

AI Adoption in the Revenue Org

A Practical Manager & Team Playbook

Most AI rollouts fail not because of the technology, but because of the adoption strategy. Teams get licenses. Managers get dashboards showing seat utilization. Nobody measures whether the AI is actually changing how well anyone does their job.

This playbook is different. For every major revenue role -- AEs, SEs, BDRs, CSMs, and Support -- you will find specific workflows with step-by-step instructions, copy-paste prompt templates you can use today, manager coaching questions for 1:1s, and clear metrics to prove adoption is working. Read it once to understand the framework. Return to it during rollout.

WHAT YOU WILL FIND IN THIS PLAYBOOK

01 Why Rollouts Fail & The Fix

The 3 root causes of failed AI adoption and the workflow-first framework that solves them.

02 Sequencing & Governance

Who to start with, how to phase the rollout, and clear rules on what AI should never do autonomously.

03 Role-by-Role Playbooks

Step-by-step workflows, copy-paste prompts, and manager coaching questions for AEs, SEs, BDRs, CSMs, and Support.

04 30-Day Sprint & Measurement

A week-by-week adoption sprint with coaching checkpoints and the metrics that prove it's working.

01 WHY MOST AI ROLLOUTS FAIL -- AND WHAT TO DO INSTEAD

High-performing sales teams are 4.9x more likely to use AI than underperforming teams (Salesforce State of Sales, 2023). Yet the majority of AI rollouts across revenue organizations produce minimal measurable impact within 90 days of launch. The reasons are consistent.

Root Cause 1: Starting with technology, not workflows

Most rollouts begin with tool procurement: "We bought Gong. Now what?" Without tying AI to a specific, high-friction task a rep does every day, adoption stalls. Reps don't adopt tools -- they adopt habits. The workflow has to be defined before the tool goes live.

Root Cause 2: No manager accountability for adoption

License utilization is not adoption. Adoption means a rep changed how they do a specific task because of AI -- and the quality or speed of the output improved. Managers need to actively review AI outputs in 1:1s, not just ask "are you using it?" without examining what was produced.

Root Cause 3: Measuring the wrong things

Logins and prompts sent tell you nothing. The right measures are task-specific: time from call end to follow-up sent, improvement in email reply rates, QBR prep time. Define one metric per role before rollout. If you can't name the metric, you are not ready to launch.

THE FIX: WORKFLOW -> HABIT -> MEASURE

For every role, define: (1) the specific task the AI supports, (2) the habit the rep builds around using and editing the output, and (3) the metric that proves the habit is working. All three must be in place before launch. The table below shows the pattern using AEs as an example.

EXAMPLE: WORKFLOW-HABIT-MEASURE FOR ACCOUNT EXECUTIVES

Workflow	Post-call follow-up email
AI Action	Generate first-draft follow-up from Gong call summary
Rep Habit	Review draft, personalize 2-3 lines, send within 2 hours of call end
Manager Action	Review one AI-assisted follow-up per rep per week in 1:1; check for personalization
Metric	Time from call end to follow-up sent (baseline vs. post-AI target: under 2 hours)
Red Flag	Reps sending raw AI output without edits; follow-ups missing call-specific details

02 SEQUENCING THE ROLLOUT & GOVERNANCE GUARDRAILS

Sequence the rollout by role friction and impact. Start with the roles that have the highest volume of repetitive tasks and the lowest AI skepticism. Build visible early wins before expanding to more skeptical roles.

RECOMMENDED ROLLOUT SEQUENCE

Phase 1 -- Weeks 1-2	BDRs and Support. Highest volume of repetitive tasks. Fastest time-to-value. Use wins here to build the business case for AE and CSM rollout.
Phase 2 -- Weeks 3-6	CSMs and SEs. High administrative burden (QBR prep, RFPs). Receptive once BDR wins

are visible. Require workflow-specific prompt training.

Phase 3 -- Weeks 7-12

AEs. Most skeptical audience. Roll out only after the prompt library is established and manager review cadence is in place. Lead with time savings, not performance monitoring.

GOVERNANCE: WHAT AI CAN AND CANNOT DO

AI CAN do (with human review before use)

Research briefs ? First-draft emails and follow-ups ? QBR narrative drafts ? Call summaries ? RFP first-pass responses ? Deal risk summaries ? Knowledge base article drafts ? Ticket classification suggestions

AI MUST NOT do without explicit manager sign-off

Send any external communication autonomously ? Make pricing commitments ? Represent legal or compliance positions ? Modify CRM records without human review ? Respond to customer escalations ? Generate content for external publication

Data handling: what goes in, what does not

Safe: deal summaries, account notes, call transcripts (in compliant tools), product documentation. Never include: customer PII beyond name/company, unreleased roadmap details, proprietary pricing, legal agreements, or personal employee data.

03 ACCOUNT EXECUTIVES

AEs are the most skeptical audience for AI and the most valuable one to win over. Lead with time savings on administrative work -- never frame AI as a performance monitoring tool. The highest-ROI workflows are pre-call research, post-call follow-up, and deal review prep.

WF 1 Pre-Call Research Brief

STEPS

1. Open the deal record in CRM. Copy: stage, ARR, known pain, champion name, notes from prior calls.
2. Find 1-2 recent company signals: news, earnings, LinkedIn activity, job postings, or executive changes.
3. Run the prompt below. Review the output -- edit any pain assumptions that are wrong before the call.
4. Use the brief during discovery to validate each assumption in the first 10 minutes.

COPY-PASTE PROMPT -- PRE-CALL RESEARCH BRIEF

I'm calling [Contact Name], [Title] at [Company] on [Date]. They are a [headcount] [industry] company. Deal context: Stage [X], ARR [\$X], close date [X]. Known pain from CRM notes: [paste]. Recent signal: [paste headline or LinkedIn trigger]. My product addresses [specific pain area]. Generate: (1) 3 pain assumptions to validate in discovery, each framed as a question. (2) 2 industry trends relevant to their business right now. (3) 1 compelling opening question that references something specific about them.

IMPACT: Saves 20-30 min of manual research per call. Reps report more focused discovery and higher confidence opening calls cold.

WF 2 Post-Call Follow-Up Email

STEPS

1. Immediately after the call, pull the AI-generated summary from your call intelligence tool (Gong, Chorus, etc.).
2. Run the prompt below, pasting the call summary and the agreed action items.
3. Read the draft. Personalize at minimum: the opening line, one reference to a specific thing they said, and the CTA.
4. Send within 2 hours of the call. Never send raw AI output -- always edit for your voice.

COPY-PASTE PROMPT -- POST-CALL FOLLOW-UP

I just had a [discovery/demo/negotiation] call with [Name] at [Company]. Call notes: [paste summary from call tool]. Action items agreed: [numbered list with owners and dates]. Write a follow-up email that: opens with a specific reference to something they said, confirms the agreed action items with owners and dates, adds one relevant piece of value (a stat, case study reference, or resource), and closes with a clear single next step. Under 200 words. Professional but direct tone.

IMPACT: 15-20 min saved per meeting. Deals where follow-up is sent within 2 hours of the call have 21% higher progression rates to the next stage.

WF 3 Deal Review Prep -- MEDDPICC Gap Analysis

STEPS

1. Before each pipeline review, pull deal details: stage, ARR, close date, champion, and current MEDDPICC status.
2. Run the prompt below. The output is your pre-read for the manager conversation, not a replacement for it.
3. Bring the AI risk summary to the pipeline review. Your manager should challenge the analysis, not accept it.
4. Log the agreed action in CRM immediately after the review so next week's AI summary reflects the update.

COPY-PASTE PROMPT -- MEDDPICC DEAL RISK SUMMARY

Deal: [Company]. ARR: [\$X]. Stage: [X]. Close date: [X]. Champion: [Name / title]. MEDDPICC -- Metrics: [confirmed/unconfirmed/unknown]. Economic Buyer: [identified/not identified/accessed]. Decision Criteria: [documented/not documented]. Decision Process: [mapped/not mapped]. Paper Process: [known/unknown]. Pain: [articulated/vague]. Competition: [identified/unknown]. Identify: (1) The top 3 risks to this deal closing on the stated date. (2) The single most important action to take this week for each risk. (3) The one MEDDPICC gap most likely to cause a push or loss if left unaddressed.

IMPACT: Reps arrive at pipeline reviews with specific risk answers, not status updates. Reviews run 30-40% shorter when this discipline is in place.

MANAGER COACHING QUESTIONS FOR AE 1:1S

- ? Show me one AI-generated follow-up from this week. What did you change before sending it? Why?
- ? What was your pre-call brief for your highest-priority call? Did it change how you opened discovery?
- ? Which deal did the AI flag as highest risk this week? Do you agree? What are you doing about it?
- ? Are there workflows where you are still prepping manually because the AI output isn't good enough? What's wrong with it?

GREEN FLAGS -- Adoption is working

- ? Reps editing AI drafts before sending -- not using raw output
- ? Pre-call briefs reference specific news or signals, not generic pain
- ? Pipeline reviews reference AI risk summaries with rep commentary added
- ? Prompts improving week-over-week as reps iterate on what works

RED FLAGS -- Intervention needed

- ? Follow-ups sent with no personalization beyond name and company
- ? Reps skipping AI research because "I already know this account"
- ? Deal risk summaries accepted in pipeline review without challenge
- ? AI used once, abandoned -- "it wasn't good enough" with no iteration

04 SALES ENGINEERS

SEs have the most to gain from AI in the revenue org. Technical documentation is time-consuming to customize; demo scripts are repetitive with minor variations; RFP responses are built from the same knowledge base every time. The key is building SE-specific prompt libraries trained on your product documentation and competitive intel.

WF 1 Demo Personalization Script

STEPS

1. Pull prospect data from CRM: company size, industry, tech stack, stated pain, deal stage.
2. Identify the 1-2 product flows most relevant to their pain -- not the full standard demo.
3. Run the prompt below to generate a customized opening and transition into the live demo.
4. Review for accuracy. Edit any product claims before using it in front of a prospect.

COPY-PASTE PROMPT -- DEMO PERSONALIZATION

My prospect: [Company], [headcount], [industry]. Their stated pain: [X]. Tech stack relevant to us: [Y]. Deal stage: [X]. Product area I'm demoing: [feature/flow]. Write a 2-minute demo opening that: (1) Opens with their specific pain in their industry's language -- not generic vendor speak. (2) Quantifies the problem using their numbers or their industry's benchmarks where possible. (3) Sets up the live demo with a "what if you could..." bridge. (4) Ends with a question that confirms we are solving the right problem before proceeding.

IMPACT: Reps report 30-40% less demo prep time. Personalized openings reduce "just take me through your standard demo" redirects in the first 5 minutes.

WF 2 RFP Response Drafting

STEPS

1. Read the RFP question carefully. What are they actually evaluating? Is scoring criteria stated?
2. Pull the most relevant section from your product knowledge base or prior RFP library.
3. Run the prompt below. Flag any claims in the output with [CHECK] before submission.
4. Verify every technical claim -- AI will sometimes hallucinate feature capabilities. The SE owns accuracy.

COPY-PASTE PROMPT -- RFP RESPONSE

RFP question: [paste the full question]. Our product's relevant capabilities: [paste from product docs or knowledge base]. Key differentiators vs. [competitor if named]: [brief list]. Write a response that: directly answers the question in the first sentence, cites specific capabilities with accurate technical detail, differentiates from standard vendor boilerplate, and stays under 250 words. Mark any claim I should verify before submission with [CHECK].

IMPACT: RFP response time per question drops from 45-60 min to 15-20 min. SEs spend time reviewing and refining, not writing from scratch.

WF 3 Technical Objection Handling

STEPS

1. Note the exact wording of the objection during or immediately after the call.
2. Identify: is this a real technical blocker, a stated preference, or competitive FUD?
3. Run the prompt below to get a structured response framework.
4. Never read AI output verbatim on a call. Use it to prepare your response, then deliver in your own words.

COPY-PASTE PROMPT -- TECHNICAL OBJECTION RESPONSE

Technical objection: "[exact wording]." Context: [company type, deal stage, who said it -- SE, IT, Security, etc.]. Competing solution referenced (if any): [X]. Our relevant capability: [brief technical description]. Give me: (1) A direct response that acknowledges the concern without being defensive. (2) A clarifying question that reveals whether this is a real blocker or a preference. (3) A proof point or reference architecture I can offer. (4) An escalation path if this requires a deeper technical conversation.

IMPACT: Reduces escalation to technical leadership for Tier 1 objections. Consistent response quality across the team, especially for newer SEs.

MANAGER COACHING QUESTIONS FOR SE 1:1S

- ? Show me the last RFP section AI generated. What was technically wrong? How did you correct it?
- ? What objection did you handle with an AI-prepared response this week? Did it land? What would you change about the prompt?
- ? Is the team's prompt library growing? Who is contributing? Who isn't, and why?
- ? Where is AI still not useful for your role? What workflow would make the biggest difference if we built a better prompt?

GREEN FLAGS -- Adoption is working

- ? SE prompt library actively growing and shared across the team
- ? RFP turnaround time measurably faster without quality decrease
- ? Technical claims in AI output being fact-checked before submission
- ? SEs iterating on prompts to get better outputs -- version 2, 3, 4

RED FLAGS -- Intervention needed

- ? AI RFP responses submitted without accuracy review -- liability risk
- ? Demo openings that are generic despite running the customization prompt
- ? No shared prompt library -- every SE reinventing the wheel
- ? SEs using AI for one workflow and ignoring the rest

05 BUSINESS DEVELOPMENT REPS

For BDRs, AI is a research and personalization multiplier. The biggest risk is AI-generated outreach that reads like AI-generated outreach -- train the team to use AI for research and structure, then add a human voice before sending. Every email should have at least one line no AI could have written without the rep reviewing it first.

WF 1 Account Research Brief

STEPS

1. Choose your target. Spend 3 minutes gathering: company size, industry, recent news, LinkedIn signals (job posts, exec activity, hiring), and their tech stack if visible.
2. Run the prompt below. The output is a starting point -- add anything you know that the AI doesn't.
3. Extract one specific personalization hook from the brief -- the one thing you will reference in outreach.
4. Do not send outreach without a real hook. "I noticed you're in [industry]" does not count.

COPY-PASTE PROMPT -- ACCOUNT RESEARCH BRIEF

Target: [Company]. Industry: [X]. Size: [X employees / \$X revenue]. Recent signal: [paste headline or LinkedIn trigger]. My product helps [specific pain area] for companies like this. Generate a 5-point account brief: (1) Most likely pain point based on the signals, with reasoning. (2) Best news or trigger hook to open outreach with. (3) Likely buying persona and their top concern. (4) One specific personalization line I can use in a cold email. (5) Best opening question for a cold call based on the research.

IMPACT: Research time per account drops from 15-20 min to 4-5 min. Specific personalization hooks are the #1 driver of cold outreach reply rates.

WF 2 Email Quality Review Before Sending

STEPS

1. Write your first draft of the outreach email as you normally would.
2. Run the quality check prompt below before adding it to a sequence.
3. Make the 2 suggested edits. Do not ask AI to rewrite the whole email -- that removes your voice.
4. If specificity scores below 7/10, do not send. Go back to the account brief and find a better hook.

COPY-PASTE PROMPT -- EMAIL QUALITY CHECK

Review this cold outreach email and score it: (1) Specificity (1-10): How specific is this to this company or person? Does it reference something real? (2) Value clarity (1-10): Is it immediately obvious why I'm reaching out and what problem I solve? (3) CTA friction (1-10): How easy is it to say yes to the ask? Give me exactly 2 specific edits -- not a full rewrite, just targeted improvements. Email: [paste your draft]

IMPACT: Open rates improve 15-25% when BDRs use AI review before sending. Reply rates improve when specificity consistently scores 8+ before send.

WF 3 Cold Call Opener Preparation

STEPS

1. Have your account brief ready before generating openers.
2. Run the prompt below. Get three options -- different angles, not variations of the same line.
3. Pick one before the call. Decide in advance, not in the moment -- hesitation kills cold call openings.
4. After the call, note which opener worked and why. Add it to the team's call opener library.

COPY-PASTE PROMPT -- COLD CALL OPENERS

I'm cold calling [Name], [Title] at [Company]. Account brief: [paste key points]. My product helps [persona like this] with [specific pain]. Generate 3 opening lines -- under 20 words each, each taking a different angle: Option 1: Lead with the news or trigger hook. Option 2: Lead with a pain assumption that references something specific about their role or company. Option 3: Lead with a direct value statement that names a concrete outcome. Do not use "I'm just reaching out" or any filler phrase.

IMPACT: Teams using prepared openers report 20-30% improvement in conversation rate -- calls that become real conversations rather than immediate hang-ups.

MANAGER COACHING QUESTIONS FOR BDR 1:1S

- ? What was the AI quality score on your best-performing email this week? What did you change after the review?
- ? Show me the account brief for the account you spent the most time on. What personalization hook did you use in the outreach?
- ? Which cold call opener landed best this week? Was it AI-suggested or your own? What made it work?
- ? Are you skipping the quality check on any emails before sending? Why?

GREEN FLAGS -- Adoption is working

- ? Account briefs reference specific signals, not industry generics
- ? Email specificity scores trending 8+ before send
- ? BDRs adding working openers to a shared team call library
- ? Reps asking for new prompt types -- sign of genuine engagement

RED FLAGS -- Intervention needed

- ? Account briefs that look identical across accounts -- prompt not customized
- ? Emails with quality score below 6 still entering sequences
- ? Reps skipping the review step because "it takes too long"
- ? Cold call openers used verbatim without adapting to actual conversation

06 CUSTOMER SUCCESS MANAGERS

CS is the most natural home for AI in the revenue org. CSMs are buried in administrative work: QBR prep, health score monitoring, renewal drafting, escalation documentation. AI gives time back without changing the relationship quality that drives retention. The goal is not to automate the CSM relationship -- it is to free up prep time so CSMs spend more time in actual conversations.

WF 1 QBR Preparation

STEPS

1. Three days before the QBR, pull the full account data: health score, usage metrics, support ticket history, goals from last QBR, renewal date, and any expansion signals.
2. Run the prompt below. The output is a first-draft QBR narrative -- not a finished deck.
3. Review the draft against your knowledge of the account. Add anything relationship-specific that the data doesn't capture.
4. Use the AI output for the exec summary slide and talking points. Spend prep time on the conversation, not the slides.

COPY-PASTE PROMPT -- QBR PREPARATION

QBR prep for [Customer]. ARR: [\$X]. Health score: [X/10]. Renewal date: [X]. Usage this quarter: [logins, active users, feature adoption %, vs. prior quarter]. Goals from last QBR: [list]. Goals achieved: [list]. Goals not achieved: [list + why]. Open support tickets or escalations: [list if any]. Generate: (1) A 3-sentence executive summary of value delivered this quarter. (2) 3 talking points for the next-quarter goals discussion. (3) 2 expansion opportunities to raise based on the usage data. (4) 1 risk or concern to address proactively before they bring it up.

IMPACT: QBR prep time drops from 4-6 hours to 60-90 minutes. CSMs arrive better prepared for the conversation rather than the slide-building.

WF 2 Weekly Health Narrative for CS Review

STEPS

1. Before the weekly CS review, identify your accounts flagged yellow or red by health score.
2. For each at-risk account, run the prompt below.
3. Bring the output to the review -- 3 sentences per account, flag color, and one recommended action.
4. Manager challenges the AI classification. CSM adds relationship context the data doesn't have.

COPY-PASTE PROMPT -- ACCOUNT HEALTH NARRATIVE

Weekly health report for [Customer]. Usage this week: [logins / active users / feature adoption % vs. prior week]. Support tickets this week: [count and overall sentiment -- frustrated / neutral / positive]. Last CSM contact: [date + brief outcome]. Health trend -- last 4 weeks: [improving / declining / flat]. NPS or CSAT if recent: [score]. Write a 3-sentence plain-English health narrative for the CS review. Classify: Green (on track), Yellow (watch closely), or Red (intervene now). Recommend one specific action for the CSM to take this week.

IMPACT: CS review meetings run 25-35% faster with pre-written health narratives. Intervention decisions are made on data, not memory.

WF 3 Renewal & Expansion Outreach

STEPS

1. Flag accounts 90 days before renewal. Pull usage trends, health score, value delivered, and expansion signals.
2. Run the prompt below. The draft is a starting point -- not a ready-to-send email.
3. Add one specific reference to a conversation you had, a result they mentioned, or a goal they set. This line cannot come from AI.
4. Send at 90 days out, not 30. Late outreach is the #1 preventable cause of renewal risk.

COPY-PASTE PROMPT -- RENEWAL OUTREACH

Renewal outreach for [Customer]. Renewal date: [X]. Current ARR: [\$X]. Value delivered this year: [2-3 specific outcomes -- outcomes, not features]. Health score trend: [improving / stable / declining]. Expansion opportunity: [usage pattern suggests they need X]. Write a renewal email that: opens with one specific outcome they achieved (not "it has been a great year"), surfaces the expansion opportunity as a natural next step based on their usage, makes the renewal ask feel like a logical continuation, not a transaction. Under 180 words. Leave a [PERSONALIZE HERE] placeholder for a relationship-specific line.

IMPACT: Renewal outreach sent at 90 days vs. 30 days correlates with 18% higher renewal rate. CSMs using AI drafts send earlier because the prep barrier is lower.

MANAGER COACHING QUESTIONS FOR CSM 1:1S

- ? What health narrative did AI generate for your riskiest account? Do you agree with the classification? What does the data miss?
- ? How much did you change the QBR exec summary from the AI draft? What was wrong with the original?
- ? Which account is 90+ days from renewal? Has outreach been sent? If not, why not?
- ? Where is AI saving you the most time this week? Where is it still not useful?

GREEN FLAGS -- Adoption is working

- ? QBR prep time measurably shorter without quality decrease
- ? Renewal outreach sent at 90 days, not 30
- ? CSMs adding relationship context to AI health narratives before review
- ? CS reviews focused on action, not status reporting

RED FLAGS -- Intervention needed

- ? QBR decks built entirely from AI output without CSM review -- accuracy risk
- ? Renewal outreach still sent at 30 days or less -- habit not changing
- ? Health narratives accepted without manager challenge
- ? AI used only for QBR prep and ignored for weekly workflow

07 SUPPORT TEAMS

Support is often the last team in AI rollouts and the first to feel the impact. AI handles Tier 1 deflection, generates resolution summaries, and routes escalations. The critical governance question -- when does AI respond autonomously vs. hand off to a human -- must be defined and documented before launch. Never automate responses to P1/P2 issues or customer escalations.

WF 1 Resolution Summary for Knowledge Base

STEPS

1. When a ticket is resolved, before closing it, paste the issue description and resolution steps into the prompt below.
2. Review the draft KB article for accuracy. Support agents own the accuracy of anything published -- not the AI.
3. Have a team lead spot-check the first 10 AI-generated articles before autopilot publishing begins.
4. Build a monthly KB review cycle to flag articles where the resolution has changed or become outdated.

COPY-PASTE PROMPT -- KB ARTICLE GENERATION

Ticket issue: [describe the problem in the customer's exact words, not internal language]. Root cause: [technical explanation for internal reference]. Resolution steps: [numbered list of what support did]. Affected product area: [X]. Generate a knowledge base article that: (1) Titles the issue in plain customer language (not "Error code XYZ"). (2) Describes the problem in the first sentence so customers confirm it's their issue. (3) Lists resolution steps a customer could follow themselves. (4) Ends with a "When to contact support" section. Mark any step requiring support access with [SUPPORT ONLY].

IMPACT: KB article creation drops from 25-35 min to 6-8 min per article. Teams that document every resolution reduce repeat ticket volume 20-30% within 90 days.

WF 2 Ticket Classification and Routing

STEPS

1. Configure AI classification to run on every incoming ticket before it enters the queue.
2. Human reviewer confirms routing before assignment -- especially for P1/P2 and enterprise accounts.
3. Track misclassification rate weekly. Update the prompt and category definitions when accuracy drops below 85%.
4. Never route P1 tickets or enterprise escalations on AI classification alone.

COPY-PASTE PROMPT -- TICKET CLASSIFICATION

Incoming ticket: [paste full ticket body]. Customer: [name]. Account tier: [enterprise / mid-market / SMB]. Classify this ticket: (1) Category: [billing / technical bug / how-to / feature request / account management / other]. (2) Priority: [P1 -- system down / P2 -- major feature broken / P3 -- degraded / P4 -- inquiry]. State your reasoning in one sentence. (3) Routing: which team or queue should handle this, and why. (4) If P1 or P2: what is the likely business impact on the customer right now? This classification will be reviewed by a human before routing.

IMPACT: Misrouted tickets drop 40-60% within 60 days. P1/P2 escalation response time improves when priority is flagged accurately at intake.

WF 3 Complex Ticket Response Drafting

STEPS

1. For tickets requiring non-standard responses (technical explanation, workaround, sensitive account), pull customer history first.
2. Run the prompt below. Do not send AI-generated responses without editing -- especially to enterprise accounts.
3. Edit for tone: AI tends toward formal. Adjust to match your support voice guidelines.
4. P1/P2 and enterprise escalations require manager review before sending.

COPY-PASTE PROMPT -- TICKET RESPONSE DRAFT

Ticket from [Customer]. Account tier: [enterprise / mid-market / SMB]. Issue: [describe the problem]. What they've already tried: [list if mentioned]. Relevant product behavior: [technical context for the agent]. Customer sentiment: [frustrated / neutral / calm / urgent]. Draft a response that: (1) Opens by acknowledging their specific issue -- not a generic greeting. (2) Explains root cause in plain language without deflecting. (3) Provides clear resolution steps or a committed next step with a timeline. (4) Closes with what happens next and who owns it. Flag any technical claim to verify with [CHECK].

IMPACT: First-response quality improves significantly. Repeat contacts on the same issue decrease when root cause explanations are clear and complete.

MANAGER COACHING QUESTIONS FOR SUPPORT 1:1S

- ? Show me the last KB article AI generated. What did it get wrong? What did you have to correct?
- ? Which ticket did AI misclassify this week? What was the routing impact? Do we need to update category definitions?
- ? Are agents editing AI-drafted responses before sending, or copy-pasting? How do you know?
- ? What ticket type is AI still not useful for? Should we build a better prompt for it?

GREEN FLAGS -- Adoption is working

- ? KB article backlog shrinking -- every resolved ticket producing an article
- ? Ticket misclassification rate below 15% and trending down
- ? Agents editing AI responses for tone and accuracy before sending
- ? Repeat ticket volume on documented issues decreasing month-over-month

RED FLAGS -- Intervention needed

- ? AI responses sent verbatim without review -- accuracy and tone risk
- ? KB articles not published because "they're not good enough" with no iteration
- ? P1/P2 tickets routed on AI classification without human confirmation
- ? Misclassification rate above 15% with no prompt review in progress

08 30-DAY ADOPTION SPRINT

Run this sprint per team after the 90-minute onboarding workshop. Each week has a specific focus, concrete tasks, and a coaching checkpoint. The goal is to build one habit per rep in 30 days -- not to deploy every workflow at once. Breadth kills adoption.

Week 1 -- Install & First Use

- ? Every rep uses the AI tool at least once per day on a low-stakes task (research brief, email draft).
- ? No quality expectations this week. The goal is habit formation, not output quality.
- ? Manager reviews outputs in 1:1: "Show me one thing AI helped you with this week."
- ? Identify the 1-2 reps with highest engagement -- they become your early advocates.

COACHING CHECKPOINT: Coaching checkpoint: "What worked? What felt clunky? What would make you use it more?"

Week 2 -- Workflow Integration

- ? Each rep commits to using AI for one specific workflow on every applicable task this week.
- ? The workflow must be one from this playbook -- not ad hoc. Pick the one with the highest weekly frequency.
- ? Manager tracks workflow usage, not quality -- the habit must come before the output review.
- ? Collect time-saved data informally: "How long did this take before? How long with AI?"

COACHING CHECKPOINT: Coaching checkpoint: "Is the workflow you picked the right one? What barriers are stopping you from using it every time?"

Week 3 -- Quality Calibration

- ? Review AI outputs as a team. What prompts are producing the best results? What needs more editing?
- ? Build a shared prompt library from the week's best examples. Rep-led, manager-facilitated.
- ? Introduce the quality review framework for each role (specificity, accuracy, tone).
- ? Compare outputs week-over-week: is output quality improving as prompts improve?

COACHING CHECKPOINT: Coaching checkpoint: "What is in the prompt library? Who is contributing? Who isn't, and why?"

Week 4 -- Measure, Celebrate & Expand

- ? Measure the metric defined before rollout. Compare to baseline. Be honest about the delta.
- ? Identify which reps are seeing the most benefit. Have them share their workflow with the team.
- ? Celebrate early wins publicly -- social proof is the most powerful adoption driver in a peer group.
- ? Identify the next workflow to add. Restart the cycle with one new workflow, not five.

COACHING CHECKPOINT: Coaching checkpoint: "What is the one result from this sprint you will take to your manager's manager?"

09 MEASUREMENT FRAMEWORK -- WHAT TO TRACK AND WHEN

Measure leading indicators in the first two weeks -- they tell you whether habits are forming. Measure lagging indicators at 30, 60, and 90 days -- they tell you whether habits are producing outcomes. If leading indicators are strong but lagging indicators are flat, the problem is workflow selection, not adoption. You are building the wrong habit.

LEADING INDICATORS -- WEEKS 1-2

Daily active use	Are reps using the tool daily? Threshold: 5+ uses per rep per week by end of Week 2.
Prompt quality	Do prompts include the required context fields? Spot-check 3 prompts per rep per week in 1:1.
Output editing rate	Are reps editing AI output before using it? Target: 100%. Raw output sent without edits = red flag.
Workflow coverage	Is the committed workflow being used on every applicable task? Confirm in 1:1 each week.

LAGGING INDICATORS -- 30 / 60 / 90 DAYS

AE: Follow-up send time	30 days: % of follow-ups sent within 2 hours of call end. Target: 60%+ at 30 days, 80%+ at 90 days.
AE: Pipeline review quality	60 days: Are deal risk summaries replacing "I think it's going well"? Qualitative manager assessment.
SE: RFP turnaround	30 days: Average time per RFP question vs. baseline. Target: 30-40% reduction at 60 days.
BDR: Email reply rate	30 days: Reply rate on AI-reviewed sequences vs. non-reviewed. Segment and compare.
BDR: Conversation rate	60 days: Cold calls that become real conversations. Target: 10-15% improvement vs. baseline.
CSM: Renewal outreach timing	30 days: % of renewals with outreach sent 90+ days out. Target: 80%+ of book at 60 days.
CSM: QBR prep time	60 days: Self-reported prep time vs. baseline. Target: 35-40% reduction.
Support: KB article volume	30 days: Articles created per resolved ticket. Target: 1 article per 3 resolved tickets.
Support: Repeat ticket rate	90 days: Repeat ticket volume on documented issues. Target: 20-25% reduction vs. baseline.

HOW TO CALCULATE ROI

Time saved x fully-loaded hourly cost = hard capacity recaptured

Example: AEs save 20 min per call on research and follow-up. 15 calls/week x 20 min = 5 hrs/week per rep. At \$85/hr fully-loaded (on \$150K OTE), that is \$425/rep/week, or approximately \$22K/rep/year. For a 20-rep AE team: \$440K in annual capacity recaptured. This is the floor. Improved output quality compounds on top of time savings and is harder to isolate but worth tracking separately.

Pipeline impact: track it, but give it 90 days before claiming causation

Compare AE deals with AI-assisted follow-up vs. without. Compare BDR sequences with AI-reviewed emails vs. standard. Compare CSM renewal rates for accounts with 90-day outreach vs. late outreach. Statistical significance takes 90 days minimum. Do not overclaim at 30 days.