

Your help desk needs help

Six steps to deploying an AI agent that actually solves help desk tickets

1

Define the scope and success metrics

Start small. Target a high-volume, low-complexity workflow (e.g., password resets or access requests). Set measurable goals like reduced SLA breaches, faster resolution time, or ticket deflection rate.



Kenny Park 11:55

Let's automate password reset requests. We need to reduce our ticket deflection rate by at least 15%.

2

Map your data and tools

List the systems the agent needs to read from and act within, such as identity, ticketing, and communication platforms. Ensure clean, accessible data and clear APIs between tools.



3

Build reasoning and actions

Use an LLM or inference engine connected to defined "tools" that execute actions such as creating, updating, or closing tickets. Focus on accuracy and reliability over conversational flair.



4

Implement governance and security

Treat your agent like any other production system. Log every action. Add human-in-the-loop approvals. Validate inputs and permissions before execution.



5

Deploy where users already work

Launch in Slack, Teams, or your existing ticketing portal. Keep the interface familiar so adoption isn't the barrier.



6

Monitor, measure, and iterate

Track tickets resolved, time saved, and deflection rate. Feed successful outcomes back into the model and expand to new use cases once performance is consistent.



Want to see what this could look like for your team?

Run your numbers with the [ITSM Benefits Calculator](#) to estimate hours saved and tickets deflected.