



AI-Assisted Support Tooling Rollout Guide

A practical sequence for introducing AI-assist tools into an existing agent workflow without disrupting live support

Why Sequencing Matters

Most failed AI rollouts in support operations aren't failures of the underlying model — they're failures of sequencing. Tools get switched on for the full queue at once, agents aren't trained on when to trust vs. override a suggestion, and the first bad automated response becomes the story everyone remembers. This guide lays out the rollout order we use to avoid that.

Stage 1: Shadow Mode (Weeks 1-2)

AI-assist tooling runs alongside agents without being visible to them yet — generating draft responses and suggestions in the background so the team can compare AI output against what agents actually sent, without any risk to live tickets.

- Sample 200–300 tickets across the target queue's common categories.
- Score AI-suggested responses against actual agent responses for accuracy and tone fit.
- Identify ticket categories where the AI is consistently strong vs. consistently weak.

Stage 2: Assisted Draft (Weeks 3-4)

For the ticket categories that scored well in shadow mode, AI-generated drafts become visible to agents as a starting point — agents edit and send, they don't send unedited. This is where most of the handle-time reduction shows up, because agents are editing instead of writing from scratch.

Guardrail: No AI-drafted response is sent without an agent reviewing it during this stage, regardless of how well shadow mode scored. Trust is earned incrementally.

Stage 3: Auto-Send for Low-Risk Categories Only (Weeks 5-6)



Only the narrowest, most repetitive, lowest-consequence ticket types (see our Human-in-the-Loop whitepaper for the classification framework) move to fully automated sending. Everything else stays in assisted-draft mode.

Selecting auto-send categories

1. Must have scored above 95% accuracy in shadow mode.
2. Must have no compliance, refund, or account-access implications.
3. Must have a simple, unambiguous escalation path if the customer isn't satisfied with the automated response.

Stage 4: Continuous Calibration (Ongoing)

Weekly QA sampling continues indefinitely on both assisted-draft and auto-send categories. As accuracy holds or improves, additional categories can move toward more automation — but the migration between stages should always be evidence-led, never assumed.

Stage	Duration	What agents see
1. Shadow mode	2 weeks	Nothing — measurement only
2. Assisted draft	2 weeks	Editable AI draft before sending
3. Auto-send (narrow)	2 weeks, then ongoing	Auto-sent for pre-approved categories only
4. Continuous calibration	Ongoing	Weekly QA sampling and category review

© Infusion BI. This guide reflects our standard AI-tooling rollout methodology and is intended as operational planning guidance.