channels on qualified pipeline from the CRM, not on platform-reported conversions.

Insight

The fastest win in paid is rarely a new campaign. It is killing the creative nobody refreshed. Put creative fatigue on a weekly agent check and let people spend their time on the reallocation decision.

Metrics to set

Cost per qualified pipeline	Primary
Creative refreshed	Every 2 to 4 weeks
Campaigns failing a fixed check	Zero live
Time from brief to paused draft	Same day

Mistakes we see

- Letting any agent change live spend without a person.
- Optimising to platform ROAS while pipeline falls.

STAGE 05 · SOCIAL AND REPUTATION

Show up, listen, *answer*

Social and reviews are where buyers check whether you are real. The old playbook posted to a calendar and checked reviews when someone complained. The AI-native play runs a loop: plan, write, check, approve, publish, listen, measure, and answer every review with facts a person approved.

 **MOXA** organic social EARLY ACCESS

 **Echo** reputation BETA

 **Alpha** short video

AI TEAMMATES	PERSON AT THE GATE	STARTS AT	NORTH STAR
MOXA, Echo, Alpha, RUPA	Brand lead approves posts and replies	Every post waits for a named person	Engaged ICP followers

The play

01 **Plan a programme, not posts.** MOXA plans across LinkedIn, Instagram, X, YouTube, Reddit, TikTok, Facebook and Pinterest. Posts wait for a named person until MOXA has earned L5. Today Nagent's team sets MOXA up with you.

02 **Give founders and leaders a voice.** A personal programme on a leader's own profile usually outperforms the company page for B2B. The agent drafts in their voice; they approve and send.

03 **Listen everywhere buyers judge you.** Echo watches local listings, review sites, app stores, marketplaces, employer reviews and social, and opens a case with an owner, risk and due time.

04 **Answer with approved facts.** Echo drafts replies and checks claims about you in AI answers against facts a person approved. It has no send path and no approve-all: a person sends.

05 **Feed it back.** Recurring complaints become product feedback; recurring praise becomes proof on your pages.

Proof

Madmonk AI's reels, made with Alpha and posted with MOXA, reached 749K views on the top reel, with four more past 200K.

Metrics to set

Engaged followers in your ICP	`Not total followers`
Review response time	`<48 hours`
Open reputation cases past due	`Zero`

Mistakes we see

- Auto-replying to reviews. Templated replies read as indifference.
- Posting from the company page only.

STAGE 06 · CAPTURE AND SPEED-TO-LEAD

Speed to touch *is the lever*

Most inbound pipeline is lost in the gap between a buyer raising a hand and a person noticing. Research has shown for over a decade that leads contacted within the hour are several times more likely to qualify, yet the typical response still takes hours or days. Software can close that gap completely.

 **Convexa** chat, voice and video

 **CORA** CRM lead

AI TEAMMATES	PERSON AT THE GATE	STARTS AT	NORTH STAR
Convexa, CORA and CRM sub-agents	Sales lead; hot leads to a human same day	L3: every email and booking approved	Time to first touch

The play

01 **Answer on the site, now.** Convexa answers visitors in chat, voice and video, books meetings from the conversation and pushes qualified leads into the CRM. It resolves, escalates and closes the loop.

02 **One CRM for every source.** Website forms, the chatbot, calendar bookings, content pages and prospecting all land as leads with ICP fit scoring, in a CRM room led by CORA.

03 **Personal first touch in minutes.** CRM sub-agents draft the welcome, enrich the lead, score it and propose a meeting. At the start, every email and booking waits for approval; as the record builds, they earn the right to send.

04 **Brief the human at the same time.** When a lead is hot, the owner gets a short brief: who, what they read, fit score, suggested next step. Positive intent reaches a person the same working day.

05 **Use magnets as value exchange.** A useful asset (like this playbook) in return for an email, delivered instantly, with the topic tagged on the subscriber so follow-up is relevant.

Proof

Bihang PDC paired DRIS with Convexa over nine months: the first clinic named by ChatGPT and Google's AI answer in its category, and about 15 qualified consultations a month.

Metrics to set

Time to first touch, hot leads	`p50 < 5 min`
Hot leads with a human same day	`100%`
Lead to meeting rate	`Weekly trend`
Conversations resolved without escalation	`Rising`

Mistakes we see

- A daily digest that is really 40 alerts. One summary, no per-event noise.
- Routing rules nobody owns. The CRM room owns them.

STAGE 07 · OUTBOUND

Fewer sends, *more meetings*

Outbound broke when volume became cheap. Teams sent more, reply rates fell, domains burned and buyers stopped reading. AI-native outbound goes the other way: a written cohort, real research per account, every claim sourced, and a north star of qualified meetings, not emails sent.

 NORA outbound research

VANI outbound email writer BETA

AI TEAMMATES	PERSON AT THE GATE	STARTS AT	NORTH STAR
NORA, VANI	Sales lead approves sends	L2: every send waits for a person	Qualified meetings booked

The play

01 **Define the cohort in writing.** Not a purchased list: a definition NORA searches against, once or as a saved daily or weekly campaign. Existing customers and suppressed contacts are removed automatically.

02 **Research before reveal.** Accounts move through Discovered, Researched, Enriched, Ready to send and Contacted. Contact details are revealed only at a fit score of 70 out of 100 or more, which protects budget and reputation.

03 **Use real signals.** Hiring for the role you replace, adopting competing tools, recent funding. A reason to write today beats a clever opener.

04 **Separate writing from checking.** VANI writes first touch, follow-up and breakup emails; a validator checks that every claim has a source. Sends go from the rep's own mailbox. Follow up at day 5 and day 12, then stop.

05 **Route replies by intent.** Replies are classified as interested, objection, scheduling or unsubscribe. Interested reaches a person the same working day; meetings go to a board with a handoff.

Proof

At Madmonk AI, NORA produced 100 qualified leads and 10 qualified meetings in 45 days.

Metrics to set

Qualified meetings per week	North star
Positive reply rate	Over raw reply rate
Sends per meeting	Falling
Daily outreach budget used	Within cap

Mistakes we see

- Measuring sends. The metric you reward is the behaviour you get.
- Revealing contacts before qualifying accounts.
- A fourth and fifth follow-up. Two is enough.

NORA sends email. It can write templates for social networks; a person sends those.

STAGE 08 · MEETINGS AND DEMOS

Close the *demo black hole*

Teams routinely lose a large share of booked demos to no-shows and post-demo drift. The meeting happens, the recap is written days later, and the next step is "circle back". Every part of that gap is coordination, which means every part of it can be closed.

 **SERA** sales chief of staff

 **CORA** and CRM sub-agents

AI TEAMMATES	PERSON AT THE GATE	STARTS AT	NORTH STAR
SERA, CRM sub-agents	Account executive	L3: recaps and nudges approved	Next step agreed in 24 hours

The play

01 **Every booking becomes a handoff.** SERA turns a booked meeting into a handoff card for its owner. One nobody acknowledges escalates after 36 hours.

02 **Brief the day before.** A prep brief lands the day before: the account, what they read, what they asked Convexa, fit, open questions, likely objections.

03 **Catch no-shows early.** If the meeting did not happen, a recovery sequence is proposed the same day instead of discovered next week.

04 **Recap with an outcome.** CRM sub-agents draft the recap and the follow-up from the rep's own mailbox. The rep marks the outcome (won, follow up, lost, no-show) and the right sequence follows.

05 **Run internal agent meetings.** Team leads can call a meeting with the agents in the room. It ends in decisions with reasons and action items on the board; anything customer-facing waits for approval.

Insight

The recap is the product. A buyer who receives a clear recap and next step within hours experiences your company as organised, which is exactly what they are buying.

Metrics to set

Demos with next step in 24h	>60%
Show rate	Weekly trend
Handoffs unacknowledged 36h	Zero
Recap sent same day	>90%

Mistakes we see

- Recaps written from memory two days later.
- No owner for a booked meeting after the SDR hands it off.

STAGE 09 · DEALS AND PIPELINE

Pipeline you can *believe*

The old pipeline review is a weekly meeting where reps defend numbers they updated an hour before. AI-native pipeline is read from the record continuously: stage, value, age and next step, with stuck deals named before the review, not during it.

 SERA pipeline answers

 DEXA deal execution SOON

 RIVA revenue operations SOON

AI TEAMMATES	PERSON AT THE GATE	STARTS AT	NORTH STAR
SERA and CORA today; DEXA, RIVA soon	AE judges and closes	L1 to L2: reads and recommends	Stalled deals with a next step

The play

01 **Ask the pipeline, do not build reports.** SERA answers pipeline questions from the CRM: stage, count, value and age. It never invents a deal; if it is not on the record, it says so.

02 **Name the next step on every stuck deal.** SERA flags deals that stopped moving and proposes the next step. The AE decides.

03 **Give each live deal a room.** From demo held onward, a deal room card holds the buying committee, notes, documents and the next commitment.

04 **Paper the deal without the chase (coming soon).** DEXA will remember, prepare, paper and follow through: buying committee, qualification, next commitment, order forms and risks. The AE talks, judges and closes. Discounts beyond a band you set go to approval.

05 **Keep the CRM honest (coming soon).** RIVA will watch hygiene, stalled deals and forecast. By design it will never edit a deal.

Insight

The division of labour is the whole idea: the agent remembers, prepares, papers and follows through; the person talks, judges and closes. Deals are relationships, and the relationship stays with your people.

Metrics to set

Deals with a dated next step	100%
Deals untouched 14+ days	Named weekly
Stage-to-stage conversion	Per stage, monthly
Forecast vs actual	Within your band

Mistakes we see

- Letting any agent change deal values or stages unreviewed.
- Pipeline reviews spent finding problems instead of deciding on them.

STAGE 10 · CUSTOMERS AND EXPANSION

The funnel does not *end at closed-won*

Most GTM playbooks stop at the signature. For most businesses, that is where the economics start: retention, expansion and referral. The AI-native team keeps the same record and the same room after the sale, so what customers say flows straight back into the top of the funnel.

 **Convexa** customer conversations

 **Echo** reputation BETA

 **CORA** CRM

LIPI CRM email writer BETA

AI TEAMMATES	PERSON AT THE GATE	STARTS AT	NORTH STAR
Convexa, Echo, CORA, LIPI	Customer success lead	L2 for outbound to customers	Resolution rate; expansions opened

The play

01 **Resolve first.** Convexa answers customers in chat, voice and video, resolves what it can, escalates what it cannot and closes the loop. Speed of resolution is the cheapest retention lever there is.

02 **Write to customers like you know them.** LIPI writes to people already in the CRM, in your workspace's voice, using what is on their record: what they bought, what they asked, what changed.

03 **Ask for the review at the right moment.** After a resolved case or a milestone, not on a timer. Echo then tracks and helps answer what they write.

04 **Turn questions into content.** The questions customers ask are the questions prospects ask AI engines. Feed them to DRIS as tracked questions and to CREA as topics.

05 **Open expansion as a conversation.** Usage, new stakeholders and new use cases become a suggested next step for the account owner, never an automated upsell blast.

Proof

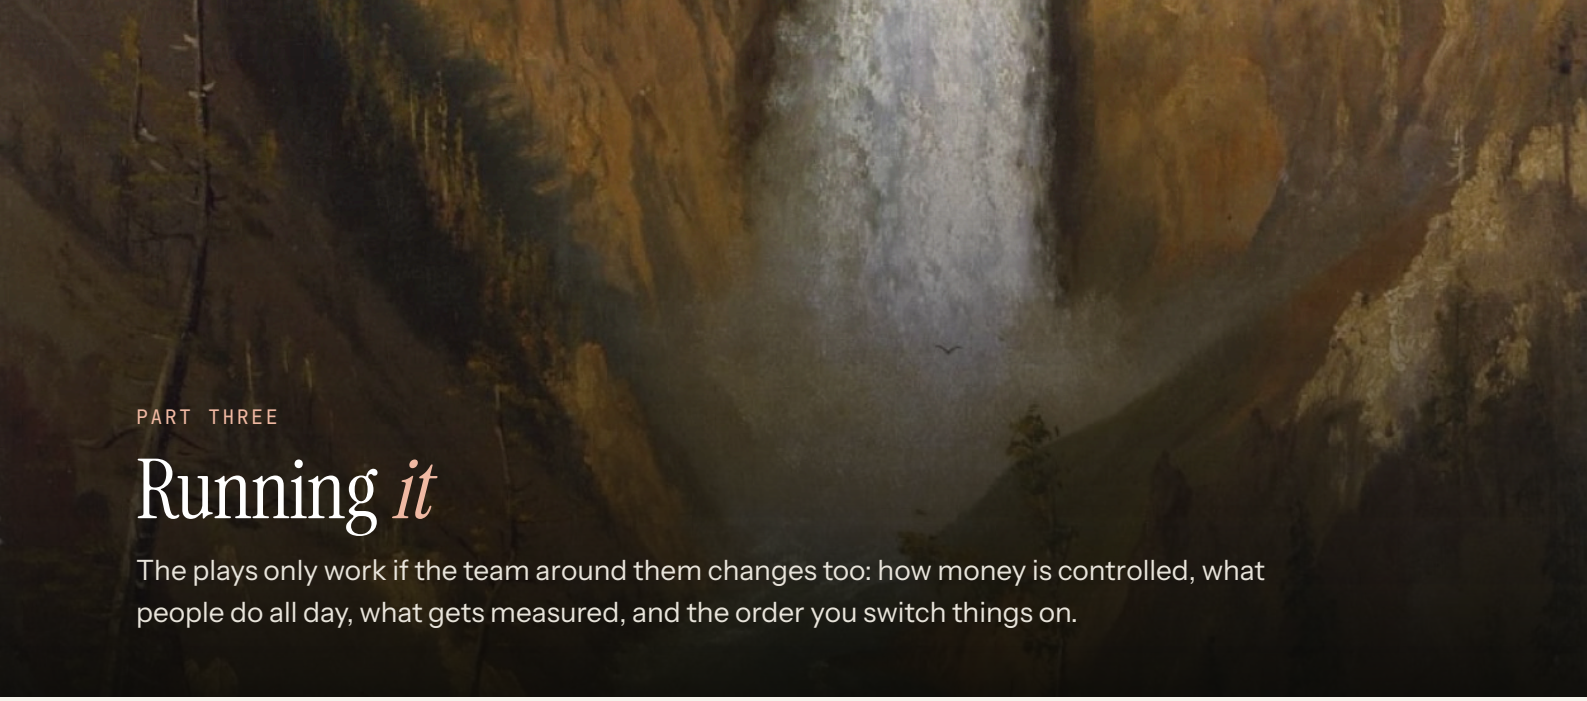
Bihang PDC holds a 4.9 rating from 51 reviews, with 50 patient questions tracked and answered across its pages and AI answers.

Metrics to set

Resolved without escalation	Rising
Time to resolution	Falling
Expansion conversations opened	Per month
Customer questions fed to content	Monthly

Mistakes we see

- Treating support as a cost centre outside GTM.
- Generic "check-in" emails that say nothing.



PART THREE

Running *it*

The plays only work if the team around them changes too: how money is controlled, what people do all day, what gets measured, and the order you switch things on.

THE WORK, IN SEVEN STEPS

Brief	One sentence in: the outcome you want.
Plan	The chief of staff breaks it down and routes it.
Run	Team leads and sub-agents do the work.
Approve	The one gate. Only the button counts.
Deliver	A finding, analysis, plan, draft or result.
Repeat	On a schedule, an event, a webhook or the thread.
Review	At one month, one quarter and one year.

WHERE PEOPLE MEET AGENTS

The room. One thread, board and record per team, with an inbox, team handbook, files and pins. Held work shows as cards to approve in the room and in one inbox.

Wherever you already are. Ask agents and approve replies from Slack with `@Nagent`. Talk to N by voice, email or the app (beta).

Approvals that expire. High and critical requests expire after 24 hours, others after 7 days, so nothing sits forgotten. Reasons go on the record.

06 · THE CONTROL PLANE

Budget, autonomy and tools, *inspectable and enforced*

Agents that act need controls that are real, not policies in a slide. The control plane is where every agent's limits live, where they are enforced before each action, and where the record of what happened can be read by anyone with the right to see it.

Budgets in credits

Every agent has a daily and monthly cap in credits; the default is **1,000 credits a day and 20,000 a month**. The cap is checked before every call, and every billable call goes on the credit ledger. When a cap is hit, the agent queues, pauses or only notifies, as you set. A breach drops its autonomy one level.

Tools, ticked one by one

Each connector action is enabled individually per agent. NIA can draft a Google Ads campaign without being able to change live budgets. Read access and write access are different ticks.

Memory with an owner

Per-agent memory, team memory in each room, and the Knowledge Hub. Extracted facts such as brand voice stay drafts until a person confirms them.

An agent, declared

Every agent is defined by ten parts you can inspect: triggers, model, memory, skills, trust, autonomy, tools, external tools, budget and surfaces. Nothing is hidden in a prompt.

The record

An append-only, exportable audit of every action, approval and autonomy change. The autonomy ledger shows why each agent is at its level.

Where it runs

Nagent cloud, a private VPC, or on premise on Enterprise, with your own model and keys. ISO/IEC 27001:2022 certified; built for DPDP and aligned with GDPR.

A governance checklist for your first agent

- Named owner and named approver
- Daily and monthly cap set, and what happens at the cap
- Tools ticked to the minimum it needs
- Starting level written down, with the reason
- Banned claims and approved proof loaded
- A review date at one month

07 · HOW THE TEAM CHANGES

Every person becomes *a manager of work*

AI-native GTM does not remove people from the funnel. It moves them from producing every artefact to directing, approving and owning the moments that need judgement. Roles change shape; the best people get more leverage, not less.

ROLE	SPENT THE WEEK ON	NOW SPENDS THE WEEK ON
Head of growth	Status meetings, chasing updates, reconciling numbers across tools	Setting outcomes per team, deciding reallocations, raising or holding autonomy, reading one record
Marketing lead	Briefing agencies, building campaigns, assembling reports	Approving plans and spend, owning the ICP and proof, deciding what not to do
Content and brand	Writing first drafts, resizing assets, scheduling posts	Editing for taste and truth, owning the voice and the banned list, approving publish
SDR or BDR	Researching accounts, writing sequences, logging activity	Approving sends, taking interested replies the same day, running first conversations
Account executive	Prep, recaps, CRM updates, chasing paperwork	Discovery, judgement and negotiation; marking outcomes so the right follow-up runs
Customer success	Answering repeat questions, writing check-ins	Escalations, renewals, expansion conversations, turning feedback into product and content
New: agent owner	Did not exist	Owns an agent's record: reviews dismissals, tunes context, proposes level changes. Often a team lead's part-time job

The head of growth's week

- **Daily:** a morning brief and a close-of-day note; approve what is waiting.
- **Weekly:** the scorecard; one reallocation decision; autonomy nominations.
- **Monthly:** review each agent's record; retire what does not move the number.

If you want help running it

Some teams want an expert in the room while they learn. Nagent's forward deployed marketers and operators can run the agent teams with you, own the plan and answer for the outcome, alongside your people.

08 · THE SCORECARD

Measure ratios, *not volume*

When work becomes cheap to produce, volume metrics stop meaning anything. A team optimising for posts published or emails sent produces a cost centre. A team optimising for ratios produces pipeline. Put this on one page and review it weekly.

STAGE	STOP REPORTING	START REPORTING	REVIEW
Foundations	Docs uploaded	Drafts approved without edits	Weekly
Be found	Keyword rank, one-off AI rank	Share of tracked questions cited, sampled repeatedly	Weekly
Content	Pieces published	Subscribers and leads per piece	Monthly
Paid	Clicks, platform ROAS	Cost per qualified pipeline	Weekly
Social	Total followers, impressions	Engaged ICP followers; review response time	Monthly
Capture	Leads captured	Time to first touch; lead to meeting rate	Daily
Outbound	Emails sent, open rate	Qualified meetings; sends per meeting	Weekly
Meetings	Demos booked	Demos with next step in 24 hours	Weekly
Deals	Pipeline value	Deals with a dated next step; stage conversion	Weekly
Customers	Tickets closed	Resolution without escalation; expansions opened	Monthly

Three numbers for the agents themselves

Approval rate

Share of an agent's cards approved without edits. The clearest signal of whether context is good enough.

Credits per outcome

Spend divided by meetings, citations or resolutions. The honest cost line, comparable across teams.

Time in queue

How long work waits for a person. If it grows, the bottleneck is approval, not production.

09 · THE 90-DAY ADOPTION PLAN

Ninety days, *in order*

The teams that succeed with agents do not switch everything on. They load context, pick the stage where their funnel leaks most, run it supervised until it earns trust, then add the stage that feeds it.

DAYS 1 TO 14

Foundations

- Load the Knowledge Hub: ICP, products, proof with sources, site, brand. Confirm the voice in Brand Lock; write the banned-claims list.
- Pick one number for the quarter and the one stage that moves it most. Write down why.
- Set up one room for that team. Name the owner, the approver and the budget for each agent.

DAYS 15 TO 30

One stage, supervised

- Switch on the first stage at a supervised level. For most B2B teams that is capture and speed-to-lead, or outbound.
- Approve every card. Edit rather than reject, so the record teaches the agent. Track approval rate daily.
- Fix context gaps the agents flag. By day 30, aim for more than half of drafts approved without edits.

DAYS 31 TO 60

Earn and extend

- Review the first agent's record at day 30. If it qualifies, raise one level. If not, fix the context, not the agent.
- Add the stage that feeds the first: if you started with capture, add search and content; if outbound, add meetings.
- Start the weekly scorecard. Retire any volume metric still being reported.

DAYS 61 TO 90

The loop

- Connect stages: customer questions to DRIS, objections to content, content intent to leads.
- Let budgets do the cost control on agents with a clean record, and spend approval time on judgement.
- Run the quarter review: credits per outcome by team, what to scale, what to stop, which levels to raise next.

> **What good looks like at day 90:** one number moving, two or three stages connected on one record, agents on the first stage at least one level above where they started, and a team that approves work in minutes rather than producing it over days.

10 · READINESS CHECK AND FAILURE MODES

Score yourself *this week*

One point for each statement that is true today. Fewer than 6: start with foundations. 6 to 9: start one stage supervised. 10 or more: run two connected stages from day one.

STATEMENT	POINT
Our ICP is written as a definition, including who is out of scope.	☐
We have at least 20 proof points, each with a source.	☐
We have a written list of claims we must never make.	☐
Every lead source lands in one CRM.	☐
We know our median time to first touch on inbound leads.	☐
We know which 50 questions our buyers ask before they buy.	☐
Each team has one number it owns, and it is not a volume metric.	☐
We can name, for each workflow, the one person who approves it.	☐
We have a monthly budget line for AI work, not just seats.	☐
Recaps go out the same day after demos.	☐
Customer questions reach our content and search plans.	☐
Someone is willing to own an agent's record for a month.	☐

Five failure modes, and the fix

The chatbot drawer

Everyone gets a private assistant; nobody shares context. **Fix:** one room per team, one record.

Autonomy on day one

An agent publishes something wrong in week one and the team never trusts agents again. **Fix:** start supervised and earn it.

The approval swamp

Five reviewers per card; queues grow; agents look slow. **Fix:** one gate per workflow, approvals that expire.

Volume theatre

Output triples, pipeline does not. **Fix:** the scorecard on page 23.

Pilot purgatory

A clever demo, no owner, no number, no budget line, quietly abandoned in month three. **Fix:** every agent gets an owner, an approver, a cap and a review date before it runs.

11 · PROOF FROM THE FIELD

What this looks like *in practice*

Four teams, four different starting stages. None of them switched on everything. Each picked the stage closest to their number and connected outward from there.

EMAMI · CONTENT AND PAID CREATIVE

25M+

YouTube views through the enterprise content studio. HOOK found the angles and Alpha made the films: 50 persona ideas, 15 shortlisted, 45 videos. The GST 2.0 film reached more than 2 million people, and eight Navratna reels drew 2.47 million views.

SOPHOZ · SEARCH AND AI VISIBILITY

6 days

From site launch on 10 September 2026 to the first Perplexity citation. CREA wrote every landing page and DRIS tracked 88 parent questions across 17 pages. 19 pages have been cited 185 times across Perplexity and ChatGPT, with 25 sign-ups in the first month.

BIHANG PDC · BE FOUND AND CAPTURE

1st

clinic named by ChatGPT and Google's AI answer in its category, after nine months of DRIS with Convexa on site. About 15 qualified consultations a month, 50 patient questions tracked, and a 4.9 rating from 51 reviews.

MADMONK AI · OUTBOUND AND SOCIAL

10

qualified meetings and 100 qualified leads from NORA in 45 days. Alpha reels posted with MOXA reached 749K views on the top reel, with four more past 200K. Convexa runs on the site and DRIS on search. A design partner since December 2024.

> The pattern in every story: one number, one starting stage, supervised first, then connected. Not "AI everywhere".

NEXT STEP

See your funnel *in the room*

Bring your number and your leakiest stage. In one session we will map your funnel onto the stages in this playbook, pick the first one to run, and show it working on your own context.

Book a session | Start a trial

cal.id/team/nagent/nagent-demo · nagent.ai/onboarding · Reply to the email this came in; MIRA reads every one.

SOURCES

1. Oldroyd, McElheran and Elkington, "The Short Life of Online Sales Leads", Harvard Business Review, March 2011.
2. SparkToro with Similarweb clickstream data, "In 2026, less than one third of Google searches still send a click", 9 June 2026.
3. Aggarwal et al., "GEO: Generative Engine Optimization", KDD 2024, arXiv:2311.09735.
4. Schulte, Bleeker and Kaufmann, arXiv:2604.07585, April 2026, on day-to-day instability of AI citations.
5. Nagent, The Answer Layer and The Outbound Thesis, nagent.ai/artefacts.
6. Customer figures from published Nagent customer stories, nagent.ai/customers.