E.R.M. ADVISORY

THE OPERATOR PLAYBOOK · VOL. 07 · 2026 EDITION

SALES ENABLEMENT & FIELD MARKETING

AI Agents for *Sales Enablement.*

Five agents and one workflow that close the handoff between marketing and the field.

5

AGENTS

1

WORKFLOW

23

PAGES

WRITTEN BY

Erik R. Miller

SENIOR B2B MARKETING EXECUTIVE · FRACTIONAL CMO

AI AGENTS FOR SALES ENABLEMENT & FIELD MARKETING

THE OPERATOR PLAYBOOK · VOL. 07

FIRST EDITION, 2026

NEW YORK CITY

ERIKRMILLER.COM

THE OPERATOR NEWSLETTER · THEOPERATOR-ERM.BEEHIIV.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.

INTRODUCTION · VOL. 07

The handoff is where marketing wins or loses.

Sales enablement and field marketing live at the seam between marketing and revenue. The seam is where most B2B companies leak the value marketing actually produced. The deck that did not get used. The briefing the AE never read. The event motion that produced names but not pipeline. The RFP that was answered with last year's template.

Five agents close the seam. Sales enablement assets that are account-specific, not generic. Meeting notes that synthesize themselves into pipeline updates the AE can actually use. Executive briefings that pull from CRM, news, and account intelligence at draft time, not at the last minute. Event coordination that runs the pre-event, on-site, and post-event motion as one workflow. RFP support that produces compliant responses in days, not weeks.

This playbook covers the five agents and one workflow that anchor the marketing-to-revenue handoff. If you sit at this seam, this Field Note is the document to plan against.

> **WHO THIS IS FOR**
>
> *For sales enablement leaders, field marketing managers, and the operators sitting at the marketing-to-revenue handoff.*

CONTENTS

Inside this *Field Note.*

5 agents and 1 hero workflow.

THE AGENT ROSTER

HERO WORKFLOW

OPERATIONS & EXECUTION

WHERE THIS FITS

The *Sales Enablement.* Agents.

AGENT 01

The Sales Enablement Agent

PURPOSE. Translates marketing campaigns, ICP shifts, and competitive intelligence into sales-ready talk tracks, battlecards, sequence templates, and objection-handling scripts. Builds the assets sales actually uses, not the assets that sit unread in a content portal.

INPUTS. Campaign briefs, ICP and persona docs, competitive intelligence, recent objections from sales calls, product positioning, current quarterly priorities, win/loss reasons.

OUTPUTS. SDR talk tracks, AE call openers, battlecards, sequence templates, objection-handling pages, monthly "what changed" digest for the sales team.

MEMORY. Sales-feedback log on which assets got used, which got ignored; objection patterns over time; win/loss themes.

HUMAN OVERSIGHT. Sales leadership co-owns this agent with marketing. Weekly sync on what is and is not landing.

COMMON MISTAKES. Building enablement assets sales never uses (volume without adoption). Disconnect from sales-call reality (the agent reads marketing's positioning rather than the field's reality). Failing to update when ICP shifts.

HOW TO DEPLOY

The Sales Enablement Agent is the marketing-to-revenue interface. It produces account-specific, situation-specific assets at a velocity the sales team can actually use. Deploy it after the ICP, persona, and content agents are live, because enablement built on generic foundations stays generic.

01 Audit the existing enablement library. Most of it will be generic decks that no AE has used in twelve months. Archive the generic. Identify the three to five asset types AEs actually use (battlecard, account-specific deck, ROI calculator, executive briefing template, RFP language).

02 Connect the agent to the Account Selection Agent's tier data, the Buyer Persona Agent's briefs, the Competitor Intelligence Agent's battlecards, and the CRM's account context.

03 Run the agent on the top 30 active Tier 1 accounts. Produce per-account assets. Have the named AE rate each asset on relevance, accuracy, and usability. Iterate until the ratings stabilize.

04 Promote to live state. The agent produces enablement assets on AE request inside the CRM. SLA is 24 hours or less for standard request types.

05 Long-run on a monthly utilization audit. Assets that go unused get retired. Assets that produce high meeting-conversion rates get reinforced as templates.

SUCCESS METRICS · WHAT TO MEASURE

AE-request fulfillment time	Median time from AE request to delivered asset. Target: < 24 hours for standard, < 72 hours for custom.
Asset utilization	Percent of agent-produced assets that get sent to a prospect within 7 days. Target: > 65%.
Meeting-conversion lift	Conversion from first meeting to second meeting on accounts using agent-produced assets vs. baseline. Target: 1.3x baseline.
AE satisfaction	Quarterly AE rating of enablement quality. Target: > 4 out of 5.

AGENT 02

The Meeting Notes & Summarization Agent

PURPOSE. Summarizes meetings (internal staff, sales calls, customer interviews, executive briefings) into structured action-item briefs. Decisions made, owners, deadlines, follow-up tasks. The most boring high-leverage agent in the system.

INPUTS. Meeting transcript or recording, attendees, meeting type and purpose, prior context for the meeting topic.

OUTPUTS. Structured summary (Decisions / Owners / Deadlines / Open questions / Follow-up tasks), meeting-type-specific format (sales call vs. customer interview vs. internal staff differ), CRM-pushable action items.

MEMORY. Meeting-type templates; recurring meeting series with prior actions; speaker-attribution patterns.

HUMAN OVERSIGHT. Owner of each meeting type validates summary format quarterly. High-stakes meetings (board, executive, customer escalations) require human review before distribution.

COMMON MISTAKES. Missing subtle disagreements that did not become explicit in the meeting. Over-aggregating, losing nuance. Pushing items to CRM without owner agreement.

HOW TO DEPLOY

The Meeting Notes & Summarization Agent converts the raw signal inside sales conversations into structured, actionable intelligence the rest of the marketing system can use. It is one of the highest-leverage agents in the stack because it surfaces voice-of-customer signal that would otherwise stay locked in the call recording.

01 Connect the agent to the call-recording platform (Gong, Chorus, or equivalent) and the CRM. The agent reads call transcripts and writes structured outputs back to the account record.

02 Define the extraction schema. Pain points (verbatim), competitor mentions, objection patterns, decision-criteria, timeline signals, persona-level quotes. The agent extracts only against the defined schema.

03 Pipe the structured output into the Buyer Persona Agent, the Competitor Intelligence Agent, and the Content Strategy Agent. Voice-of-customer signal that does not flow into other agents is wasted.

04 Run the first 30 days in shadow mode. Sales and marketing leadership review weekly summaries for accuracy and value. Adjust the schema based on what is most useful.

05 Promote to live state. The agent now produces daily summary briefs to AE and account-team Slack channels and feeds the marketing system's feedback loops automatically.

SUCCESS METRICS · WHAT TO MEASURE

Extraction accuracy	Percent of extracted pain points, objections, and competitor mentions verified as accurate by the AE on the call. Target: > 90%.
Signal-to-content cycle time	Median time from a recurring objection appearing in three or more calls to that objection being addressed in published content. Target: < 30 days.
AE time saved	AE hours per week saved on note-taking, vs. baseline. Target: > 4 hours per AE per week.
Cross-agent flow	Percent of extracted signals that successfully populate downstream agents. Target: 100%.

AGENT 03

The Executive Briefing Agent

PURPOSE. Drafts executive-ready briefs, board updates, customer-meeting prep documents, and pre-read packets. Voice-matched to the named executive. Designed for exec review, never auto-sent.

INPUTS. Topic, recipient profile, prior correspondence with the recipient, the executive's voice samples, the strategic narrative, recent performance data.

OUTPUTS. Draft briefs in three lengths (one-page, three-page, board-pack), executive voice-matched, with embedded data and recommended decisions.

MEMORY. Executive voice exemplars, prior board/customer meeting briefs, recipient-specific context.

HUMAN OVERSIGHT. Always-on. The executive personally reviews every briefing under their byline. Voice drift undermines trust faster here than anywhere else.

COMMON MISTAKES. Voice drift. Generic executive tone that does not sound like the named person. Burying the recommended decision instead of leading with it.

HOW TO DEPLOY

The Executive Briefing Agent produces the one-page brief an AE walks into an executive meeting with. It pulls from CRM, news, account intelligence, and persona research, and synthesizes it into the document the AE actually opens in the conference room.

01 Define the brief template. One page. Account context, recent triggers, named-stakeholder intelligence, suggested conversation arc, suggested next step. The agent works against this template, not freelance.

02 Connect the agent to the CRM, the Meeting Notes & Summarization Agent's output, the Account Selection Agent's tier data, and a news/intelligence feed for the named accounts.

03 Run the first 20 briefs in advisory mode. AEs read them before live meetings. AEs rate the brief on usefulness, accuracy, and freshness. Adjust the agent against the ratings.

04 Promote to on-demand state. AEs request a brief 24-48 hours before any executive meeting. The agent produces it on request.

05 Long-run on a quarterly brief-usefulness audit. Sections of the brief that AEs consistently ignore get cut. Sections they consistently flag as missing get added.

SUCCESS METRICS · WHAT TO MEASURE

Brief-request rate	Percent of executive meetings that have an agent-produced brief in advance. Target: > 80%.
AE usefulness rating	Average AE rating of brief usefulness immediately after the meeting. Target: > 4.2 out of 5.
Brief-to-second-meeting rate	Conversion from first executive meeting to second meeting where the brief was used vs. baseline. Target: 1.3x baseline.
Freshness	Percent of intelligence in the brief sourced within 30 days of the meeting. Target: > 80%.

AGENT 04

The Event Marketing Coordination Agent

PURPOSE. Plans and executes virtual and in-person events end-to-end — from objective and audience definition through run-of-show, content, comms, and follow-up sequences. Connects events to pipeline rather than treating them as standalone marketing motions.

INPUTS. Event objective, target ICP, target persona, budget, prior event data, current quarterly priorities, brand voice samples for comms.

OUTPUTS. Event run-of-show, content plan (sessions, abstracts, speakers, panel topics), comms calendar (pre, during, post), targeted invitation list, follow-up nurture sequences keyed to attendee behavior.

MEMORY. Past events with attendance, engagement, and pipeline outcomes; what content drew which audience; what follow-up sequences converted.

HUMAN OVERSIGHT. Field-marketing lead owns this agent. Weekly check-ins during the run-up to a major event. Post-event pipeline audit within 60 days.

COMMON MISTAKES. Disconnecting event ROI from pipeline. Optimizing for attendance instead of pipeline-qualified audience. Treating follow-up as one sequence instead of segmented by behavior.

HOW TO DEPLOY

The Event Marketing Coordination Agent runs the pre-event, on-site, and post-event motion as one coordinated workflow. It is the agent that turns event spend from a leads-list-of-questionable-quality into a measurable pipeline contributor.

01 Build the event-motion playbook. Pre-event outreach to known attendees in the ICP. On-site coordination with sales. Post-event follow-up sequences segmented by attendee behavior. Each event motion follows the playbook.

02 Connect the agent to the Account Selection Agent, the Email Nurture Sequence Agent, the event-platform attendee data, and the CRM. Pre-event targeting is impossible without ICP-graded attendee data.

03 Run the agent on one tier-one event end-to-end. Compare against the previous-quarter's equivalent event. The comparison is the agent's first proof point.

04 Promote to default coordination mode for every event over a defined spend threshold. Smaller events stay in advisory mode unless explicit scope expansion.

05 Quarterly, audit the agent's motion against pipeline outcomes. Events that produced pipeline above benchmark get reinforced. Events that did not get re-scoped or retired.

SUCCESS METRICS · WHAT TO MEASURE

Pre-event ICP-meeting set rate	Percent of Tier 1 attendees with a pre-scheduled on-site meeting. Target: > 30%.
Post-event conversion	Percent of attendees converting to a follow-up sales conversation within 30 days. Target: 1.5x previous-event baseline.
Event-to-pipeline cycle time	Median time from event to first pipeline entry for attendees. Target: < 45 days.
Cost per pipeline dollar	Total event cost per dollar of attendee-sourced pipeline. Target: 25% improvement vs. baseline.

AGENT 05

The RFP / Proposal Support Agent

PURPOSE. Drafts proposal sections from the approved content library — security, integration, pricing, customer-success, vertical-specific case studies. Accelerates RFP turnaround without sending boilerplate that fails to answer the actual RFP. The agent that earns its keep on a single won deal.

INPUTS. The RFP document, prior winning proposals, customer-specific context, the content library, current product positioning, pricing rules.

OUTPUTS. First-draft proposal sections, response matrix mapping each RFP question to a draft answer, gaps requiring SME input, customer-specific tailoring.

MEMORY. Past RFPs with win/loss, the answers that won and the answers that lost, vertical-specific positioning that worked.

HUMAN OVERSIGHT. Sales lead and SMEs review every section before submission. The agent never submits an RFP — it accelerates the SME's first draft, no more.

COMMON MISTAKES. Boilerplate that does not address the specific RFP. Failing to tailor for the customer's actual context. Producing volume that hides gaps.

HOW TO DEPLOY

The RFP / Proposal Support Agent collapses RFP response time from weeks to days. It maintains the canonical response library, generates first-draft responses against the buyer's questions, and flags any new compliance language that needs human review. It pays for itself in the first response cycle.

01 Catalogue every previously won and lost RFP from the last twelve months. Extract the response-language library, question patterns, and compliance language. This becomes the agent's working corpus.

02 Establish the response-quality bar. Every agent-drafted response gets reviewed against the bar before submission. The bar is owned by the head of solutions or the equivalent role, not the agent.

03 Run the first three RFP responses in co-drafting mode. The agent produces the first draft, a human solutions writer revises, the agent learns from the revisions.

04 Promote to first-draft default state for RFPs under a defined complexity threshold. Higher-complexity RFPs stay in co-drafting mode permanently.

05 Long-run a quarterly library-hygiene cycle. Outdated response language gets retired. New product capabilities get added. Compliance language gets updated. The corpus only stays useful if it stays current.

SUCCESS METRICS · WHAT TO MEASURE

RFP cycle-time reduction	Median time from RFP receipt to submission with the agent vs. baseline. Target: > 60% reduction.
Win rate maintenance	Win rate on agent-supported RFPs vs. historical baseline. Target: at parity or higher.
Compliance accuracy	Number of compliance-language errors in submitted responses. Target: 0.
Library reuse rate	Percent of response content reused from the library across multiple RFPs. Target: > 60%.

The Hero *Workflow.*

HERO WORKFLOW

SDR-marketing alignment cycle.

How the SDR-marketing alignment cycle stops being a meeting and starts being a system.

The SDR-marketing alignment cycle closes the feedback loop between what marketing produces and what sales actually uses. Most companies run this as a monthly meeting. The agent stack runs it as a continuous workflow: meeting notes flow into voice-of-customer, voice-of-customer flows into content and enablement, enablement flows back into sales, sales feedback flows into the next iteration.
Five agents, one human checkpoint, a closed-loop cycle that keeps marketing's view of the ICP and field reality from drifting apart. The most important alignment workflow in the system, and the one most teams skip.

AGENTS INVOLVED

- Meeting Notes & Summarization Agent (extracts signal from calls)
- Buyer Persona Agent (consumes voice-of-customer)
- Content Strategy Agent (consumes objection patterns)
- Sales Enablement Agent (produces and updates per-rep assets)
- Executive Briefing Agent (supplies pre-meeting briefs)
- Competitor Intelligence Agent (updates battlecards from field signals)

HUMAN CHECKPOINTS

- *AE rating of enablement assets after every use*
- *Monthly SDR-marketing alignment review (the meeting still exists, but it is short and decision-focused)*

THE WORKFLOW, STEP BY STEP

01 Call transcript ingestion. The Meeting Notes & Summarization Agent reads every sales conversation transcript and extracts pain points, objection patterns, competitor mentions, decision-criteria, timeline signals, and verbatim quotes.

02 Voice-of-customer routing. Extracted signals flow to the Buyer Persona Agent, the Content Strategy Agent, and the Competitor Intelligence Agent. Each agent updates its working corpus.

03 Content response. The Content Strategy Agent identifies any objection appearing in three or more calls in the past 30 days. Each pattern becomes a planned topic for the next editorial cycle.

04 Enablement update. The Sales Enablement Agent updates the standard battlecard, ROI calculator, and objection-handler library off the new signals.

05 Brief production. AEs request executive briefs 24-48 hours before pipeline meetings. The Executive Briefing Agent assembles each brief from CRM, recent transcripts, and account intelligence.

06 Sales feedback capture. After every meeting using agent-produced enablement, AEs rate the asset on relevance, accuracy, and usability. The rating feeds back to the Sales Enablement Agent.

07 Monthly alignment review. Marketing and sales leadership review what content was produced, what enablement was used, what signals were caught, and what was missed. The meeting is now 30 minutes, not 90, because the agents have already done the synthesis.

SUCCESS METRICS

- Signal-to-content cycle time. Target: < 30 days from an objection appearing in three calls to it being addressed in published content.
- Enablement asset utilization. Target: > 65% of agent-produced assets sent to a prospect within 7 days.
- AE asset rating. Target: > 4 out of 5 average rating.
- Executive brief usage rate. Target: > 80% of pipeline meetings have an agent-produced brief in advance.
- Battlecard freshness. Target: every battlecard updated within 30 days of the last field signal change.

COMMON FAILURE MODES

- Notes captured but signal not routed. Fix: the Meeting Notes Agent must pipe structured outputs to downstream agents automatically, not produce summaries that sit in a folder.
- Enablement assets shipped but never used. Fix: utilization is measured. Assets that go unused get retired. The utilization audit runs monthly.
- Monthly alignment meeting becomes performative. Fix: the agenda is restricted to three questions: what signals did we miss, what enablement did not get used and why, what is the highest-priority pattern for next month. Anything else is out of scope.

Pre-flight *checklist.*

The marketing-to-sales handoff is where most B2B companies leak the value marketing produced. The agents in this Field Note exist to close the leak. Resolve every item below before activating.

- [x] The ICP and persona work upstream is current. Enablement built on stale ICP and stale personas produces stale enablement.
- [x] The marketing-sourced and marketing-influenced pipeline definitions are documented and adopted by both functions.
- [x] AE adoption commitments are in place. The Sales Enablement Agent produces assets. If AEs do not use them, the agent stops being useful. Adoption discipline is a precondition.
- [x] The call-recording platform is connected and the AEs are using it. The Meeting Notes & Summarization Agent reads from call recordings.
- [x] The CRM is the canonical account record. Every other system reads from it. Without a canonical CRM, the Executive Briefing Agent and the Sales Enablement Agent produce conflicting briefs.
- [x] A defined set of asset types is named (battlecard, account deck, ROI calculator, executive briefing template, RFP language). The Sales Enablement Agent operates against this defined set, not freelance.
- [x] The event-marketing motion has a documented playbook. The Event Marketing Coordination Agent runs the playbook, not the other way around.

The 30-60-90 *Roadmap.*

The handoff-layer rollout sequence. Each agent in this Field Note operates at the marketing-to-revenue interface. Their value compounds when deployed together, but the order still matters.

DAYS 1-30

Stand up the Foundation. Audit the existing enablement library. Archive the generic. Identify the asset types AEs actually use. Document the event-marketing playbook. Connect the call-recording platform. Validate that the CRM is the canonical record. Deploy the Meeting Notes & Summarization Agent. It runs against the call-recording corpus from day one, producing structured outputs to the account record. Sales and marketing leadership review weekly summaries for accuracy.

DAYS 31-60

Deploy the Sales Enablement Agent on the top 30 active Tier 1 accounts. Produce per-account assets. Named AEs rate every asset on relevance, accuracy, and usability. Iterate until the ratings stabilize. Promote to broader coverage. Deploy the Executive Briefing Agent in advisory mode. AEs request briefs 24-48 hours before live executive meetings. Rate the briefs immediately after the meetings.

DAYS 61-90

Deploy the Event Marketing Coordination Agent. Run it on one Tier 1 event end-to-end. Compare against the previous-quarter's equivalent event. Deploy the RFP / Proposal Support Agent in co-drafting mode. The agent produces first drafts, the solutions writer revises, the agent learns from the revisions. After three RFP responses, the co-drafting cycle becomes the default workflow.

DAY 91+

Expand the Event Marketing Coordination Agent to default-mode for every event above the defined spend threshold. Promote the RFP / Proposal Support Agent to first-draft default for RFPs below the complexity threshold. Long-run the governance discipline: monthly utilization audit on the enablement library, quarterly brief-usefulness audit, quarterly RFP library hygiene.

Pitfalls & *remediation.*

PITFALL 01

Producing assets AEs do not use.

The Sales Enablement Agent's output volume is impressive. Volume without utilization is waste. The utilization audit must run from day one. Assets that go unused get retired without sentimentality. Templates that produce high meeting-conversion rates get reinforced.

PITFALL 02

Treating meeting notes as note-taking, not signal.

The Meeting Notes & Summarization Agent's real value is the signal it produces for the upstream marketing agents. Pain points become content topics. Objections become enablement updates. Competitor mentions become battlecard refreshes. Teams that use the agent only for AE convenience capture a tenth of the value.

PITFALL 03

Skipping the playbook for events.

The Event Marketing Coordination Agent runs the event playbook. Without a playbook, the agent invents one and it will not match how your team actually wins at events. Document the playbook first, deploy the agent second.

WHERE THIS FITS

Where this fits.

The five agents and one workflow in this playbook are the field and revenue slice of a larger marketing operating system. The handoff only works if everything upstream of it also works, the strategy, the demand, the content, the ops.

The flagship covers the full operating system end to end: **AI Agents for Marketing Teams**. Free, no email gate, made to forward.

— 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/ERMAdvisory

Instagram: instagram.com/ermadvisory

Advisory: erikrmiller.com/#contact

THE OPERATOR PLAYBOOK · VOL. 07

The seam between marketing and revenue is where most B2B companies leak the value.

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	erikrmiller.com/resources
NEWSLETTER	theoperator-erm.beehiiv.com
FACEBOOK	facebook.com/ERMAdvisory
INSTAGRAM	instagram.com/ermadvisory
ADVISE	erikrmiller.com/#contact

E.R.M. ADVISORY · NEW YORK CITY · 2026