

FREE RESOURCE · FOR HEADS OF SUPPORT

20 Tips for a Successful AI Deployment for CX

A practical field guide for Heads of Support and CX Leaders who are evaluating, piloting, or scaling AI in support operations.

Reach us at

chandan.maruthi@twig.so

SOC 2 Type II · 1,200+ support teams

Proven Results

80%**Ticket Resolution**

Autonomous handling rate

90%**Faster Response**

From minutes to seconds

40%**Lower Costs**

Avg. across 1,200+ agents

Part 01

Strategy & Planning

Tips 01-07

Foundation, scoping & process design

Part 02

Launch & Tuning

Tips 08-14

RAG quality, feedback loops & QA

Part 03

Scale & Sustain

Tips 15-20

Agent training, compliance & ROI

INTRODUCTION

Why This Guide Matters

Deploying AI in customer experience is one of the highest-stakes decisions a support leader makes. Get it right and you accelerate resolution, empower your team, and delight customers at scale. Get it wrong and you erode trust, frustrate agents, and spend months unwinding a deployment that was never set up to succeed.

This guide is built for Heads of Support and CX leaders who are evaluating, piloting, or scaling AI in their operations. These 20 tips are drawn from the patterns that separate successful deployments from expensive lessons – the questions you need to ask, the traps you need to avoid, and the decisions that compound over time.

The tips are organised across three stages that mirror your deployment journey: Strategy & Planning (01–07), Launch & Tuning (08–14), and Scale & Sustain (15–20). Each tip is self-contained. Read them in order or jump to where you are in your journey.

The leaders who deploy AI successfully are not the ones with the biggest budgets or the newest technology. They're the ones who plan deliberately, involve their people, and treat AI as a living system rather than a finished product. This guide helps you be one of them.

PART 01

Strategy & Planning

Foundation, metrics, pilots & process design

PART 02

Launch & Tuning

Retrieval, feedback, QA & adversarial testing

PART 03

Scale & Sustain

Agent training, compliance & ROI scaling

PART 01 · Tips 01–07

Strategy & Planning

The foundation decisions that determine long-term success

01

Start With the Customer Journey, Not the Technology

Before evaluating any AI platform, map your current customer journeys end-to-end. Identify the top 10–15 contact drivers by volume and pinpoint where resolution is delayed by knowledge gaps, routing failures, or agent ramp-up time. AI that solves real friction generates measurable ROI from day one. Deploy it because your customers will feel the difference – not because the technology is available.

02

Define Success Metrics Before You Go Live

Agree on KPIs before a single conversation is handled. Key metrics: first-contact resolution (FCR), average handle time (AHT), CSAT/NPS impact, containment rate, escalation rate, and cost per contact. Without a pre-defined baseline you cannot prove – or improve – performance. Set a 90-day review cadence from the outset and align your executive sponsor to these numbers early.

03

Treat Your Knowledge Base as the Foundation, Not an Afterthought

AI is only as good as the knowledge it draws from. Audit your knowledge base before deployment: stale articles, inconsistent answers, coverage gaps, and tribal knowledge that lives only in agent heads will degrade AI accuracy. Invest 4–6 weeks in KB remediation before going live. A well-structured, verified knowledge base is the single highest-leverage investment you can make.

04

Involve Agents Early and Often

Agents are your best source of ground truth about what customers actually ask, how questions are phrased, and where processes break down. Run structured workshops with your top performers before designing AI workflows. When agents feel ownership, adoption skyrockets and resistance drops. They should see AI as something that makes their job better – not a threat to their role.

05

Choose a Narrow Scope for Your Pilot

The most successful AI deployments start focused: one queue, one channel, one contact reason. Trying to boil the ocean creates complexity that is hard to debug or measure. A well-executed pilot on password resets or order status creates confidence, proof points, and learning that compounds when you expand. Plan your pilot to run 6–8 weeks before any scaling decisions.

06

Design for Graceful Escalation

Every AI interaction needs a clear, frictionless path to a human. Design your escalation logic deliberately: when should AI escalate automatically? When can customers self-escalate? And crucially — what context gets passed to the agent? Warm handoffs with full conversation history are the difference between customers feeling helped and customers feeling abandoned.

07

Don't Automate a Broken Process

AI amplifies your existing processes — both good and bad. If your returns workflow requires 4 manual steps and 3 system lookups, automating it will just fail faster. Before deploying AI to any workflow, ask: 'If this were perfect, what would it look like?' Redesign the process first, then automate the redesigned version. This mindset separates transformational deployments from expensive disappointments.

PART 02 · Tips 08–14

Launch & Tuning

Go-live execution, quality control and early optimisation

08

Prioritize Retrieval Quality Over Model Sophistication

For most CX use cases, how the AI retrieves and applies knowledge matters more than which large language model powers it. Focus your evaluation on retrieval accuracy, citation reliability, and hallucination rates — not headline benchmarks. A well-tuned RAG system on your verified content will outperform a cutting-edge model paired with poor knowledge architecture, every time.

09

Build a Feedback Loop From Day One

Your AI system should be a living, improving asset — not a static deployment. Build mechanisms for agents to flag incorrect AI answers, supervisors to review low-confidence responses, and customers to signal dissatisfaction. Route flagged answers to a weekly knowledge review. Teams that build this infrastructure from launch consistently outperform those who treat AI as set-it-and-forget-it.

10

Set Realistic Expectations With Stakeholders

AI for CX is not a switch you flip. Expect 30–60 days of tuning before performance stabilises. Plan for initial CSAT fluctuations as customers adjust to new interaction patterns. Stakeholders who understand the learning curve become patient partners; those with unrealistic expectations become detractors when Week 2 doesn't match Week 12 projections. Over-communicate the maturity curve from kickoff.

11

Build Your 'AI Ready' Content Standard

Not all knowledge base content is AI-readable. Long narrative articles with embedded caveats or conflicting information mid-section cause inconsistent AI responses. Create an 'AI Ready' standard: short sentences, one answer per section, clear conditional logic ('If the customer is on Plan X...'), and explicit verification dates. Enforce this standard rigorously before and after deployment.

12

Instrument Everything From the Start

You cannot optimise what you cannot see. Ensure your AI platform logs every interaction with full context: query, retrieved sources, response generated, customer action taken, and whether escalation triggered. Build dashboards that surface your top failure modes weekly. Teams with strong instrumentation fix problems in days; teams flying blind discover them in customer complaints months later.

13

Address the Persona and Tone Question Early

Will your AI have a name? A personality? A defined escalation voice? These shape customer expectations and brand perception at scale. Define your AI persona in writing before launch: how it introduces itself, how it declines out-of-scope questions, how it transitions to humans. Review with brand, legal, and CX leadership. Inconsistency in persona is a trust issue, not a design issue.

14

Plan for the Long Tail of Queries

The top 20 contact reasons may represent 60% of volume, but the remaining 40% – hundreds of edge-case query types – is where AI breaks down most visibly. Design an explicit fallback strategy: 'I don't have enough information, let me connect you with a specialist' is far better than a confident wrong answer. Test edge cases exhaustively in your QA process before launch.

PART 03 · Tips 15–20

Scale & Sustain

From working pilot to compounding platform value

15

Train Agents on Working With AI, Not Around It

Agents who don't understand how your AI works will either distrust its outputs or accept them blindly – both are dangerous. Run training covering how the AI retrieves answers, what confidence indicators mean, how to override AI suggestions, and how to escalate feedback. AI-literate agents catch errors faster, improve the system more effectively, and deliver consistently better customer outcomes.

16

Test Adversarially Before You Go Live

Before launch, dedicate a team to breaking your AI. Attempt prompt injection, test off-topic queries, explore every edge case, and run a red-team exercise with your most experienced agents playing difficult customers. Production will always surface surprises – adversarial pre-launch testing dramatically shrinks the blast radius of Day 1 failures. This step is non-negotiable.

17

Align Legal and Compliance Early – Not at Launch

AI in customer service raises real compliance questions: data retention, PII handling, regulatory disclosures (especially in financial services and healthcare), and liability for incorrect answers. Bring legal and compliance into vendor evaluation, not deployment week. Build compliant guardrails from the start. Retrofitting compliance is expensive and disruptive; designing for it from day one is not.

18

Communicate Transparently With Customers

Customers have a right to know when they're interacting with AI. Beyond legal requirements in many jurisdictions, transparency builds trust. Design clear, honest disclosure language: 'Hi, I'm an AI assistant.' Avoid personas that obscure the AI nature. Teams that communicate transparently consistently report higher CSAT with AI interactions than those who attempt to simulate human responses.

19

Scale What Works, Retire What Doesn't

After 90 days, you will have data. Use it ruthlessly. Double down on workflows where AI achieves high containment and strong CSAT. Retire or redesign workflows where AI consistently escalates, generates complaints, or consumes agent time to correct. Treat your AI deployment as a portfolio: some bets pay off, some don't, and capital should flow toward what is working. Review quarterly.

20

Think Platform, Not Project

The most common failure mode for AI in CX is treating it as a one-time project with a finish line. Teams that unlock compounding value treat AI as a living platform: continuously updating knowledge, expanding to new channels, incorporating new capabilities as they mature, and measuring ROI on an ongoing basis. Assign a dedicated AI programme owner. Make continuous improvement a standing meeting.

TWIG · AI FOR CX

Ready to Put These Tips Into Practice?

Twig helps support teams deploy AI that agents trust and customers love — built on verified knowledge, grounded in accuracy, designed for CX.

Reach us at

chandan.maruthi@twig.so

70%

Autonomous Resolution

90%

Faster First Response

40%

Lower Support Costs