



Why AI Agents Work: A Practical Guide for Success

»

 opusresearch

July 2025

Ian Jacobs, VP & Lead Analyst, Opus Research

Opus Research, Inc.
893 Hague Ave.
Saint Paul, MN 55104

www.opusresearch.net

Published July 2025 © Opus Research, Inc. All rights reserved.



Beyond the Hype—What AI Agents Actually Mean for Customer Support

Let's skip the breathless AI headlines and get to the point: this guide is for companies who are ready to do something with AI in customer service—not just think about it. Specifically, we're talking about AI Agents for the contact center, whether that's a VoiceAI Agent answering calls, a ChatAI Agent handling inbound messages, or a multimodal approach doing both. AI Agents aren't scripted chatbots in disguise—they're conversational assistants powered by LLMs and constructed using prompts. Think of them as software agents trained to carry the support load, handling everyday questions with the consistency of a robot and enough smarts to keep customers happy.

The goal here isn't just automation for automation's sake. It's to shift how support works: resolve routine issues instantly, surface smarter insights for human agents, and create the kind of experience customers actually want. And it starts with being strategic from day one: setting expectations, choosing the right problems to solve, and avoiding the common pitfalls that turn promising projects into expensive misfires.

This guide is designed to help teams move from interest to action. If you're exploring AI-powered automation for the contact center, this is your starting point. Minimal buzzwords, just a clear path to getting started.

Getting Your House in Order: Data and Call Drivers

Before you even think about deploying an AI Agent, you need to understand what it will hear and how it will respond. That means tackling two foundational questions:

1. Do we have the right data to train and guide the system?
2. Do we know which customer interactions are actually worth automating?

What “Data Readiness” Really Looks Like

You don't need a perfect data lake or a decade of analytics maturity—but you do need relevant inputs. Start with what you already have:

- Call and chat transcripts are goldmines for understanding how your customers communicate.
- Best-in-class interactions, whether from live agents or early AI pilots, help models understand what “good” looks like.
- Both structured data (CRM entries, disposition codes) and unstructured data (free-text notes, voice logs) matter—and they work best in tandem.

Data integration can get messy, especially if you're stitching together legacy systems. But modern platforms offer middleware, APIs, or services that can speed this up. Don't let perfect be the enemy of progress.

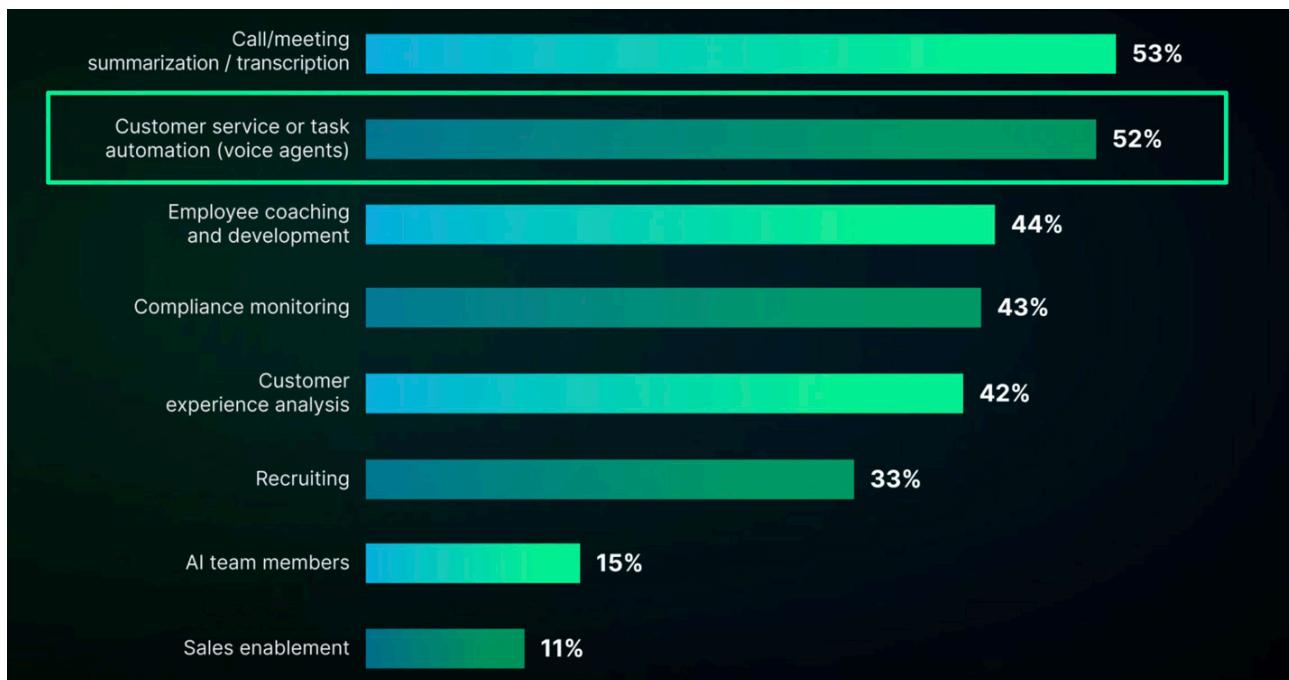
Also critical: compliance. If your industry is regulated (finance, healthcare, etc.), privacy requirements need to be accounted for up front—not patched on later. VoiceAI agents in particular must respect consent laws, retention policies, and auditability standards from day one.



Customer Service Automation: From IVRs to VoiceAI

Leveraging voice and call data can create new business opportunities for many organizations. A recent survey conducted by Opus Research found much effort is aimed at driving automation, with 52% of respondents believing “customer service or task automation” to be the most transformative use case for VoiceAI technology. (figure below) This speaks to the clear value found in automating repeatable tasks but a focus on experience and effectiveness must remain.

Figure 1: Most Transformative Use Case for Voice AI Technology



Source: Opus Research Voice AI Survey, n=400 (2025)

Zeroing in on the Right Use Cases

Not every conversation should—or can—be handled by an AI Agent. That’s why the smartest starting point is identifying your top call drivers, especially Level 1 (L1) interactions: high-volume, low-complexity inquiries where automation adds value without risking customer trust.

Here are some common examples of call types that fit this bill:

- “What’s the status of my order?”/ “Where is my order?”
- “How do I reset my password?”
- “Can I reschedule my appointment?”

Segment these from more complex L2 or L3 scenarios where human judgment still rules. Then look for patterns where AI can resolve the issue outright, or tee things up with all the context pre-filled and seamlessly provided to a human agent. (Figure 2 below)



This is also where business rules come into play—helping determine when to escalate, when to transfer, and when to simply say, “We’ll get back to you.” Done right, this groundwork sets the stage for everything that follows. Skip it, and you’re building your AI Agent on a shaky foundation.

Figure 2: Automation Potential by Call Type

Call Type	Examples	Automation Potential	Recommended AI Agent Role
Level 1 (L1) High-volume, low-complexity	Order status, password resets, appointment scheduling	● High	Full automation with VoiceAI or ChatAI Agent
Level 2 (L2) Moderate complexity, some judgment required	Billing discrepancies, product troubleshooting, plan changes	● Moderate	AI Agent gathers context → human takes over
Level 3 (L3) Low-volume, high-complexity	Escalations, complaints, sensitive or high-value cases	● Low	AI Agent triages and routes to human expert

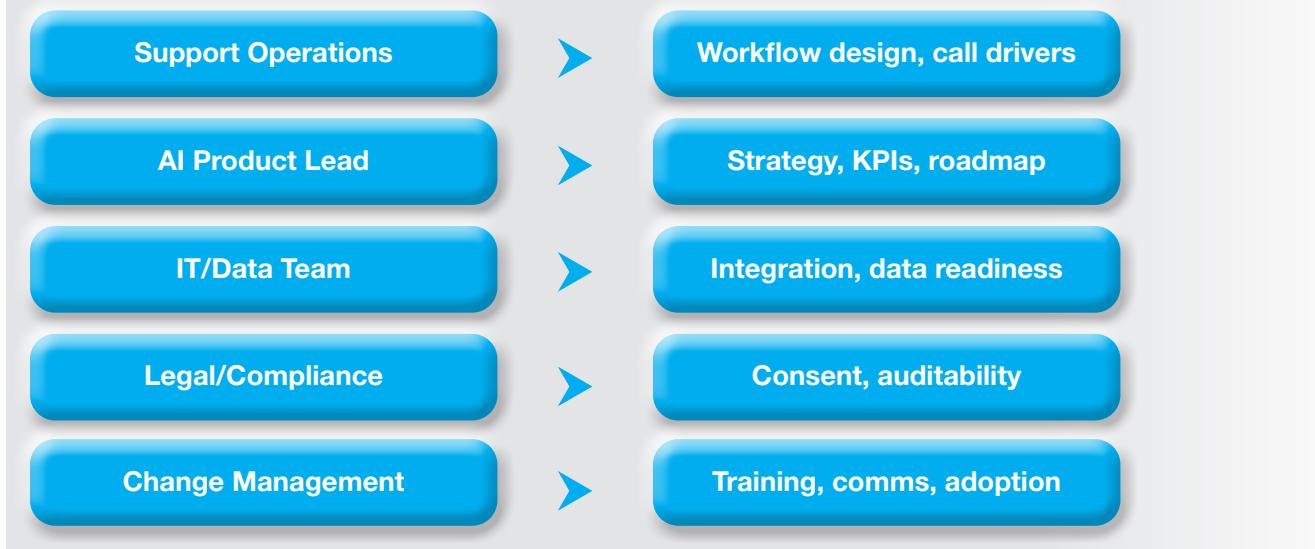
Building the Right Team and Strategy

Deploying AI Agents in customer support isn’t just a tech decision—it’s an organizational decision. Getting the right people involved early can mean the difference between a pilot that quietly fizzles out and a rollout that sticks. Spoiler: you’ll need more than just IT and ops.

Start by identifying a cross-functional team that spans support, product, legal, and yes, engineering. These folks don’t all need to be AI experts, but they do need to be aligned on outcomes. What’s the AI Agent supposed to accomplish? What are we measuring? Who owns what when things go sideways?

A few key roles tend to show up over and over in successful deployments:

Figure 3: How Cross-Functional Teams Support Outcomes





You don't need a huge team—just the right mix of buy-in and ownership. A great real-world example: one national dental healthcare company with 425+ locations retrained frontline agents to become AI QA reviewers and prompt engineers, while redirecting others into inside sales roles. Not only did this improve AI outcomes, but the strategy also opened up new career paths internally, making the company a more attractive option for potential employees.

And then there's the external team: your vendor. Given how transformative AI Agents can be, you'll want to choose one that does more than just sell licenses. You want a partner that shows up with implementation playbooks, shares benchmarking data, and sticks around after go-live to tune performance and troubleshoot. Most technology today looks promising at the POC stage, but a trusted partner is needed to bring it to production.

Here's some key things to look for:

- **Scalability and integration:** Can their solution plug into your existing systems—and grow with you as needs evolve?
- **Industry expertise and customization:** Do they understand your sector's quirks and offer prebuilt workflows or domain-tuned models?
- **Security, compliance, and reliability:** Are they equipped to meet your regulatory requirements and support enterprise-grade uptime?

Finally, set expectations. AI Agents won't transform your entire contact center overnight—and that's a good thing. Start with one or two scoped use cases. Test, tune, and only then scale. That's how successful teams build both momentum and trust.

Figure 4: Vendor Evaluation Criteria for AI Agents

Criteria	Why It Matters for Large Enterprises	Key Questions to Ask Vendors
Scalability & Integration	Supports growth, avoids workflow disruption	Can you handle peak loads? How do you integrate with our current stack?
Industry Expertise & Customization	Ensures relevance, handles both simple and complex cases	Do you have similar clients? How customizable is the solution?
Security, Compliance & Reliability	Protects sensitive data, ensures uptime and trust	What certifications do you have? What are your uptime guarantees?

Starting Small: Pilot to Scale

Let's assume you've done the prep work: you've identified solid L1 use cases as outlined above, your data is in good shape, and your team (internal and vendor partner) is on board. Now comes the fun part: proving the AI Agent can deliver real value—and then scaling it without creating chaos.



Begin with a Focused Pilot

Your first deployment should be narrow by design. Think proof-of-concept, not transformation. A few characteristics of a good pilot:

- The use case is high volume and low risk (order status > billing disputes).
- The success metric is clear and measurable (containment rate, AHT reduction, improvement in CSAT).
- You can run it in parallel with existing systems and workflows—without breaking anything.

This phase isn't just about testing the AI. It's about pressure-testing the whole system: integrations, handoffs, analytics, and agent feedback loops.

Design for the Handoff, Not Just the Bot

Even the best AI Agent won't get it right 100% of the time. And that's fine—if your escalation strategy is nailed down and smooth. A warm handoff to a human agent (with full context intact) keeps the customer experience from feeling robotic when it matters most.

This is where human-AI collaboration shines. Imagine a VoiceAI Agent handling a credit card activation call that turns into a fraud concern—the moment the caller says “I didn’t apply for this,” the AI tags it as potential fraud, summarizes the context, and transfers the call to a live agent with the transcript and key flags. The agent doesn't waste time asking the customer to repeat themselves, and the customer doesn't have to re-explain a potentially stressful situation. That's not just fallback—it's a coordinated response that combines AI efficiency with human empathy.

Your pilot should build in these kinds of fallback paths, along with agent assist tools and post-interaction QA. The more your AI Agent can work with humans instead of trying to replace them, the more trust (and value) you'll gain.

Scale in Phases, Not in Panic

Once the pilot shows promise, the natural urge is to roll it out everywhere. Resist that urge. Instead, scale in waves:

1. Expand to similar intents or channels (e.g., from password resets in chat to password resets in voice).
2. Move into new verticals or customer segments (e.g., high-value accounts).
3. Add more complexity—only when your AI Agent has proven it can handle nuance.

Keep monitoring KPIs, run retrospectives, and tune regularly. AI adoption is a process, not a one-off project.

Measuring Success and ROI

Once your AI Agent is live—whether in a pilot or broader rollout—it's time to move from setup to scorekeeping. But forget vanity metrics. Measuring success means tracking what actually matters to your business, your customers, and your support team.

Start with a handful of core KPIs that reflect value at different levels. Some examples of common KPIs:

- **Containment rate:** What percentage of interactions are resolved without human handoff?



- **Average Handle Time (AHT):** Has it gone down for agents? Or up, but for good reason (i.e., only complex cases remain or increased upsell opportunities)?
- **CSAT or NPS delta:** Are customers as happy with AI-led interactions as they are with human-led ones—or happier?
- **Escalation rate:** How often is the AI Agent handing off to a human? And is it doing so at the right moments?
- **Intent coverage:** Of all customer inquiries, how many can the AI reliably recognize and respond to? This helps gauge the system's breadth and readiness to scale.

You'll also want to track post-interaction QA, especially early on. Are agents flagging escalations that could've been avoided? Is the AI making bad assumptions or missing key intents? Your mindset should be that these aren't failures—they're learning opportunities. And the data should feed directly back into your tuning loop.

And then there's the big one: ROI. This isn't just cost savings from deflection (though that's often a big part). You'll want to look at total impact:

- Are you saving agent time—or reallocating it to higher-value tasks?
- Are you increasing first-contact resolution?
- Are customers getting faster answers without the wait?

Most teams won't get a clean "we saved \$X" in month one. But if your KPIs are moving in the right direction post-deployment—and the customer and agent experience aren't suffering in the process—you're on the right track. Fairly quickly after that you'll be able to tot up the hard numbers of ROI.

Future-Proofing Your Investment

AI Agents aren't just another tech rollout—they're part of an ongoing shift in how support operations run. So, while your first phase might focus on basic L1 automation, it's worth thinking ahead to where this can (and should) go next.

The first thing to plan for is adaptability. Customer behavior changes. Business rules evolve. AI models improve. Make sure both your architecture and vendor partner can flex with those changes—whether that means new channels, new languages, new use cases, or new integrations.

Next is continuous innovation. Never treat your AI Agent like a finished product. You'll need a plan to regularly refresh training data, revisit your use case roadmap, and bring in emerging capabilities like generative summarization, multimodal interactions, or multi-intent recognition when they make sense. A year from now, your customers may well expect things your current system isn't built to handle—so start creating feedback loops now.

A year from now, your customers may well expect things your current system isn't built to handle—so start creating feedback loops now.



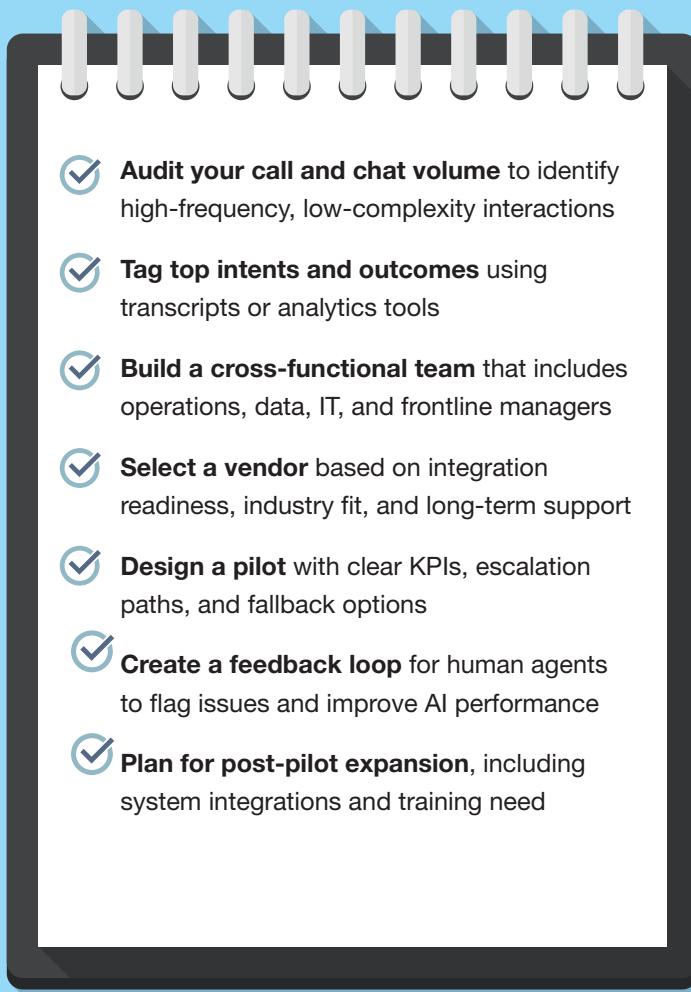
Also, as your AI Agents take on more responsibility, the role of human agents will evolve. Support teams won't just handle escalations—they'll act as reviewers, trainers, and high-touch service specialists. A future-proof strategy includes upskilling your people, not just your tech. For example, agents who previously fielded basic password reset calls might be trained to review AI Agent transcripts for accuracy, tag new intents for future automation, or manage real-time AI escalations in sensitive situations. These roles don't replace today's frontline work—they provide a path to evolve it.

And finally, don't fall into the "automate everything" trap. The sweet spot remains a hybrid model—AI handling the heavy lifting, humans stepping in when nuance or empathy is required. That balance is what keeps the customer experience strong even as automation scales.

Getting Started and Next Steps

Here's a simple starter checklist to get things moving:

Figure 5: A Checklist for Getting Started





Deploying AI Agents isn't about placing a big bet—it's about running a well-structured pilot with clear business goals. If you've made it this far, you already know the essentials:

- What to automate first (and what to leave alone)
- How to align your internal team and vendor support
- What success looks like—and how to measure it
- How to scale without breaking what already works

Now it's time to move from planning to action. Start with a real use case. Define the KPIs. Build a pilot that can grow. If you're deliberate now, you'll set yourself up for wins that are not just measurable—but sustainable.

Deploying AI Agents isn't about placing a big bet—it's about running a well-structured pilot with clear business goals.

About Opus Research

Opus Research is a diversified advisory and analysis firm providing critical insight on software and services that support digital transformation. Opus Research is focused on the merging of natural language understanding, machine learning, conversational AI, LLMs, conversational intelligence, intelligent authentication, and digital commerce.

For sales inquires please e-mail info@opusresearch.net or call +1(415) 904-7666

This report shall be used solely for internal information purposes. Reproduction of this report without prior written permission is forbidden. Access to this report is limited to the license terms agreed to originally and any changes must be agreed upon in writing. The information contained herein has been obtained from sources believed to be reliable. However, Opus Research, Inc. accepts no responsibility whatsoever for the content or legality of the report. Opus Research, Inc. disclaims all warranties as to the accuracy, completeness or adequacy of such information. Further, Opus Research, Inc. shall have no liability for errors, omissions or inaccuracy in the information contained herein or interpretations thereof. The opinions expressed herein may not necessarily coincide with the opinions and viewpoints of Opus Research, Inc. and are subject to change without notice.

Published July 2025 © Opus Research, Inc. All rights reserved.