

SOLUTION BRIEF

Call Center Omnichannel AI

AI-powered voice, SMS, chat, and social deployed on top of existing operations, without ripping out what already works.

**The call center has
new architecture.**

Restructured, not rebuilt. AI handles the routine volume. Your people handle the calls that need a person. And it all runs on the systems you already have.

70% of inbound call volume is routine & repeatable

4 unified channels voice, SMS, chat, social

3 labor pools in concert local, shared, AI

High-volume call centers are stuck between two pressures. Staff costs keep climbing and callers can't be handed a self-service app and left to figure it out. Earlier fixes fell short. Centralized labor pools and SaaS telephony pulled jobs away from the communities they served, then stalled on legacy systems that were never built to talk to AI. So the real question isn't whether AI can handle the calls. It's how to deploy AI in a way that respects community labor, serves complex callers, connects to the data you already have, and earns confidence before you scale.

One Agent, Every Channel

Most AI call center deployments start and end with voice. We treat voice as one channel in an omnichannel agent that maintains context across every medium a caller touches.

Voice Real-time speech-to-text, sentiment, and intent resolution.

SMS Two-way texting of reminders, confirmations & cancellations.

Chat Web & in-app flows on the same agent logic.

Social Messages routed to the same workflows.

What the Platform Does

Multi-LLM Reliability. Best model per task, auto-routing around outages, and no single-model dependency.

Multilingual Neutral Voice. Calibrated across diverse caller populations.

Agentic Actions. Goes beyond answers to handle updates, cancellations, and queries end-to-end.

Command Center. Real-time view of call volume, resolution rates, escalation patterns, and AI performance.

Parallel Scaling. Thousands of concurrent calls and messages, handled with ease.

Smart Routing & Sentiment. Routes every interaction to the right resource before frustration sets in.

Safe Failover. Configurable triggers for warm hand-off to humans with full context.

Agent Assist. Access history and answers so human interactions are faster and better informed.

The Right Resource

POOL 1: Local High-Touch Team. Community-embedded team that handles relationship-critical interactions, complex cases, and escalations.

POOL 2: Shared National Team. Available off-hours peak coverage and business continuity. Enables labor sharing across multiple locations without displacing local high-touch teams.

POOL 3: AI Omnichannel Agent. Handles routine interactions at scale. Multilingual and always available. Routes to human pools when complexity or sentiment triggers escalation.

Pilot first. Scale with confidence.

Two-week sprints that end with a gated checkpoint and your right to pause or stop with no obligation. It's not just a commercial term, it's a delivery philosophy. For teams burned by big tech commitments that underdelivered, every gate is a real decision point.

PHASE 1: Pilot (3 months, 6 sprints)

Build a gated model. Deploy AI agent, build data layer, establish hybrid labor, deploy text channel.

PHASE 2: Expand (3 months)

Increase agent actions. One interface across databases. Inbound calls consolidated to single endpoint.

PHASE 3: Automate (3 months)

Extend AI into operations. Dispatch decisioning, automated cancellation and no-show workflows. Data layer value increases.

PHASE 4: Scale (ongoing)

Multi-location rollout. Expand operations, add task agents, outreach and subscription management.

Why CloudFactory?

CloudFactory helps enterprises run AI reliably in production, especially in high-stakes use cases where errors create real business risk. By combining policy enforcement, human validation and workflow control, we ensure AI produces consistent, auditable outcomes. CloudFactory integrates with existing systems so teams can move from experimentation to production and scale AI without losing control.

