



AI growth operations

A practical guide to designing AI-enabled growth operations with accountable commercial review.

WORKS WITH THE ASSISTANTS TEAMS ALREADY USE



EDITION 26 / EMBER & ROSE

EXECUTIVE WORKING EDITION - 08 OCT 2026

26

AT A GLANCE

A practical guide to designing AI-enabled growth operations with accountable commercial review.

Growth teams carry a heavy load of repeated work: reading market signals, building campaigns, producing variants, learning from customers and reporting results. AI can compress much of that work and widen how much a team can test. It can also flood the market with generic content, bias decisions with weak data and put brand and compliance at risk. This playbook describes an operating model for AI-enabled growth: signals, campaign operations, customer learning loops, commercial review and measurement, with people accountable for what goes out and what gets funded.

WHO IT IS FOR

Heads of growth, marketing, sales operations and revenue operations, plus product marketing, customer success and finance partners who share responsibility for commercial results.

TIME TO USE

70 minutes to read; a half-day workshop to design the first operating loop

What you will have at the end

- A map of growth tasks suited to AI support, with owners
- A weekly operating loop from signal to decision to learning
- A commercial review gate for claims, brand, compliance and spend
- A measurement plan that separates activity from incremental results

Contents

01	What AI-enabled growth operations means	03
02	Where AI fits in growth work and where it does not	04
03	The operating loop, in plain language	06
04	Roles, cadence and the commercial review gate	07
05	Risks, controls and customer trust	09
06	Measurement: activity, outcomes and incrementality	11
07	Building the first operating loop	13
08	Common pitfalls and an illustrative scenario	14
09	Questions to ask before you buy or build	16
10	A first quarter that produces evidence	17
11	Terms used in this publication	18

EXECUTIVE BRIEF

What AI-enabled growth operations means

01



The goal is a faster, better-governed loop from signal to action to learning, not a larger volume of content.

Growth operations is the system behind commercial results: how a company notices demand, decides where to act, produces and runs campaigns, learns from customers and reallocates effort. Teams have always done this with spreadsheets, meetings and judgment. AI changes the economics of the work in the middle. Reading a hundred sales call notes, grouping them into themes, drafting five angles for a campaign, building variants for each segment and summarizing last month's results can all be done in hours instead of days.

Speed on its own is a trap. A team that produces ten times as many emails that nobody wanted has made its brand worse more efficiently. The advantage goes to teams that use the saved time to learn more per cycle: more hypotheses tested, more customer evidence read, more decisions made with data and fewer based on opinion. That requires a clear operating loop and a commercial review step with named owners who are accountable for claims, spend and customer treatment.

- Signals: gather and make sense of market, competitor and customer evidence
- Decisions: choose where to focus using evidence, not volume of activity
- Production: create and adapt campaign assets within brand and compliance rules
- Delivery: run experiments with controlled targeting and spend
- Learning: feed results and customer language back into the next cycle
- Review: keep people accountable for claims, data use and commitments



PRINCIPLE

AI widens how much you can try. People decide what is worth trying, what is true enough to say, and what is fair to the customer.



This playbook is written for commercial leaders who want a practical design, not a tool list. Tools change; the loop, the controls and the measurement discipline last.

WORKING QUESTION

If your team doubled its output tomorrow, would results double, or would you just have more to review?

WHERE IT FITS

Where AI fits in growth work and where it does not

02

AI is strong at reading, summarizing, drafting and variation. It is weak at knowing your customer's real situation, your brand's limits and what is true about your product.

GROWTH TASK	FIT	REASON
Summarizing sales calls, support tickets, reviews	Strong	Large text volumes, theme extraction, human spot checks
Competitor and market scanning	Strong with verification	Faster collection; claims must be checked at source
Drafting campaign copy and variants	Strong with review	Fast first drafts; brand and claims review required
Segmenting audiences and personalizing messages	Moderate	Depends on data quality and consent; avoid creepy or unfair targeting
Lead scoring and prioritization	Moderate	Needs historical outcomes; watch for bias and drift
Forecasting pipeline	Moderate	Statistical methods often outperform language models; keep human judgment
Pricing and discount decisions	Weak	Strategic, legal and margin implications; humans decide

GROWTH TASK	FIT	REASON
Brand positioning and strategy	Weak as decider	Use as a sparring partner; leadership owns the choice
Customer commitments, contracts and regulated claims	Not suitable, unattended	Accountable people and legal review



Match the method to the question

Not every growth problem is a language problem. Questions about which segment converts best or whether a campaign caused an uplift need analytics and experimental design, not a chat assistant. Questions about what customers actually say, how competitors position themselves or how to adapt a message to five audiences suit language models well. Good teams use each method where it fits and use AI to make the analysts faster, not to replace measurement.



COMMON MISTAKE

Asking a language model to explain why sales dropped. It will produce a convincing story. Ask it to list hypotheses and the data that would test each, then test them.

WORKING QUESTION

Which of your growth tasks are language tasks, and which are measurement tasks that need analysts?



HOW IT WORKS

The operating loop, in plain language

03

A weekly loop turns scattered inputs into decisions and learning. AI speeds up each stage while people hold the decisions.

01 Sense

Collect signals: sales notes, support themes, search and social trends, competitor changes, win and loss reasons, product usage. AI clusters and summarizes; analysts verify the strongest patterns.

02 Decide

A weekly growth review chooses two or three bets, each with a hypothesis, audience, success measure and budget limit.

03 Produce

AI helps draft assets and variants from an approved message house and claims library. Humans edit and approve.

04 Run

Launch with defined audiences, holdouts where feasible, frequency caps and spend limits.

05 Learn

Compare results with hypotheses. Extract customer language and objections. Record what to repeat, change or stop.

Foundations the loop depends on

- A message house: positioning, proof points and approved claims, kept current
- A claims library: statements that can be made, with evidence and expiry dates
- A voice guide with examples of good and poor copy
- A clean customer and campaign data layer with consistent definitions of lead, opportunity and customer
- Consent and preference records that govern how customer data may be used
- A shared experiment log so results are not forgotten

Retrieval is what makes AI useful here. If the assistant can look up your claims library, voice guide and past campaign results, its drafts are grounded. Without that, it will produce generic copy and may invent features, customer logos or statistics. A good control is a rule that any factual claim in a draft must map to an entry in the claims library, or be flagged for verification.



LOOP STAGE	AI HELPS WITH	PERSON DECIDES
Sense	Clustering feedback, summarizing calls, scanning sources	Which signals matter
Decide	Drafting hypotheses, options and trade-offs	Which bets to fund
Produce	First drafts, variants, adaptations	What is accurate, on brand and compliant
Run	Setup checklists, anomaly alerts	Targeting, budgets, pauses
Learn	Result summaries, theme extraction	What the results really mean



WORKING QUESTION

Do you have a single current source of approved claims and positioning that an assistant could draw on?

OPERATING MODEL

Roles, cadence and the commercial review gate

04



Name who owns signals, assets, data and approval. A short, steady cadence beats a large, occasional review.

ROLE	RESPONSIBILITY	DECISION THEY KEEP
Head of growth	Owens the loop, priorities and results	Which bets get funded or stopped
Growth analyst	Validates signals, designs tests, measures results	What counts as evidence



ROLE	RESPONSIBILITY	DECISION THEY KEEP
Content and campaign lead	Directs production with AI support	Message quality and brand fit
Revenue operations	Owns data definitions, routing and systems	Data rules and integrations
Sales and customer success leads	Supply field evidence and test messages	Whether messages match real conversations
Legal and compliance reviewer	Reviews claims and data use for risk	Whether a claim or practice is acceptable
Finance partner	Challenges spend and return	Budget release and reallocation

The commercial review gate

Every asset or decision with external impact passes a light review before launch. The gate should be proportionate. A social post variant needs a quick check; a pricing page change, a claim about performance or a campaign that uses personal data needs more. Define tiers so reviewers are not overloaded and so risky items do not slip through.

TIER	EXAMPLES	REVIEW
Low	Social variants, subject lines with approved claims	Content lead check, spot audits
Medium	New landing pages, email sequences, sales enablement	Content lead plus claims check against library
High	Performance claims, comparisons with competitors, pricing, personal data use, regulated sectors	Named approver plus legal or compliance

Review gate checklist

- Every factual claim maps to the claims library or has a cited source
- No invented customers, quotations, statistics or testimonials
- Comparisons with competitors are accurate, current and fair
- Audience selection respects consent, preferences and fairness
- The message does not exploit sensitive circumstances or mislead
- Required disclosures and unsubscribe options are present
- The asset follows the voice guide and sounds like the company
- A named person has approved it and the record is saved

Cadence: a weekly growth review of 45 minutes, a monthly look at the experiment log and claims library, and a quarterly look at the whole system, including tools, data quality and results by channel.

WORKING QUESTION

For your highest-risk type of external content, who exactly approves it, and what do they check?

RISKS AND CONTROLS

Risks, controls and customer trust

05

The most serious growth risks are false claims, misuse of customer data and unfair treatment. Controls are mostly discipline in review and data handling.

RISK	HOW IT SHOWS UP	CONTROL
Unsupported or false claims	Draft states performance or features that are not true	Claims library; review gate; no claim without evidence
Invented proof	Fake testimonials, logos, statistics	Prohibit in instructions; reviewer checks; source for every proof point
Misuse of personal data	Personalization that crosses consent or feels intrusive	Consent records; data minimization; review by privacy owner; check current requirements
Bias and unfair targeting	Segments or scores that disadvantage groups	Review inputs and outcomes by group; human review of rules
Generic, off-brand output	Content sounds like everyone else	Voice guide with examples; edit for substance, not only polish
Overreliance on weak signals	Decisions based on a few vivid anecdotes	Analyst validation; sample sizes; hypotheses stated up front

RISK	HOW IT SHOWS UP	CONTROL
Platform and regulatory rules	Ads or messages breach rules on advertising, email or data	Check current requirements for each channel and region; keep legal informed
Confidential data leakage	Pipeline or customer data in unapproved tools	Approved tools only; role-based access

Laws and platform rules on advertising, email, telemarketing, data protection and disclosure of AI-generated content vary by place and change. This playbook does not provide legal advice. Build a habit of checking current requirements for each region and channel, and involve legal or compliance reviewers in the design of your high-tier gate.

TRUST TEST

Would a customer be comfortable if they saw exactly how we chose this message for them and what data we used? If not, change the practice before you scale it.

- Be honest about automation where customers interact directly with AI, and offer a path to a person
- Do not imply a person wrote or endorsed something if they did not
- Respect opt-outs and preference changes quickly and completely
- Keep sensitive attributes out of targeting unless there is a clear, lawful, reviewed reason
- Keep a log of inputs, outputs and approvals so you can investigate complaints

WORKING QUESTION

Which of your current personalization or targeting practices would you be least comfortable explaining to a customer?

MEASURING VALUE

Measurement: activity, outcomes and incrementality

06



Track what changed because of the work, not merely what the team produced.

Growth teams can drown in dashboards. Separate three kinds of measure. Activity measures show what the team did, such as assets produced or experiments launched. Outcome measures show market response, such as qualified leads, conversion and pipeline. Incremental measures show what would not have happened without the work, which requires holdouts, comparisons or controlled tests. AI will raise activity easily. Only the other two tell you whether it matters.



MEASURE	TYPE	NOTE
Time from signal to launched test	Operating speed	Should fall without quality loss
Experiments completed per month and share with a clear conclusion	Learning rate	More tests are useful only if they are decisive
Qualified lead rate and Outcome conversion by stage	Outcome	Needs consistent definitions
Pipeline and revenue influenced	Outcome	Attribution is imperfect; treat with care
Uplift versus holdout	Incremental	Strongest evidence of impact
Cost per qualified opportunity	Efficiency	Include tools and review time
Claim errors caught at review and after launch	Quality	Post-launch errors are the key alarm
Complaints, opt-outs and unsubscribe rate	Customer trust	Rising rates signal a problem



Run clean experiments



01 State the hypothesis

Write what you expect, for whom, and which metric will move. Decide the minimum change that would matter.

02 Control what you can

Use a holdout or A/B split. Change one main thing at a time. Fix the duration before starting.

03 Check the data

Verify tracking, sample sizes and segment definitions before reading results.

04 Decide and record

Scale, change or stop. Log the result, including failures, in the experiment log.



Illustrative example, not a benchmark: a team uses AI to produce five variants of a nurture email for each of three segments. Without a holdout, a rise in clicks may reflect seasonality or list changes. With a holdout, the team can say whether the sequence added qualified conversations. Count review time and tool cost in the efficiency measure, so that savings are real.

BEWARE FALSE PRECISION

Language models can summarize results fluently but may misstate significance or invent a cause. Have an analyst verify calculations and conclusions before decisions are made.



WORKING QUESTION

Which of your current growth metrics could rise even if the business gained nothing?



IMPLEMENTATION GUIDE

Building the first operating loop

07



Begin with one loop for one segment or product, running weekly, and improve it before adding channels.

01 Pick the scope

Choose one product, segment or region with enough volume to learn from. Name the loop owner.



02 Build the foundations

Create the claims library, message house and voice guide. Fix the worst data definition issues in the CRM.

03 Set up customer learning

Route call notes, tickets and win and loss reasons into a regular AI-assisted theme summary, validated by an analyst.

04 Define the review gate

Agree tiers, reviewers and turnaround times. Publish the checklist.



05 Run weekly

Hold the 45-minute review. Pick bets, launch tests, review results and update the log.

06 Evaluate at 90 days

Compare speed, learning rate, outcomes and quality errors with your baseline, then decide how to extend.

Customer learning loop in practice

- Collect: recorded calls where consent allows, notes, tickets, survey comments, reviews
- Protect: remove personal identifiers before analysis where the task does not need them
- Cluster: use AI to group by theme, then sample the raw text to check the groups
- Quantify: count themes by segment and stage with an analyst, not by impression
- Translate: turn findings into message tests, product feedback and sales enablement
- Close the loop: tell sales and product what was learned, and what changed



Launch readiness

- Loop owner and reviewers are named
- Claims library and voice guide exist and are current
- Approved tools and data access are set up
- Review tiers and turnaround times are agreed
- Baseline for speed, conversion and quality is captured
- Experiment log is created and shared
- Escalation route for complaints and suspected errors is clear

Train the team on both the tools and the judgment. Show examples of strong and weak AI drafts, practice editing for substance, and rehearse how to challenge a convincing but unsupported summary.

WORKING QUESTION

Which single product or segment is the best place to prove the loop, and who will own it?

PITFALLS

Common pitfalls and an illustrative scenario

08

Most failures come from volume without learning, claims without evidence and measurement without control.

- Publishing more because it is cheap, until the brand sounds generic
- Allowing assistants to cite results, customers or features that do not exist
- Personalizing with data customers did not expect to be used
- Reading dashboards without holdouts and crediting AI for seasonal gains
- Skipping the review gate when deadlines press
- Treating call summaries as facts without checking the underlying notes

- Letting the claims library go stale

ILLUSTRATIVE SCENARIO, HYPOTHETICAL

A team launches an AI-drafted landing page that says the product 'integrates with all major CRMs.' It supports two. Sales spends a quarter correcting expectations. The claims library now lists exact integrations, and the gate requires every capability claim to match it.

SYMPTOM	LIKELY CAUSE	RESPONSE
Rising output, flat pipeline	Activity without learning or targeting	Cap volume, strengthen hypotheses, use holdouts
Sales distrusts marketing content	Claims drift from reality	Involve sales in review; refresh claims library
Unsubscribes and complaints rise	Over-frequency or intrusive personalization	Cap frequency; revisit data use
Insights feel repetitive	Summaries based on the same sources	Add new evidence streams; interview customers
Reviewers overloaded	No tiering	Introduce tiers; automate low-risk checks

A subtler pitfall is organizational. If growth, sales, product and customer success each run their own AI experiments, customers meet inconsistent messages and the company learns the same lesson four times. A shared claims library, a shared experiment log and a monthly cross-team review keep the company coherent.

Recovery follows the same pattern as in other operations: pause, find the cause, fix the foundation and tell people what changed. Teams that handle errors openly keep the trust of sales, legal and customers.

WORKING QUESTION

What would be the first visible sign that your team was producing more but learning less?

BUYING AND SETUP

Questions to ask before you buy or build

09

Ask about data, claims control, integration and evidence, and pay as much attention to the review workflow as to the generation features.

Questions for any growth technology or agency partner

- Can the system draw on our claims library and flag unsupported statements?
- How is customer and prospect data stored, used and protected, and is it used to train shared models?
- Can we apply consent and suppression rules, and are they enforced across channels?
- What approval workflow, roles and audit trail are built in?
- How does it integrate with our CRM, marketing platform, analytics and call recording tools?
- How does it support holdouts and incrementality, and who verifies the measurement?
- What evidence of results can you show that is relevant to a company like ours, and can we run a pilot on our own data?
- What are the full costs, including seats, usage, integration, content review time and training?
- How do we export our data, prompts and assets if we leave?

OPTION	BEST WHEN	WATCH FOR
Add AI features to existing platforms	Your stack is already well integrated	Feature depth and data terms
Specialist point tools	A clear gap such as call analysis	Fragmented data and governance
Custom assistant on your content	You have strong foundations and specific needs	Build and maintenance effort
External agency with AI-enabled workflow	You need capacity quickly	Control of claims, data and quality

Be skeptical of demos that show volume. Ask to see the review and measurement features, and run a pilot with a hypothesis, a holdout and a pre-agreed pass mark. Make sure someone on your side owns the claims library and the experiment log, whatever you buy.

If you want an outside view of where AI would help your growth operations, Halden & Company offers free assessments at halden.company/assessments, and you can contact the team at halden.company/contact. The approach stays the same: build the loop, control the claims, measure what changed and keep people accountable.

**WORKING QUESTION**

If a vendor's AI wrote a false claim in your name, which part of your process would have caught it?

30 / 60 / 90

A first quarter that produces evidence

10



Prove one reviewed, measured growth loop on a single segment, then extend with evidence.

First 30 days: ground it

- Name the loop owner and choose one product or segment
- Create the claims library, message house and voice guide
- Agree review tiers, reviewers and turnaround times
- Capture a baseline for speed, conversion and quality errors

**Days 31 to 60: test it**

- Run the weekly growth review and launch the first tests with holdouts
- Start the AI-assisted customer theme summary, validated by an analyst
- Track claim errors caught at review and after launch
- Start the shared experiment log



Days 61 to 90: decide

- Compare results with the baseline and the holdouts
- Refresh the claims library and tighten the review gate
- Decide whether to extend to another segment or channel
- Present learning and next bets to commercial leadership and finance

GLOSSARY

Terms used in this publication

11

TERM	MEANING
Growth operations	The systems, data and routines that connect market signals, campaigns, sales handoff and learning.
Message house	A structured statement of positioning, key messages and proof points.
Claims library	A maintained list of statements the company can make, each with evidence and an owner.
Holdout	A group deliberately not exposed to a campaign, used to measure what the campaign added.
Incrementality	The additional result caused by an action, beyond what would have happened anyway.
Attribution	Methods for assigning credit for outcomes to channels or touchpoints, which are always approximate.
Lead scoring	Ranking leads by estimated likelihood to convert, based on past data.
Voice of customer	Customer language and feedback gathered from calls, tickets, reviews and surveys.
Consent	A person's permission for specified use of their data, which must be recorded and respected.
Revenue operations	The function that aligns data, process and systems across marketing, sales and customer success.
Experiment log	A shared record of hypotheses, tests and results, including failures.

TERM	MEANING
Personalization	Tailoring messages or offers to a person or segment, within consent and fairness limits.



A note on use

This playbook is general guidance and not legal or compliance advice. Rules on advertising, email, data protection and AI disclosure vary by region and change; check current requirements and involve legal or compliance reviews for high-risk content.

TAKE THE NEXT STEP

Take the free Halden assessment at halden.company/assessments, or talk to us at halden.company/contact about turning this into a plan for your organization.

