

CONTENT MARKETING

AI Agents for *Content Marketing.*

Seven agents and one workflow that take content velocity from human-throughput to operating-system speed.

WRITTEN BY

Erik R. Miller

SENIOR B2B MARKETING EXECUTIVE · FRACTIONAL CMO

7

AGENTS

1

WORKFLOW

28

PAGES

AI AGENTS FOR CONTENT MARKETING
THE OPERATOR PLAYBOOK · VOL. 03

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.

Content velocity broke human throughput. This is the rebuild.

Content is the function where the gap between what leadership expects and what humans can produce got widest, fastest. The teams winning AI search, ranking in AEO answers, and showing up in the LinkedIn feeds that buyers actually read are not the teams writing harder. They are the teams running a content operating system.

A content agent stack does three things humans cannot. It holds brand voice consistently across every output. It produces at fifteen to twenty pieces a week without quality collapse. It learns from what actually ranked, converted, or got shared, and adjusts what it produces next.

This playbook covers the seven content and brand agents, plus the cornerstone content workflow that runs across all of them. Read it as a complete map of how content marketing should be staffed in 2026.

WHO THIS IS FOR

For VPs of Content, content marketing managers, SEO leads, and the operators responsible for brand voice and editorial throughput.

CONTENTS

Inside this *Field Note*.

7 agents and 1 hero workflow.

THE AGENT ROSTER

- AGENT 01** The Brand Voice Agent
 - AGENT 02** The Content Strategy Agent
 - AGENT 03** The Editorial Planning Agent
 - AGENT 04** The SEO Agent
 - AGENT 05** The AEO Agent
 - AGENT 06** The LinkedIn Social Agent
 - AGENT 07** The Ad Copy Agent
-

HERO WORKFLOW

- WORKFLOW** Cornerstone content production
-

OPERATIONS & EXECUTION

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

WHERE THIS FITS

- Where this fits.
-

The *Content Marketing*. Agents.

AGENT 01

The Brand Voice Governance Agent

PURPOSE. Reviews any AI-generated or human-written content for brand voice consistency before it ships. Flags violations, scores fit, and recommends ship/revise/rebuild. Does not rewrite. The single most important agent for protecting voice quality across a scaled content motion.

INPUTS. The Brand Voice Document (with 30+ exemplars), the draft submitted for review, channel and audience context, the persona the piece is written for.

OUTPUTS. Voice fit score (1-10), line-by-line flags with rule references, summary recommendation (ship / revise / rebuild), specific examples from the Brand Voice Document for the writer to compare against.

MEMORY. All previously approved exemplars (positive examples), all previously rejected drafts (negative examples), recent voice updates and version notes.

HUMAN OVERSIGHT. Score 7 or below requires human writer review and resubmit. Score 8+ ships with spot-check review. Monthly recalibration: editor reviews 10 random flags to validate strictness.

COMMON MISTAKES. Becoming too lenient and letting voice drift through. Becoming too strict and blocking legitimate creative variance. Both fail modes are addressed by monthly recalibration against fresh editor judgment.

HOW TO DEPLOY

The Brand Voice Governance Agent is the quality floor. It is not a creative agent. It rejects, it edits, it enforces. Deploy it before any content-producing agent goes live, because content velocity without voice governance is a brand-dilution machine.

- 01** Build the Foundation voice document. This is not a paragraph of adjectives. It is the brand vocabulary, the banned phrases, the tonal range, the cadence rules, and a corpus of fifteen to twenty exemplar pieces of written work that demonstrate the voice in practice.
- 02** Run the agent against archive content first. Feed it ten previously published pieces and ask for voice-consistency scoring. Investigate the lowest-scoring pieces. They will reveal where the voice document is unclear or where the existing content actually does not match the brand.
- 03** Connect the agent to every content-producing agent as a mandatory review step. No human-written or AI-written content publishes without a voice-consistency pass.
- 04** Calibrate the rejection threshold for the first 60 days. Too high and the agent becomes a bottleneck. Too low and the system regresses to generic AI output. The Head of Content tunes the threshold weekly off the rejection log.
- 05** Long-run the agent with quarterly retraining off newly published, voice-passed work. The voice document evolves as the brand evolves. The agent must evolve with it.

SUCCESS METRICS · WHAT TO MEASURE

Voice-pass rate on first review	Percent of submitted content that passes voice review without revision. Target: > 70% after 90 days of agent operation.
Off-voice incidents post-publication	Number of pieces flagged as off-voice by senior brand readers after publication. Target: < 1 per quarter.
Vocabulary discipline	Percent of banned-phrase usage caught before publication. Target: 100%.
Editor time saved	Reduction in editor hours spent on voice fixes vs. baseline. Target: 50% reduction within 90 days.

The Content Strategy Agent

PURPOSE. Translates business priorities, ICP segments, and search/AEO opportunities into a quarterly content plan. Owns topic selection, narrative tie-in, format choice, and publishing cadence. The agent that makes editorial scale.

INPUTS. ICP document, last 90 days of pipeline data, top-performing recent content, search and AEO query data, sales call themes, current quarterly priorities, content gap analysis.

OUTPUTS. Quarterly content calendar (M/W/F cadence), per-piece briefs with strategic-narrative tie-in, target keyword and AEO query for each, internal linking recommendations.

MEMORY. Prior 12 weeks of published content with performance data, voice notes from feedback cycles, current strategic priorities, the topics that are working and the topics that are not.

HUMAN OVERSIGHT. Quarterly plans reviewed by Head of Content before activation. Briefs spot-checked weekly for narrative tie-in and ICP alignment.

COMMON MISTAKES. Recommending generic SEO topics with no narrative tie. Drifting toward content that ranks easily rather than content that converts. Treating content as a topic problem instead of an audience and angle problem.

HOW TO DEPLOY

The Content Strategy Agent decides what gets written, why, and against which pillar. It does not decide how content gets written. Deploy it after the ICP, Persona, and Brand Voice agents are operating, because strategy without those three is content astrology.

- 01** Establish the pillar architecture. Three to five strategic content pillars, each tied to a specific buyer outcome the brand competes on. The agent works within this architecture, not against it.
- 02** Connect the agent to the ICP Research Agent's grades, the Buyer Persona Agent's briefs, the Intent Signal Analysis Agent's reports, and the Analytics & Reporting Agent's content-performance data. Strategy without measurement context defaults to taste.
- 03** Produce the first quarterly content plan in shadow mode. Compare to the existing plan. Every diff is a strategic conversation, not an agent error.
- 04** Promote to live state. The agent produces the rolling 90-day plan with rationale per topic. The Head of Content approves the plan, the Editorial Planning Agent sequences it.
- 05** Run a monthly retrospective. Topics that ranked, converted, or drove pipeline get reinforced. Topics that did not get cut. The agent's long-run accuracy depends on this discipline.

SUCCESS METRICS · WHAT TO MEASURE

Plan-to-publish ratio Percent of planned topics that publish within the planned window. Target: > 80%.

Topic-to-pipeline correlation Percent of pipeline-creating content traceable to a planned topic in the prior 90 days. Target: > 40%.

Pillar balance Distribution of published content across the defined pillars. Target: within 10% of the planned distribution.

Strategic plan velocity Time from market signal to topic-on-plan. Target: < 14 days.

The Editorial Planning Agent

PURPOSE. Sequences content across the calendar to build narrative arcs (cluster strategy), respect channel cadence, and align with launch and event milestones. The Editorial Planning Agent runs alongside Content Strategy: Strategy decides what; Editorial decides when, in what order, and how it ladders.

INPUTS. Content Strategy Agent's quarterly plan; product launch calendar; event calendar; field marketing motion calendar; SDR campaign cadence; SEO/AEO cluster maps.

OUTPUTS. 12-week sequenced editorial calendar; cluster maps (cornerstone + cluster posts that internally link); flight-plan dependencies (which pieces must publish before which others to land the narrative arc).

MEMORY. Past clusters with traffic and pipeline data, channel cadence patterns that worked, sequencing decisions that backfired.

HUMAN OVERSIGHT. Editorial director approves cluster strategy and cornerstone selections. Weekly stand-up reviews shifts to the calendar.

COMMON MISTAKES. Over-stuffing the calendar past sustainable cadence. Treating each post as standalone instead of part of a cluster. Failing to coordinate editorial sequence with field motion.

HOW TO DEPLOY

The Editorial Planning Agent sequences the calendar. It is the operating mechanism behind the Content Strategy Agent's plan. It does not decide what to write. It decides when, in what order, and against which dependencies.

- 01** Build the dependency model first. Cluster strategy means cluster pieces ship in a defined sequence, not randomly. The agent needs to know which pieces support which, which precede which, and what the cadence rules are per channel.
- 02** Connect the agent to the Content Strategy Agent's 90-day plan and to the launch and event calendar maintained by marketing and product. Editorial sequencing in isolation from launches is a coordination failure.
- 03** Produce the first six-week sequenced calendar in shadow mode. Editorial leadership reviews for cluster integrity, channel cadence, and launch alignment.
- 04** Promote to live state. The agent maintains the rolling calendar, flags conflicts, and proposes re-sequencing when launches or events shift.
- 05** Audit monthly. Where the calendar slipped, identify whether the slip was an agent error, a content production bottleneck, or a strategic shift the agent could not have anticipated. The agent learns from the audit log.

SUCCESS METRICS · WHAT TO MEASURE

On-cadence publish rate	Percent of planned pieces published within the agent's scheduled window. Target: > 85%.
Cluster integrity	Percent of cluster pieces that ship in the defined sequence. Target: > 90%.
Conflict-flag lead time	Median time the agent flags a calendar conflict before the conflict materializes. Target: > 7 days.
Replanning rate	Number of unscheduled replanning cycles per quarter. Target: < 3.

The SEO Agent

PURPOSE. Optimizes content for traditional search engines: keyword research, on-page structure, internal linking, schema, technical hygiene. Operates in tight pairing with the AEO Optimization Agent – the two are complementary, not interchangeable.

INPUTS. Search query data, competitor SERP analysis, keyword difficulty data, internal site analytics, the editorial calendar, the brand's authority signals.

OUTPUTS. Per-piece SEO brief (target keyword, semantic cluster, recommended H-structure, FAQ candidates, schema requirements, internal linking targets), monthly site-health audit, quarterly content gap analysis.

MEMORY. Keyword performance over time, ranking-to-pipeline correlation, internal linking topology of the site.

HUMAN OVERSIGHT. SEO lead approves cornerstone-level keyword targets. Spot-checks per-piece briefs weekly. Quarterly pipeline-correlation audit against keyword choices.

COMMON MISTAKES. Optimizing for ranking ease over pipeline contribution. Keyword stuffing that degrades quality. Internal linking that follows topology rather than reader journey.

HOW TO DEPLOY

The SEO Agent owns search visibility against the pillar architecture. It is not a keyword research tool. It is a long-running strategist that watches what is winning in your category and recommends topic, structure, and internal-link decisions.

- 01** Establish the keyword-cluster framework. Each strategic pillar gets a defined keyword cluster, mapped against intent stage (informational, comparison, transactional). The agent works within this framework.
- 02** Connect the agent to Google Search Console, your rank-tracking tool of choice, and competitor SEO data. Without competitor visibility, the agent is optimizing in a vacuum.
- 03** Produce the first monthly SEO opportunity report in shadow mode. Investigate the top five recommendations manually. Most will be sound. The ones that are not will reveal where the cluster framework needs refinement.
- 04** Promote to live state. The agent now feeds topic recommendations to the Content Strategy Agent, internal-link recommendations to the Editorial Planning Agent, and on-page recommendations directly to the content producer.
- 05** Run a quarterly SEO audit. The agent's top recommendations should map to the quarter's rank gains. Where they do not, retrain.

SUCCESS METRICS · WHAT TO MEASURE

Pillar-cluster ranking velocity

Net rank improvement across the defined cluster keywords per quarter. Target: 15+ position improvement per quarter on tracked keywords.

Topic-to-ranking conversion

Percent of agent-recommended topics that achieve a page-one rank within 6 months. Target: > 40%.

Internal-link discipline

Percent of new content shipping with the agent's recommended internal link structure. Target: > 90%.

Cluster authority growth

Domain rating or equivalent topical authority signal per pillar. Target: monthly positive movement.

The AEO Optimization Agent

PURPOSE. Optimizes content for retrieval and citation by AI search engines and large language models – Answer Engine Optimization. The newest and fastest-growing agent in the cluster, because the buyer increasingly starts research inside an AI interface rather than a search box.

INPUTS. AI search citation data (Perplexity, ChatGPT, Claude, Google AI Overviews where measurable), structured-data audit, FAQ schema coverage, entity clarity audit, "answer-first" copy patterns, canonical question taxonomy for the category.

OUTPUTS. Per-piece AEO brief (canonical question phrasing, answer-first lede, semantic depth requirements, FAQ schema block, entity tagging, citation-worthy stat list), monthly AEO citation report, quarterly entity-graph audit.

MEMORY. Pieces that have earned citations in AI surfaces, citation-decay patterns, the answer phrasings that AI engines reliably retrieve.

HUMAN OVERSIGHT. Editorial director approves canonical question taxonomy. SEO/AEO lead reviews citation reports monthly.

COMMON MISTAKES. Treating AEO as a flavor of SEO instead of a distinct retrieval target. Stuffing FAQ blocks with low-quality questions. Failing to tag entities so the AI can resolve who the brand is and what it is authoritative on.

HOW TO DEPLOY

The AEO Optimization Agent owns AI search visibility. Answer Engine Optimization is now the primary distribution channel for B2B content, surpassing traditional SEO in influence for many comparison and evaluation searches. This is the agent that wins or loses you that surface.

- 01** Catalogue every AI search surface that matters for your buyer. ChatGPT, Perplexity, Claude, Google's AI Overviews, Bing Copilot. The agent works against the defined set. Generic AEO without a defined surface set is wasted effort.
- 02** Establish the AEO content patterns. Definitions written with citation-friendly structure, comparison tables that travel well across models, expert-quote density, source citations. The agent enforces these patterns.
- 03** Run weekly AEO visibility checks. The agent queries each defined surface with a defined query set and captures citation presence, position, and competitive coverage. The output is the AEO Visibility Report.
- 04** Connect the agent to the SEO Agent's topic plan. AEO and SEO topics overlap heavily, but AEO requires structural patterns SEO does not. The agent ensures both are satisfied.
- 05** Iterate on the content patterns monthly. Models change. What got cited last month may not get cited next month. The agent's long-run job is to keep up.

SUCCESS METRICS · WHAT TO MEASURE

Citation rate	Percent of defined query set where the brand is cited by at least two AI engines. Target: > 30% after 90 days.
Comparison-query share of voice	Percent of brand-vs-competitor queries where the brand is positioned favorably. Target: > 60%.
AEO content velocity	Number of new pieces shipping per month with full AEO pattern compliance. Target: > 50% of all new content.
Surface coverage	Percent of defined AI surfaces where the brand is meaningfully visible. Target: 100% within 180 days.

The LinkedIn / Social Media Agent

PURPOSE. Drafts native-format LinkedIn posts (and other social media), distinct executive-voice posts, and re-share hooks that land in B2B feeds. Optimizes for hook quality, native pacing, and executive POV — not for recycled blog summaries, which is the most common failure mode.

INPUTS. Recent blog posts with extractable POVs; the executive's voice samples; news in the category; engagement patterns by post format; the editorial calendar.

OUTPUTS. Hook-first LinkedIn drafts (executive and company-page variants); thread-format drafts where appropriate; reshare hooks for the team to engage; monthly "what's working in the feed" briefing.

MEMORY. Past posts with engagement data, hook patterns that landed, executive voice exemplars, topic-engagement correlation.

HUMAN OVERSIGHT. Executive reviews any post under their byline before posting. Editor reviews company-page drafts. Monthly engagement audit.

COMMON MISTAKES. Generic "thought leader" tone that sounds like every other post. Recycling blog content into LinkedIn instead of writing native-format from the same insight. Over-posting and exhausting the audience.

HOW TO DEPLOY

The LinkedIn / Social Media Agent owns the social-media muscle of the marketing organization. It produces content adapted to the platform's reading patterns, not blog posts copied into a status update. It is one of the highest-impact agents for B2B brand visibility.

- 01** Define the social voice. It is related to the brand voice but not identical. LinkedIn rewards specific formats, hooks, cadence, and POV that pure brand voice often does not capture. The agent has its own voice sub-document.
- 02** Connect the agent to the Content Strategy Agent's long-form pipeline. Every published long-form piece produces a defined number of social adaptations. The agent does not invent unconnected social content.
- 03** Establish the executive co-presence model. The CEO, the CMO, and selected operator voices each get their own posting cadence and content allocation. The agent operates against named voices, not a generic brand account.
- 04** Run weekly content batches in shadow mode for three weeks. The Head of Content reviews voice fidelity, hook quality, and POV strength. Adjust the voice sub-document accordingly.
- 05** Promote to live state. The agent now batches content weekly, schedules through the publishing platform, and produces a monthly performance report with hook-pattern analysis and next-period recommendations.

SUCCESS METRICS · WHAT TO MEASURE

Engagement rate vs. baseline	Average engagement rate (likes, comments, reshares per impression) vs. pre-agent baseline. Target: 2x baseline within 90 days.
Reach growth on named voices	Quarterly follower growth on the executive and operator voices the agent supports. Target: 5%+ per quarter.
Inbound from social	Number of inbound demo or advisory inquiries attributable to social content per quarter. Target: trend up, with a measurable baseline by month 3.
Voice fidelity	Percent of posts approved on first review without voice-related revision. Target: > 80%.

The Ad Copy Agent

PURPOSE. Drafts paid media ad copy across LinkedIn, Google, Meta, programmatic, and connected TV – adapted to channel-specific format, length, and audience. Generates A/B test variants with documented hypotheses, not just text changes.

INPUTS. Persona briefs, ICP fit reasoning, recent landing-page performance, prior winning ad copy, current campaign objective, channel-specific format constraints.

OUTPUTS. Channel-specific ad copy variants (typically 5-7 per ad set), A/B test hypothesis matrix, copy-to-landing-page consistency check, per-variant audience match note.

MEMORY. Winning copy patterns by channel and persona, copy-to-pipeline correlation, channel-format updates as platforms change them.

HUMAN OVERSIGHT. Paid media lead approves all paid copy before launch. Brand Voice Governance Agent scores copy before launch.

COMMON MISTAKES. Optimizing for click-through over qualified pipeline. Failing to align ad copy with the landing page experience. Generating variants that test letters and not hypotheses.

HOW TO DEPLOY

The Ad Copy Agent produces and iterates the variants for every paid channel. It is one of the easiest agents to deploy and one of the easiest to deploy badly.

Performance discipline matters more than creative volume.

- 01** Build the variant taxonomy first. Headline variant, body variant, CTA variant. Channel variant. Persona variant. Stage variant. The agent does not generate randomness, it generates intentional variants against a defined taxonomy.
- 02** Connect the agent to the Buyer Persona Agent's briefs and the Brand Voice Governance Agent. Persona language drives the variant. Voice governance prevents quality collapse at volume.
- 03** Establish the testing discipline. Two-cell or three-cell A/B/C tests with statistical-significance thresholds, defined test duration, and a single primary metric per test. The agent does not run uncontrolled experiments.
- 04** Run the first 30 days against existing campaigns. Compare the agent's top-performing variants against the human-written controls. Promote the agent on the channels where it wins, keep it in shadow on the channels where it loses, and investigate the losses.
- 05** Long-run on a weekly variant-batch cadence with a monthly performance debrief. The agent's job is to keep the winning patterns winning and to retire the losing patterns before they cost too much spend.

SUCCESS METRICS · WHAT TO MEASURE

CTR vs. control	Average click-through rate of agent variants vs. human-written controls. Target: 1.2x control within 60 days.
Cost per qualified lead	CPL on agent-led variants vs. baseline. Target: 15% reduction within 90 days.
Variant burn rate	Number of variants that hit fatigue (CTR decay > 25%) per month. Target: high in the first month, then declining, indicating the agent has learned the durable patterns.
Approval-pass rate	Percent of generated variants that pass voice and compliance review on first submission. Target: > 85%.

The Hero *Workflow.*

HERO WORKFLOW

Cornerstone content production.

How one strategic theme becomes ten distributed assets.

The cornerstone workflow takes a single strategic theme and turns it into a cornerstone long-form piece, supporting cluster content, atomized social, sales-enablement asset, and a nurture sequence. The same intellectual capital, distributed across every channel marketing owns. Six agents, four human-in-the-loop checkpoints, one cornerstone piece. Time-to-publish drops from 8-12 hours of human time to 3-4 hours, with quality going up because every step is governed.

AGENTS INVOLVED

- Content Strategy Agent (selects the theme and pillar)
- SEO Agent (validates keyword cluster opportunity)
- AEO Optimization Agent (ensures AI-search compliance)
- Editorial Planning Agent (sequences supporting pieces)
- Brand Voice Governance Agent (quality floor at every step)
- LinkedIn / Social Media Agent (atomizes for executive voices)
- Sales Enablement Agent (produces field-ready summary)
- Email Nurture Sequence Agent (builds the nurture asset)

HUMAN CHECKPOINTS

- *Theme approval by the Head of Content before drafting begins*
- *Cornerstone draft review before publication*
- *Final social-batch sign-off on executive voices*

THE WORKFLOW, STEP BY STEP

- 01** Theme selection. The Content Strategy Agent proposes the quarterly themes against pillar coverage and SEO opportunity. The Head of Content approves one cornerstone theme per quarter.
- 02** Brief assembly. The Editorial Planning Agent assembles the cornerstone brief: target keyword cluster, AEO patterns to enforce, supporting-piece outline, distribution plan. The brief is the single source of truth for the workflow.
- 03** Cornerstone drafting. A writer (human or AI) drafts the cornerstone piece against the brief. The Brand Voice Governance Agent reviews. Revisions cycle until the voice-pass threshold is met.
- 04** Supporting cluster production. The Editorial Planning Agent sequences four to six supporting pieces over the following six weeks, each one targeting a sub-cluster keyword and linking to the cornerstone.
- 05** Social atomization. The LinkedIn / Social Media Agent produces a defined number of social posts across named executive voices, scheduled across four weeks post-publication.
- 06** Sales asset generation. The Sales Enablement Agent produces a one-page summary of the cornerstone framed for the field. It is added to the standard enablement library.
- 07** Nurture sequence build. The Email Nurture Sequence Agent constructs a five-email post-download sequence using the cornerstone as anchor content.

SUCCESS METRICS

- Cornerstone ranking velocity. Target: page-one rank on the target cluster within 90 days.
- Cluster authority signal. Target: 5-10 supporting pieces published in the 6 weeks after cornerstone.
- Cornerstone-attributable pipeline. Target: at least one closed-won account attributable to cornerstone content per quarter.
- Social atomization yield. Target: 12-18 social posts produced per cornerstone, across two or more executive voices.

COMMON FAILURE MODES

- Theme drift. The cornerstone gets diluted into a generic explainer because the brief was too broad. Fix: every brief specifies the contrarian POV and the explicit thesis.
- Atomization without atomization. Posts copy-paste from the cornerstone instead of making cornerstone arguments at platform-appropriate length. Fix: the social agent has its own voice sub-document, not just the cornerstone text.
- No nurture follow-through. The cornerstone produces downloads but no nurture sequence captures them. Fix: nurture sequence ships in the same sprint as the cornerstone, not as a follow-up.

Pre-flight *checklist.*

Content agents are the highest-velocity agents in the stack. They are also the agents where brand dilution shows up first if the Foundation is weak. Resolve every item below before activating.

- The brand voice document exists, is owned by a single human, and has been validated against fifteen exemplar pieces of published work. Voice governance without a real Foundation document is theater.

- Three to five strategic content pillars are defined, each tied to a specific buyer outcome the brand competes on.

- The persona briefs that drive content language are validated, with verbatim quotes from real buyer conversations. Personas built on demographic guesses produce content that sounds like everyone else.

- Editorial workflow is documented from idea to publication, with every approval step named and owned. The Editorial Planning Agent will run this workflow. Without it, the agent invents one and it will not match your reality.

- SEO and AEO tracking infrastructure is in place. Without rank tracking on the SEO side and AI-citation tracking on the AEO side, the relevant agents cannot prove value.

- Distribution channels are inventoried with channel-specific cadence rules. The LinkedIn / Social Media Agent runs against named voices, not a generic brand account.

- Performance feedback loop is wired to the Analytics & Reporting Agent. Content velocity without performance feedback is content for content's sake.

The 30-60-90 *Roadmap.*

The content-stack rollout sequence. Order matters. Skipping phases produces brand drift, not leverage.

DAYS 1-30

Stand up the Foundation. The brand voice document is the highest-priority artifact. Audit existing content for voice consistency, identify the voice exemplars, and validate the voice document against them. Re-audit the persona briefs and refresh any that are more than nine months old. Define the pillar architecture if it does not exist. Inventory the existing content library, mark archive vs. active, and identify the SEO/AEO content gaps the strategy will address. No content agent ships in this phase.

DAYS 31-60

Deploy the Brand Voice Governance Agent and the SEO Agent. Voice governance is the quality floor. It must operate before any volume-producing agent. The SEO Agent runs in advisory mode this phase, feeding topic and on-page recommendations to the human content team. The Editorial Planning Agent stands up in shadow mode, sequencing the human-authored calendar so the team can validate the agent's sequencing logic before the calendar becomes agent-authored.

DAYS 61-90

Deploy the Content Strategy Agent and the AEO Optimization Agent. Promote the Editorial Planning Agent to live state. The Content Strategy Agent now produces the rolling 90-day plan, approved by the Head of Content, sequenced by the Editorial Planning Agent. The AEO Optimization Agent runs the weekly AI-citation visibility check. Cornerstone Content workflow goes live, taking one strategic theme per quarter from brief to published cornerstone to social atomization to nurture asset, all coordinated by the agent stack.

DAY 91+

Deploy the LinkedIn / Social Media Agent and the Ad Copy Agent. Both are downstream agents that depend on the upstream content infrastructure being in place. Long-run the governance discipline: monthly voice-pass-rate audits, quarterly cornerstone retrospectives, ongoing keyword-cluster and AEO-pattern recalibration.

Pitfalls & *remediation.*

PITFALL 01

Deploying volume agents before voice governance.

The fastest way to dilute a brand is to ship AI-generated content at volume with no voice review. The Brand Voice Governance Agent must operate before the Content Strategy Agent, the SEO Agent, or any volume-producing agent.

PITFALL 02

Treating AEO as SEO with a new label.

AI search rewards different content patterns than traditional search. Comparison tables that travel well across models. Citation-friendly definitions. Expert-quote density. Teams that bolt AEO onto an existing SEO workflow without restructuring the content patterns lose AI citations to competitors who restructured.

PITFALL 03

Letting the agent stack optimize for output volume.

Forty pieces of average content lose to four pieces of cornerstone content every time. The Content Strategy Agent must optimize for strategic-pillar coverage and cornerstone development, not output count.

WHERE THIS FITS

Where this fits.

The seven agents and one workflow in this playbook are the content-marketing slice of a larger marketing operating system. The full architecture covers strategy, demand, measurement, revenue, ops, and governance.

The flagship pulls all of it together: **AI Agents for Marketing Teams**. Seventy-two pages, free, no email gate.

— Erik R. Miller

Erik R. Miller

SENIOR B2B MARKETING EXECUTIVE · FRACTIONAL CMO

Web: erikrmiller.com

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

LinkedIn: linkedin.com/in/erikrm

Facebook: facebook.com/ERMA Advisory

Instagram: instagram.com/ermadvisory

Advisory: erikrmiller.com/#contact

Content velocity broke human throughput. An operating system is the rebuild.

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