

ICP & GTM STRATEGY

# AI Agents for ICP & GTM Strategy.

*The four agents that keep your ideal customer profile, personas, and account selection live, accurate, and defensible.*

WRITTEN BY

*Erik R. Miller*

SENIOR B2B MARKETING EXECUTIVE · FRACTIONAL CMO

4

AGENTS

0

WORKFLOWS

21

PAGES

AI AGENTS FOR ICP & GTM STRATEGY  
THE OPERATOR PLAYBOOK · VOL. 06

---

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.

# Strategy is upstream of everything.

The ICP is the single most leveraged document in a B2B marketing organization. Get it wrong and every dollar of demand, content, ABM, and sales enablement compounds the wrong way. Get it right and every dollar compounds correctly.

For most companies, the ICP is a static slide. Written in a strategy off-site. Updated once a year if at all. Disconnected from what is actually winning in the market this quarter. The result is predictable. Marketing chases the wrong segments. Sales loses faith. The CMO defends a strategy no one believes.

An agent-driven strategy layer changes this. The ICP becomes a live, versioned document that updates off closed-won and closed-lost data continuously. Personas refresh off actual call recordings, not surveys from 2023. Account selection grades every account against a current definition of fit, not a frozen one. Competitor intelligence flows into the strategy in real time.

This playbook covers the four agents that anchor the strategy layer. If you are doing ICP work, this is the document to build the next version of your strategy against.

## WHO THIS IS FOR

*For Heads of Marketing Strategy, GTM leaders, founders, and fractional CMOs building or rebuilding the ICP from scratch.*

## CONTENTS

# Inside this *Field Note*.

4 agents.

### THE AGENT ROSTER

<b>AGENT 01</b>	The ICP Research Agent .....
<b>AGENT 02</b>	The Buyer Persona Agent .....
<b>AGENT 03</b>	The Account Selection Agent .....
<b>AGENT 04</b>	The Competitor Intelligence Agent .....

---

### OPERATIONS & EXECUTION

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

---

### WHERE THIS FITS

–	Where this fits. ....
---	-----------------------

---

# The ICP & GTM Strategy. Agents.

AGENT 01

## The ICP Research Agent

**PURPOSE.** Maintains a live, multi-dimensional Ideal Customer Profile across firmographic, behavioral, technographic, and situational dimensions, and produces ICP fit scores for any account on demand. The ICP Research Agent is the canonical source of "who is the right customer" for every other agent in the system.

**INPUTS.** The current 4D ICP document; closed-won and closed-lost data with reasoning; ten or more recent customer interview transcripts; sales call notes; web behavior data where accessible; intent feeds; [an explicit warning that firmographic-only ICPs are insufficient](#).

**OUTPUTS.** Versioned ICP definitions; per-account ICP fit scores (A/B/C/D grade with per-dimension scores); quarterly ICP refresh report flagging shifts in fit drivers; a "what changed" memo whenever the underlying ICP definition is updated.

**MEMORY.** The last 50 accounts scored, accuracy of grades against eventual outcomes, ICP definition version history with reasoning for each shift, the customer-language phrasings that distinguish A-grade from B-grade fit.

**HUMAN OVERSIGHT.** The Head of Strategy approves any change to the ICP definition. Quarterly accuracy audit comparing predicted ICP grades against closed-won/closed-lost reality; anomalies trigger root-cause review.

**COMMON MISTAKES.** Over-weighting firmographic data because it is the easiest to source. Letting situational dimensions go stale (situational signals decay within 30 days). Treating the ICP as static when buyer reality shifts quarterly. The fix is in [the four-dimension ICP framework](#).

You are the ICP Research Agent for [Company]. You hold the canonical 4D ICP definition (Firmographic, Behavioral, Technographic, Situational). When asked to evaluate an account, score each dimension 0-10, return an overall A/B/C/D grade, and list the top three reasons for the grade. Rules: - Never grade A on firmographic data alone. A-grade requires evidence in at least three of the four dimensions. - If situational signals are older than 30 days, mark Situational as "stale" and adjust grade accordingly. - Use the customer-language library for fit-driver phrasings. Do not invent. - If data is insufficient to grade, return "insufficient data" with the missing inputs listed. Do not guess. - Output must be in the standard ICP Fit Brief template, not prose. Escalate to the Head of Strategy when: - An account scores A in three dimensions but D in the fourth – these are the highest-judgment-required accounts. - A grade contradicts the prior quarterly trend for the segment.

#### HOW TO DEPLOY

*The ICP Research Agent is the strategy layer. Deploy it first. Every other agent in your stack depends on the ICP grades this agent maintains.*

- 01** Audit your current ICP document. Most companies have a single-dimensional firmographic filter labeled ICP. Reject that as the starting point. The Foundation document this agent will use must define ICP across four dimensions, firmographic, behavioral, technographic, and situational, with the scoring weights explicit.
- 02** Pull twelve months of closed-won and closed-lost outcomes with reasoning attached. The agent has no judgment without outcome data. If outcome data does not exist in the CRM, the agent should refuse to grade until the data is supplied.
- 03** Run the agent in audit-only mode for two weeks. It scores accounts and produces ICP Fit Briefs, but no segmentation decisions are made off its output. Sales leadership reads the briefs and flags graded outputs that contradict field instinct.
- 04** Reconcile contradictions weekly with the Head of Strategy. Where the agent and the field disagree, one of them is missing context. Add that context to the Foundation document or to the agent's memory. Both are valid resolutions.
- 05** Promote to live grading once the ICP definition has stabilized for two consecutive weeks without revision. Quarterly accuracy audit becomes the long-running discipline.

## SUCCESS METRICS · WHAT TO MEASURE

<b>ICP grade accuracy</b>	Percent of A and B grades that close-won within the typical sales cycle. Target: A grade win rate > 30%, B grade win rate > 15%.
<b>Account-level fit consistency</b>	Same account graded the same way by the agent and by the Head of Strategy. Target: agreement > 80%.
<b>Definition stability</b>	Number of revisions to the ICP definition per quarter. Target: 0-2 revisions after first 90 days.
<b>Adoption depth</b>	Percent of new marketing-sourced accounts entering the pipeline with an ICP grade attached. Target: 100%.

## AGENT 02

# The Buyer Persona Agent

**PURPOSE.** Maintains buyer persona profiles inside ICP accounts and produces persona-level briefs that distinguish economic buyer, technical evaluator, end user, and procurement gatekeeper. Persona work and ICP work are different problems. [The three-layer targeting model](#) separates them deliberately.

**INPUTS.** Sales call recordings, customer interviews, win/loss notes, public LinkedIn signals, review-site verbatims, support ticket themes by role.

**OUTPUTS.** Persona briefs (one to two pages each), buying-committee maps per account tier, message-resonance scoring per persona, "what this persona reads / watches / trusts" intelligence.

**MEMORY.** Persona profile versions, the verbatim customer-language for each persona's pains and desired outcomes, recent feedback from sales on persona accuracy.

**HUMAN OVERSIGHT.** The Head of Content reviews persona briefs before they shape editorial direction. Sales calls pulled into persona memory are anonymized.

**COMMON MISTAKES.** Drifting toward generic personas without ICP context (an "Operations Leader" persona at a 50-person SaaS is not the same as one at a 5,000-person bank). Inventing motivations that have no customer-language evidence. Over-engineering buying committees beyond what the deal size actually warrants.

#### HOW TO DEPLOY

*Personas are downstream of ICP. Deploy this agent only after the ICP Research Agent has stabilized. Personas built on a moving ICP definition are wasted work.*

- 01** Catalogue every persona document the marketing organization currently has. Most of them will be three years stale. Mark them as archive, not source. The agent does not work from stale source.
- 02** Feed the agent fifteen to thirty recent sales call recordings, customer interviews, and win-loss notes from accounts inside the ICP definition. The persona profiles the agent produces must use verbatim customer language from these sources, not generic role descriptions.
- 03** Run the agent against three personas first, the most strategically important roles in the buying committee. Resist the temptation to map all twenty roles immediately. Three personas done well outperform ten personas done generically.
- 04** Validate every persona brief with two interviews of actual people in that role. The agent's output is a hypothesis until human validation closes the loop.
- 05** Publish the validated personas into the Foundation document. The agent maintains them on a quarterly cadence, refreshing the verbatim-language library and flagging when the underlying patterns shift.

## SUCCESS METRICS · WHAT TO MEASURE

<b>Verbatim-quote density</b>	Percent of persona brief content that traces to a specific customer interview or call recording. Target: > 60%.
<b>Message resonance lift</b>	Open and reply rates on messages built from agent-produced persona briefs vs. baseline. Target: 1.4x baseline within 60 days.
<b>Stakeholder confidence</b>	Sales leadership rating of the persona briefs as accurate descriptions of buyers they meet. Target: 4 out of 5 or higher.
<b>Refresh discipline</b>	Number of personas refreshed in the prior quarter against the agreed cadence. Target: 100%.

## AGENT 03

# The Account Selection Agent

**PURPOSE.** Builds and refreshes target account lists by tier, drawing on the ICP Research Agent's grades and the Intent Signal Agent's signal stack. Outputs the weekly Tier 1, Tier 2, and Tier 3 lists that drive the ABM motion. Most ABM programs fail because account selection is wrong; [this is where that failure starts](#).

**INPUTS.** ICP grades; first-party intent data; third-party intent feeds; funding, leadership, and hiring news; public-record signals; CRM activity from the last 90 days; current sales capacity by tier.

**OUTPUTS.** Tier-segmented account lists; weekly tier-movement report (accounts entering/leaving each tier with reasoning); a "watchlist" of accounts approaching tier-promotion thresholds.

**MEMORY.** Account-level history (when entered each tier, motion run, outcome); tier-promotion accuracy by signal type; recent disqualification reasoning.

**HUMAN OVERSIGHT.** ABM lead and sales leadership review the Tier 1 list weekly. Sales has veto power on Tier 1 entry to keep the list operationally honest.

**COMMON MISTAKES.** Over-rotating Tier 1 accounts week-to-week, exhausting sales attention. Letting tier-promotion be triggered by single-source intent (third-party intent alone is the noisiest predictor). Failing to cap Tier 1 at a number sales can actually run motion against.

#### HOW TO DEPLOY

*The Account Selection Agent operationalizes the ICP into tier-segmented account lists. It runs weekly and produces the only account list marketing and sales should be working off.*

- 01** Define the tier model before deploying. Tier 1, 2, 3, 4 with explicit criteria for each. Tier 1 typically means concentrated motion, named ABMer, custom content. The agent does not get to invent the tier definitions.
- 02** Connect the agent to the ICP Research Agent's grading output, first-party intent data sources, and at least one third-party intent feed. Without all three, the agent will over-rotate on whichever signal is available.
- 03** Produce the first list in audit mode. Compare against the list sales is currently working. Investigate every account where the agent and sales disagree. The disagreements are where the value is.
- 04** Hand the agent's Tier 1 list to the ABM lead and Sales VP for joint weekly review. Sales has veto power on Tier 1 promotions. This is not a policy debate, it is a deployment requirement.
- 05** Promote to live state once the joint review has produced two consecutive weeks of clean lists. Long-run the weekly tier-movement report as the standing artifact.

## SUCCESS METRICS · WHAT TO MEASURE

<b>Tier 1 win rate</b>	Percent of accounts the agent promoted to Tier 1 that closed-won within the typical sales cycle. Target: at least 2x average company win rate.
<b>Tier accuracy at promotion time</b>	Tier-promotion accuracy measured against eventual outcome. Target: 70%+ for Tier 1, 50%+ for Tier 2.
<b>Sales adoption</b>	Percent of sales-pursued accounts that appear on the agent's current tier list. Target: > 85%.
<b>Veto rate</b>	Percent of agent-promoted Tier 1 accounts vetoed by sales each week. Target: < 15%. A higher rate means the ICP or the tier criteria are off.

## AGENT 04

# The Competitor Intelligence Agent

**PURPOSE.** Tracks competitor positioning, product, pricing, content, and field motion. Produces battlecards, threat alerts, and quarterly competitive briefs. The agent's most important rule is to never invent. Hallucinated competitor claims are the fastest way to destroy sales trust in the system.

**INPUTS.** Competitor websites, public earnings, news, hiring signals, review-site verbatims, win/loss interviews, sales-team field notes, social media positioning shifts.

**OUTPUTS.** Per-competitor positioning briefs; quarterly battlecard refresh; threat alerts (new product launches, pricing moves, leadership changes); message-overlap analysis showing where the brand and a competitor are saying the same thing.

**MEMORY.** Competitor positioning over time; verified vs. unverified claims (with sources); historical accuracy of threat alerts.

**HUMAN OVERSIGHT.** Sales enablement reviews battlecards before distribution. Any competitor claim must include a verifiable source citation. Unsourced claims are flagged, never published.

**COMMON MISTAKES.** Hallucinating competitor features, pricing, or customer wins. Treating press coverage as truth. Over-indexing on one loud competitor and missing the quiet one taking share.

#### HOW TO DEPLOY

*The Competitor Intelligence Agent has one rule above all others. It does not invent. Hallucinated competitor claims destroy sales trust faster than any other agent failure. Deploy this agent with the most cautious oversight model in the stack.*

- 01** Define the competitor set explicitly. Top three direct, two adjacent, one emerging. The agent works only against this defined set. Anything outside the set requires a deliberate scope change.
- 02** Connect the agent to public sources only at first. Competitor websites, public earnings, news, hiring signals, review-site verbatims. Do not feed the agent unsourced field rumor as input.
- 03** Establish the never-invent rule in the agent's system prompt. If a fact cannot be cited to a source, the agent must either flag missing data or refuse to write. Test the rule before promoting the agent to production.
- 04** Produce the first quarterly battlecard refresh in shadow mode. The Head of Strategy compares the agent's output against the existing sales-trusted positioning. Document every divergence.
- 05** Promote to live state. The agent now produces weekly threat alerts and quarterly battlecard refreshes. The Head of Strategy approves any quarterly narrative shifts before publication.

## SUCCESS METRICS · WHAT TO MEASURE

<b>Citation rate</b>	Percent of competitor claims in agent output that include a verifiable source link. Target: 100%.
<b>Hallucination incidents</b>	Number of unsourced or fabricated competitor claims found in any agent output. Target: 0. This is a kill criterion.
<b>Threat-alert lead time</b>	Median time from a competitor move (launch, pricing, leadership change) being publicly visible to the agent flagging it. Target: < 5 business days.
<b>Sales adoption</b>	Percent of active sales reps who reference the current battlecard in their last five calls. Target: > 70%.

# Pre-flight *checklist.*

---

*Strategy is upstream of everything. The agents in this Field Note produce the documents every other agent reads from. Resolve every item below before activating.*

- Twelve months of closed-won and closed-lost data exists with reasoning attached. Without outcome data, the ICP Research Agent has no signal to grade against.

---

- Sales call recordings are accessible and tagged. The Buyer Persona Agent requires real buyer language. Recordings are the source.

---

- Customer interviews are scheduled on a regular cadence (quarterly minimum). Persona briefs decay without ongoing voice-of-customer input.

---

- The competitor set is named explicitly. Top three direct, two adjacent, one emerging. The Competitor Intelligence Agent works only against the defined set.

---

- The strategy review cadence is established. Monthly or quarterly leadership reviews where the ICP and persona work gets pressure-tested.

---

- Cross-functional adoption commitments are in place. The strategy work fails if marketing produces it and sales ignores it. Sales leadership must commit to using the ICP and personas the agents produce.

---

- The Head of Strategy or equivalent role is named. The ICP Research Agent operates with human approval on definition changes. Without a named approver, the discipline collapses.

# The 30-60-90 *Roadmap.*

---

*The strategy-layer rollout sequence. The four agents in this Field Note must be deployed in the order below. Each enables the next.*

## **DAYS 1-30**

Stand up the Foundation. Audit the existing ICP, personas, and competitor positioning. Build the four-dimensional ICP definition if it does not exist. Identify the call-recording library, the customer-interview cadence, and the competitor source set. Deploy the ICP Research Agent in audit-only mode. The agent scores accounts and produces ICP Fit Briefs. No segmentation decisions are made yet off the agent's output. Sales leadership reads the briefs and flags graded outputs that contradict field instinct.

## **DAYS 31-60**

Promote the ICP Research Agent to live grading. The definition has stabilized over two weeks without revision. Deploy the Buyer Persona Agent. Feed it the recent call recordings, interviews, and win-loss notes. Build the first three personas (the strategically important roles in the buying committee). Validate each persona brief with two interviews. Publish the validated personas to the Foundation.

## **DAYS 61-90**

Deploy the Account Selection Agent and the Competitor Intelligence Agent. The Account Selection Agent works off the ICP grades and the tier model. The Competitor Intelligence Agent operates under the never-invent rule, producing the first quarterly battlecard refresh in shadow mode. Sales leadership reviews. Promote to live state once the battlecard refresh has been validated.

## **DAY 91+**

Expand persona coverage to additional roles. Long-run the strategy-layer discipline: quarterly ICP accuracy audit, quarterly persona refresh, quarterly battlecard refresh, weekly account-selection joint review.

# Pitfalls & *remediation.*

---

## PITFALL 01

### *Promoting the ICP Research Agent to live grading too early.*

The most common failure pattern. Teams ship the agent before the four-dimensional definition has stabilized. The agent grades against a moving target. Every downstream agent inherits the instability. Wait for two weeks of definition stability before promotion.

## PITFALL 02

### *Building personas without verbatim language.*

The Buyer Persona Agent produces real personas only if it reads from real buyer conversations. Teams that feed it survey data or demographic profiles get back personas that read like every other persona document. Real conversations only.

## PITFALL 03

### *Treating the Competitor Intelligence Agent as a content-generation tool.*

It is not. It is a research and synthesis agent under a never-invent rule. The moment teams use it to produce competitive content without source citations, hallucinations enter the field, sales trust collapses, and the program ends.

## WHERE THIS FITS

# Where this fits.

The four agents in this playbook are the strategy layer of a larger marketing operating system. Strategy is upstream. The demand, content, ABM, and ops functions all consume what these four agents produce.

For the complete picture, read the flagship playbook: **AI Agents for Marketing Teams**. Free, no signup, forwardable. If you want a deeper look at the ICP framework itself, see the post [ICP Development: The Framework I Use Every Time](#).

— Erik R. Miller

## Erik R. Miller

SENIOR B2B MARKETING EXECUTIVE · FRACTIONAL CMO

**Web:** [erikmiller.com](http://erikmiller.com)

**Newsletter:** The Operator, [theoperator-erm.beehiiv.com](http://theoperator-erm.beehiiv.com)

**LinkedIn:** [linkedin.com/in/erikrm](https://linkedin.com/in/erikrm)

**Facebook:** [facebook.com/ERMAdvisory](https://facebook.com/ERMAdvisory)

**Instagram:** [instagram.com/ermadvisory](https://instagram.com/ermadvisory)

**Advisory:** [erikmiller.com/#contact](http://erikmiller.com/#contact)

# Strategy is *upstream* of everything. The ICP is the *upstream* document.

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](http://erikmiller.com/resources)
- NEWSLETTER** [theoperator-erm.beehiiv.com](http://theoperator-erm.beehiiv.com)
- FACEBOOK** [facebook.com/ERMAdvisory](https://facebook.com/ERMAdvisory)
- INSTAGRAM** [instagram.com/ermadvisory](https://instagram.com/ermadvisory)
- ADVISE** [erikmiller.com/#contact](http://erikmiller.com/#contact)