

MARKETING OPERATIONS & REVOPS

AI Agents for *Marketing* *Operations.*

The agent stack that turns the ops function into the leverage point of the entire GTM org.

WRITTEN BY

Erik R. Miller

SENIOR B2B MARKETING EXECUTIVE · FRACTIONAL CMO

5

AGENTS

1

WORKFLOW

23

PAGES

AI AGENTS FOR MARKETING OPERATIONS
THE OPERATOR PLAYBOOK · VOL. 05

FIRST EDITION, 2026

© 2026 ERIK R. MILLER · E.R.M. ADVISORY

NEW YORK CITY

ERIKRMILLER.COM

THE OPERATOR NEWSLETTER · THEOPERATOR-ERM.BEEHIV.COM

This Field Note is one volume in the Operator Playbook series.

Adapted from the flagship essay AI Agents for Marketing Teams, originally published on erikrmiller.com.

Permission granted to forward, print, and share this document in full. Excerpts welcomed with attribution.

Marketing operations is the leverage point.

Every other function in a marketing org runs on what marketing operations produces. The data. The reporting. The workflows. The integrations. The knowledge base. When ops works, every other function gets faster. When ops fails, every other function lies to itself.

Five agents change the shape of the function. Analytics that produces narrative weekly reports, not just dashboards. Customer journey mapping that updates itself off real touchpoint data. A knowledge base that answers every internal question with the actual canonical answer. CRM hygiene that catches data decay before it kills attribution. Workflow automation that builds and audits the wiring between systems.

This playbook covers the five operations and intelligence agents, plus the weekly performance intelligence workflow. If you run marketing ops or RevOps, the agent stack in this Field Note is your starting point.

WHO THIS IS FOR

For Heads of Marketing Operations, RevOps leaders, and the operators responsible for data quality, attribution, and the marketing-tech stack.

CONTENTS

Inside this *Field Note*.

5 agents and 1 hero workflow.

THE AGENT ROSTER

- AGENT 01** The Analytics Reporting Agent
 - AGENT 02** The Customer Journey Agent
 - AGENT 03** The Knowledge Base Agent
 - AGENT 04** The Workflow Automation Agent
 - AGENT 05** The CRM Hygiene Agent
-

HERO WORKFLOW

- WORKFLOW** Weekly performance intelligence
-

OPERATIONS & EXECUTION

- Pre-flight checklist
 - The 30-60-90 Roadmap
 - Pitfalls & remediation
-

WHERE THIS FITS

- Where this fits.
-

The *Marketing Operations.* Agents.

AGENT 01

The Analytics & Reporting Agent

PURPOSE. Produces weekly, monthly, and quarterly performance reports tied to revenue metrics. Identifies anomalies, explains them, and recommends specific next actions. Replaces the dashboard-walking ritual with structured intelligence.

INPUTS. CRM pipeline data, web analytics, paid platforms, content performance, ABM dashboard, sales activity data, the strategic narrative for context.

OUTPUTS. Standardized W/M/Q reports (under 600 words weekly; under 1,500 monthly; under 3,000 quarterly), anomaly explanations, recommended next actions, executive-friendly summary block at top.

MEMORY. Last 12 reports for trend analysis, recurring themes, prior recommendations and whether they were acted on.

HUMAN OVERSIGHT. CMO reviews weekly reports. Anomalies (10%+ deviation from baseline) trigger root-cause analysis before publishing. Monthly accuracy audit on whether recommended actions correlated with subsequent improvement.

COMMON MISTAKES. Defaulting to vanity metrics when revenue data is incomplete. Misattributing outcomes. Failing to recommend specific next actions, leaving leaders to interpret on their own.

HOW TO DEPLOY

The Analytics & Reporting Agent is the intelligence layer. It does not run a dashboard. It writes the weekly narrative that turns data into a decision the marketing team can act on. Deploy it before any optimization agent goes live.

- 01** Build the canonical metrics dictionary. Every metric the agent reports against must have a single definition the entire org agrees on. Marketing-sourced pipeline, marketing-influenced revenue, MQL definition (if you still use one), ICP-fit pipeline. Disambiguation is the first job.
- 02** Connect the agent to every data source it needs (CRM, marketing automation, ad platforms, product analytics, attribution model). Disconnected agents produce confident-sounding narratives over incomplete data. That is the failure mode to prevent.
- 03** Produce the first three weekly narrative reports in shadow mode. The Head of Marketing Ops and the CMO read the reports and flag every claim that is wrong, ambiguous, or unsourced. The agent retrains off the flags.
- 04** Promote to live state. The agent now publishes the weekly performance narrative every Monday morning. The CMO uses it as the starting point for the leadership Monday review.
- 05** Long-run a monthly accuracy audit. Where the agent's narrative diverged from eventual outcomes, the divergence gets documented. The agent gets retrained against the audit.

SUCCESS METRICS · WHAT TO MEASURE

Narrative accuracy	Percent of weekly narrative claims that hold up against the next month's actuals. Target: > 90%.
Decision velocity	Time from data event to decision action, vs. pre-agent baseline. Target: 50% reduction.
Adoption depth	Percent of marketing leadership meetings that reference the current week's narrative as primary input. Target: 100%.
Metric-dispute frequency	Number of cross-functional disputes about what a metric means per quarter. Target: declining quarter over quarter.

The Customer Journey Mapping Agent

PURPOSE. Continuously maps the customer journey from first signal through closed-won and into expansion. Identifies friction points, drop-off cohorts, and the high-leverage moments where AI agents can amplify human handoff. Operates as the connective intelligence across Marketing, SDR, Sales, and CS.

INPUTS. CRM stage data; web behavior; product-usage signals; sales activity; support tickets; renewal data; customer interview themes.

OUTPUTS. Quarterly customer-journey map with quantified friction points, monthly cohort drop-off briefs, named high-leverage moments where intervention has the highest expected pipeline impact.

MEMORY. Historical journey maps with intervention-to-outcome correlation; named friction points and how they were resolved.

HUMAN OVERSIGHT. RevOps and CS leadership review the quarterly map. Specific intervention proposals require functional sign-off.

COMMON MISTAKES. Confusing the funnel with the journey. Mapping the buyer's path instead of the buyer's experience. Failing to integrate post-sale data, which leaves expansion blind spots.

HOW TO DEPLOY

The Customer Journey Mapping Agent maintains the buyer's journey as a live, behaviorally validated model rather than a static slide. It is downstream of the Buyer Persona Agent and feeds the Editorial Planning Agent and the Email Nurture Sequence Agent.

- 01** Define the journey stages explicitly. Awareness, evaluation, comparison, decision, post-purchase. Stage transitions must have observable behavioral markers, not opinions. The agent works against the defined stages.
- 02** Connect the agent to product analytics, content-engagement data, CRM stage transitions, and the Buyer Persona Agent's briefs. Without behavioral data, the journey map becomes a fantasy.
- 03** Produce the first per-persona journey map in shadow mode. Validate the maps with the Sales team. Their day-to-day reality is the closest available ground truth.
- 04** Promote to live state. The agent now refreshes the journey maps quarterly and produces monthly journey-friction reports flagging stage-transition drop-off changes.
- 05** Long-run on a quarterly recalibration. Buyer behavior changes. The map must change with it.

SUCCESS METRICS · WHAT TO MEASURE

Stage-transition accuracy

Agreement between the agent's observed transitions and sales-reported reality. Target: > 80%.

Friction-flag actionability

Percent of agent-flagged friction points that result in a corrective action within 60 days. Target: > 60%.

Cross-functional adoption

Percent of content, demand, and sales planning that references the current journey map. Target: > 75%.

Map freshness

Days since last per-persona refresh. Target: < 100 days.

The Internal Marketing Knowledge Base Agent

PURPOSE. Maintains the marketing organization's institutional memory in queryable form: prior campaigns and outcomes, decisions and reasoning, voice exemplars, persona briefs, the running record of "what we tried and what we learned." The agent that turns marketing operations into a learning organization.

INPUTS. Campaign retros, decision memos, customer interviews, sales notes, post-mortems, the editorial archive, the asset library, the case-study library.

OUTPUTS. On-demand answers to internal questions ("did we run a campaign like this in 2024?"), quarterly "what we learned" digest, decision-precedent briefs for new campaigns.

MEMORY. Tagged and indexed campaign archive; decisions with reasoning; outcomes with lessons; people-to-context attribution for institutional knowledge that lives in heads.

HUMAN OVERSIGHT. Marketing ops owns the indexing and tagging discipline. Quarterly memory audit.

COMMON MISTAKES. Letting the index decay because no one owns it. Surfacing answers without precedent reasoning, leaving teams without the why. Treating the knowledge base as storage rather than retrieval.

HOW TO DEPLOY

The Internal Marketing Knowledge Base Agent is the institutional memory of the marketing function. Every other agent draws from it. Without it, agents repeat the same questions to humans, lose context across handoffs, and produce inconsistent output across the team.

- 01** Inventory every artifact that defines how marketing runs. Brand guidelines, ICP, personas, positioning, voice document, campaign templates, attribution model, definitions. This is the agent's training corpus.
- 02** Establish the canonical-source convention. Each artifact has exactly one canonical source. If the brand voice document exists in three places, two of them get deleted. The agent works only off canonical sources.
- 03** Run the agent in question-answering mode for two weeks. Every team member sends questions they would normally ask a colleague. The agent answers from the corpus. Wrong answers reveal where the corpus is missing or contradictory.
- 04** Promote the agent to be the default-first source. Slack integration, internal search, in-tool widgets. Everywhere a marketer would otherwise ask a colleague, the agent is the first stop.
- 05** Long-run a weekly corpus-hygiene review. The Head of Marketing Ops owns this. Stale documents get retired. New documents get canonicalized. The agent only stays useful if the corpus stays clean.

SUCCESS METRICS · WHAT TO MEASURE

Question-answer accuracy	Percent of agent answers verified as correct by the relevant subject-matter expert. Target: > 90%.
Time-to-answer	Median time from question asked to answer received vs. asking a colleague. Target: < 10 seconds. Colleague baseline is typically hours.
Source coverage	Percent of canonical sources fully indexed and queryable. Target: 100%.
Replacement adoption	Percent of questions that previously went to colleagues now going to the agent. Target: > 70% within 90 days.

The Workflow Automation Agent

PURPOSE. Designs, maintains, and audits marketing automation flows — lead routing, scoring, lifecycle staging, list management, attribution logic. Works at the intersection of MAP, CRM, and the rest of the agent system. Less glamorous than content agents; far more leveraged.

INPUTS. Current workflow specs, conversion data, lead-routing rules, CRM stage definitions, ICP grades, exit conditions.

OUTPUTS. Workflow specs (visualized), trigger and exit definitions, change-request documents with risk analysis, monthly automation audit report.

MEMORY. Historical workflow performance, change history with outcomes, current automation map.

HUMAN OVERSIGHT. Marketing operations approves any workflow change. Automation changes follow a documented change-control discipline.

COMMON MISTAKES. Over-engineering workflows so that the failure surface grows faster than the value. Letting workflow changes accumulate without documentation. Confusing complexity with sophistication.

HOW TO DEPLOY

The Workflow Automation Agent builds and audits the wiring between the systems marketing runs on. It is the agent that prevents the marketing stack from becoming an unmaintained tangle of brittle integrations. Deploy it after the Knowledge Base agent is operational.

- 01** Inventory every workflow currently running. Marketing-automation flows, ad-platform integrations, CRM triggers, attribution pipelines. Most of them will be undocumented. The agent's first job is the inventory.
- 02** Categorize each workflow by criticality. Critical, important, nice-to-have, broken. The agent's next job is to flag the broken ones and recommend retirement or rebuild.
- 03** Run the agent in audit mode for the first 30 days. It produces a weekly workflow-health report. Marketing Ops fixes what the agent flags. Trust accumulates over the fix cycle.
- 04** Promote the agent to build mode for new low-risk workflows. The agent proposes the workflow design, builds it in a sandbox, runs tests, and submits for human approval before production deploy.
- 05** Long-run a quarterly stack-audit cadence. The agent reports on workflow drift, broken integrations, and recommended consolidation. Marketing Ops acts on the report.

SUCCESS METRICS · WHAT TO MEASURE

Workflow-failure detection lead time

Median time the agent flags a workflow failure before it produces downstream impact. Target: < 4 hours.

Build-cycle time

Median time from new-workflow request to production deploy with the agent vs. without. Target: 50% reduction.

Stack health score

Composite score across workflow uptime, documentation completeness, and integration discipline. Target: monthly positive trend.

Broken-workflow remediation rate

Percent of agent-flagged broken workflows fixed within 14 days. Target: > 90%.

The CRM Hygiene & Data Enrichment Agent

PURPOSE. Maintains CRM data quality at scale: deduplication, enrichment, role-tagging, ICP-grade tagging, account-hierarchy correctness, lifecycle-stage hygiene. The unsexy agent that quietly powers every other agent in the system. Garbage data poisons the entire operating system.

INPUTS. CRM exports, enrichment-vendor data, ICP grades, role taxonomies, the buying-committee map, the workflow automation specs.

OUTPUTS. Daily duplicate report, weekly enrichment report, monthly data-quality scorecard, change-request packets for ops review.

MEMORY. Data-quality trend lines, enrichment-source quality history, recurring data hygiene failure modes.

HUMAN OVERSIGHT. Marketing operations owns every change to the CRM. Bulk changes are reviewed in batch before commit.

COMMON MISTAKES. Letting enrichment vendors degrade silently. Bulk-applying changes without review. Treating CRM hygiene as a one-time cleanup rather than an ongoing discipline.

HOW TO DEPLOY

The CRM Hygiene & Data Enrichment Agent is the unglamorous foundation of the entire marketing intelligence stack. Without it, every other agent makes decisions on stale or contradictory data. Deploy it early. Run it forever.

- 01** Audit the CRM. Most CRMs have a duplicate-account rate above 8%, a missing-field rate above 30% on critical fields, and an out-of-date contact rate above 50%. The audit is the starting point.
- 02** Define the hygiene rules. Required fields. Duplicate-detection criteria. Enrichment sources and trust ranking. Decay rules. The agent works against the defined rules.
- 03** Run the agent in audit-only mode for 14 days. It produces a daily hygiene report. Marketing Ops reviews and approves the proposed merges, enrichments, and field-completions.
- 04** Promote to bounded autonomy. Low-risk hygiene operations (field-completion from trusted sources, deduplication of high-confidence matches) execute autonomously. High-risk operations (account merges, contact reassignment) require human approval.
- 05** Quarterly, retrain the agent on its own audit log. Where it made errors, the rules get refined. Where it caught patterns humans were missing, the patterns get formalized.

SUCCESS METRICS · WHAT TO MEASURE

Duplicate-account rate	Percent of active accounts with a duplicate in the system. Target: < 2%.
Critical-field completion	Percent of accounts with all required fields populated. Target: > 95%.
Data-decay rate	Percent of contacts identified as stale (no engagement, no enrichment match) per quarter. Target: < 10%, with quarterly remediation.
Downstream-agent error reduction	Percent reduction in agent failures attributable to CRM data quality vs. pre-agent baseline. Target: > 50%.

The Hero *Workflow.*

HERO WORKFLOW

Weekly performance intelligence.

How the weekly narrative turns data into a Monday decision.

The weekly performance workflow produces the Monday-morning narrative that the CMO and the leadership team open before the leadership review. Not a dashboard. Not a metric dump. A short narrative answering three questions: what happened last week, what does it mean, and what should we do this week.

Five agents, one human checkpoint, the weekly cadence the CMO would otherwise spend a half-day producing manually. The compounding is in what the CMO does with the freed half-day, not in the time savings themselves.

AGENTS INVOLVED

- Analytics & Reporting Agent (synthesizes the narrative)
- Customer Journey Mapping Agent (provides stage-transition context)
- Internal Knowledge Base Agent (resolves metric-definition disputes)
- Workflow Automation Agent (flags any data-pipeline failures)
- CRM Hygiene & Data Enrichment Agent (validates underlying data quality)

HUMAN CHECKPOINTS

- *Head of Marketing Ops reviews the narrative draft Sunday evening*
- *CMO reads the narrative Monday morning before leadership review*

THE WORKFLOW, STEP BY STEP

- 01** Data pull. Every Sunday 06:00 ET, the Analytics & Reporting Agent pulls metrics from CRM, marketing automation, ad platforms, and the attribution model. The CRM Hygiene Agent validates that the underlying data passes quality thresholds before the pull is used.

- 02** Variance analysis. The agent compares last-week metrics to the prior 4-week average, identifies statistically meaningful moves, and labels each move as expected or unexpected.

- 03** Narrative drafting. The agent writes a 600-800 word narrative covering pipeline performance, channel performance, ABM motion status, content velocity, and anomalies worth attention.

- 04** Cross-agent context. The Customer Journey Mapping Agent provides stage-transition context for any anomalies. The Knowledge Base Agent answers definitional questions to prevent metric ambiguity.

- 05** Human review. The Head of Marketing Ops reviews the narrative Sunday evening and signs off or sends back for revision.

- 06** Publication. The narrative is sent to the CMO Monday at 07:00 ET, ready for the 08:30 leadership review.

- 07** Audit logging. Every narrative is logged. End of quarter, the team reviews the narrative log to identify where the agent saw something the team missed and where the team caught something the agent missed. The agent retrains off the audit.

SUCCESS METRICS

- Narrative on-time rate. Target: > 95% delivered Monday 07:00 ET.

- Claim accuracy. Target: > 90% of narrative claims hold up against next-month actuals.

- Decision velocity. Target: 50% reduction in time from data event to decision action vs. pre-agent baseline.

- Leadership adoption. Target: 100% of Monday leadership reviews reference the current narrative as the starting input.

COMMON FAILURE MODES

- Narratives that read like dashboards. Fix: the agent has a strict 600-800 word limit and must answer three questions, not list metrics. Discipline is enforced at the prompt level.
- Confident narratives over bad data. Fix: the CRM Hygiene Agent must validate before the pull is used. If data quality fails, the narrative is delayed or annotated with the data gap.
- Narrative becoming wallpaper. Fix: quarterly retrospectives on the narrative log keep the agent learning. Without retrospectives, the agent drifts toward generic.

Pre-flight *checklist.*

Marketing operations is the substrate layer. Every other function compounds on what marketing ops produces. Resolve every item below before activating the ops agent stack.

- The metrics dictionary exists, is canonical, and is owned by a single role (Head of Marketing Ops or RevOps). Every metric in it has one definition the entire org agrees on.

- The attribution model is documented, the methodology written, and leadership has signed off. Ambiguous attribution is the source of every cross-functional revenue dispute.

- CRM data quality has been audited. Duplicate rate, field-completion rate, decay rate are all measured. The CRM Hygiene & Data Enrichment Agent will be deployed against the baseline established here.

- Every system in the marketing stack is inventoried with its purpose, owner, integration points, and dependencies. Most marketing stacks have undocumented integrations. The Workflow Automation Agent will discover them. The audit gets ahead of the discovery.

- The internal marketing knowledge corpus is identified. Brand guidelines, ICP, personas, voice, attribution model, campaign templates, definitions. These become the Knowledge Base Agent's training corpus.

- A canonical-source convention is established. Each artifact has exactly one canonical source. Conflicting versions are deleted.

- Leadership reporting cadence and format is defined. The Analytics & Reporting Agent will produce the weekly narrative. The format and audience must be set before the agent goes live.

The 30-60-90 *Roadmap.*

The marketing-ops rollout sequence. Operations is the foundation layer. If the ops agents fail, every downstream agent fails with them.

DAYS 1-30

Stand up the Foundation. Build the metrics dictionary. Document the attribution model. Audit the CRM. Inventory the stack. Identify and canonicalize the knowledge corpus. Define the reporting cadence and format. Deploy the Internal Marketing Knowledge Base Agent in question-answering mode against the canonicalized corpus. The team uses it daily. Wrong answers reveal corpus gaps.

DAYS 31-60

Deploy the Analytics & Reporting Agent and the CRM Hygiene & Data Enrichment Agent. The Analytics & Reporting Agent runs in shadow mode, producing the weekly narrative for review by the CMO and the Head of Ops. Adjustments to the agent reflect adjustments to the metrics dictionary. The CRM Hygiene & Data Enrichment Agent runs in audit-only mode, producing daily hygiene reports. Marketing Ops fixes what the agent flags.

DAYS 61-90

Promote the Analytics & Reporting Agent to live state. The weekly narrative becomes the canonical leadership input. Promote the CRM Hygiene & Data Enrichment Agent to bounded autonomy on low-risk hygiene operations. Deploy the Customer Journey Mapping Agent and the Workflow Automation Agent. The Customer Journey Mapping Agent produces per-persona journey maps. The Workflow Automation Agent inventories and audits the existing workflow stack.

DAY 91+

Promote the Workflow Automation Agent to build mode for new low-risk workflows. Long-run the governance discipline: weekly hygiene reports, monthly accuracy audits on the Analytics & Reporting Agent, quarterly journey-map recalibration.

Pitfalls & *remediation.*

PITFALL 01

Deploying intelligence agents on top of bad data.

The Analytics & Reporting Agent produces confident-sounding narratives over whatever data it has. If the CRM is dirty and attribution is contested, the agent produces confident, wrong narratives. The CRM Hygiene & Data Enrichment Agent must be operational before the Analytics & Reporting Agent goes to live state.

PITFALL 02

Skipping the canonical-source convention.

The Knowledge Base Agent fails the moment the corpus has contradictory versions of the same artifact. Three voice documents, two ICPs, four attribution definitions. The agent inherits the chaos. Canonicalize first.

PITFALL 03

Treating ops as cost center, not leverage point.

The marketing-ops function is where every other function compounds. Treating it as a cost center starves it of the investment needed to support the agent stack. The agent stack does not survive in a starved ops function.

WHERE THIS FITS

Where this fits.

The five agents and one workflow in this playbook are the marketing-operations slice of a larger system. Ops is the substrate. It does not produce demand. It makes demand possible.

For the full architecture across strategy, content, demand, ABM, governance, and the 90-day rollout, read the flagship: **AI Agents for Marketing Teams**. Free, no email gate.

— Erik R. Miller

Erik R. Miller

SENIOR B2B MARKETING EXECUTIVE · FRACTIONAL CMO

Web: erikmiller.com

Newsletter: The Operator, theoperator-erm.beehiiv.com

LinkedIn: linkedin.com/in/erikrm

Facebook: facebook.com/ERMAdvisory

Instagram: instagram.com/ermadvisory

Advisory: erikmiller.com/#contact

Marketing operations is the function where every other *function* compounds, or fails to.

This Field Note is one volume in the Operator Playbook series. The full system has 25 agents, four hero workflows, a governance model, and a 90-day rollout.

Six other volumes cover Demand, Content, ABM, Operations, ICP, and Sales Enablement. The flagship pulls all of it together.

NEXT STEP

Get the *full playbook*.

- FLAGSHIP** Read the full 65-page playbook
- SERIES** erikmiller.com/resources
- NEWSLETTER** theoperator-erm.beehiiv.com
- FACEBOOK** facebook.com/ERMAdvisory
- INSTAGRAM** instagram.com/ermadvisory
- ADVISE** erikmiller.com/#contact