



WHITE PAPER

The Zero Contact Center

Five years is the perceived window for contact center transformation. In reality, leading organizations began this journey 2–3 years ago and are already midway. The question isn't whether it will change, it's whether you are accelerating fast enough to keep up.



What You'll Find Inside

Executive Summary 03

The Five-Year Thesis 04

The Zero Contact Center Maturity Model 05

The Dual-Track Operating Model 06

Key Enabler: Coforge Signal 08

From Call Types to Task-Level Granularity 10

The Economics 11

Three Mindset Shifts 12

A 90-Day Starting Roadmap 13

Executive Summary

The contact center industry is on the brink of a structural shift, not an incremental one. The pace of AI capability means that within the next 2–3 years, AI will resolve a majority of interactions, with leading organizations already progressing toward this state. The overwhelming majority of contacts that today flow through human agents — not just the simple queries, but the moderate and difficult ones too — will be handled by AI. What remains for humans will be a small, high-value tier: complex, VIP, and empathic-moment interactions.

We call this end-state the **Zero Contact Center**. The leaders who reach it first will operate with a 40-50%+ lower cost-per-contact, higher CSAT, and a fundamentally different P&L. Those who delay will carry a structurally higher cost base while competitors transition faster.

The key takeaway:


If you are not already running an explicit Zero Contact Center program - with task-level use case decomposition, a layered AI architecture, and a 24-36 month roadmap - you are not behind on AI. You are behind on strategy.



40-50%+

**Cost-per-contact
reduction**

Achievable within 24-36 months



3 Layers

Of AI-led transformation

Front door, copilot, human



2-3 Years

**Compressed
transformation window**

The leaders are already at the
halfway point

The Five-Year Thesis

Why Contact Centers as We Know Them are Ending

AI capability in customer service is on a steep curve. Today, most production deployments handle “easy” queries reasonably well: password resets, balance inquiries, order status. Within 18 months, the same systems will reliably handle moderate-complexity queries: multi-step transactions, light reasoning, and contextual policy decisions.

Within 36 months, they will move into difficult territory: judgment calls, exception handling, multi-system orchestration.

By the five-year mark, the residual workload for human agents will be a thin top layer: genuinely complex investigations, high-stakes financial or medical decisions, sensitive escalations, and VIP relationships where the human itself is the product.




Figure 1. Projected share of contact volume that AI can resolve autonomously, by query complexity.

Why This Matters Now, Not Later

Three forces are compounding. First, foundation models keep getting better at the things contact centers do: multi-turn dialog, tool use, structured reasoning over policy documents. Second, the integration tooling around them (orchestration platforms, voice stacks, agent frameworks) has matured to production-grade. Third, the economics have crossed an important threshold: the cost of an AI-handled contact is now meaningfully below the cost of a human-handled one for any task in the easy-to-difficult band.

The result is that the gap between leaders and laggards is no longer measured in features. It is measured in the operating model — which takes 24-36 months to redesign and deploy. By the time the capability gap is undeniable, you cannot catch up by buying more software. This means most organizations are not starting a five-year journey — they are catching up on one already underway.

Before beginning any major initiative like this, it is critical to define the starting point. Below, we will outline a maturity model that helps honestly assess your current state and identify the operational and architectural shifts required to make a Zero Contact Center a reality.

The Zero Contact Center Maturity Model

Where you are vs. Where you need to be




Figure 2. The Zero Contact Center Maturity Model. Cost reduction figures are typical ranges observed across mid-to-large operations.

- 01** | Most contact centers today sit at Phase 1, which is characterized by point automation. They have a chatbot for FAQs, a callback scheduler, perhaps a real-time transcription tool. These deliver 5-10% efficiency gains and stop there.
- 02** | Phase 2 is where the curve gets steeper, as AI acts as the “front door” and human agents are assisted by an AI co-pilot.
- 03** | The defining feature of Phase 3 is orchestrated, task-level AI across channels. This is where most leaders want to be in around 18 months.
- 04** | Phase 4 is a truly Zero Contact Center, the strategic destination and end-state on this journey.

How to Use This Model

To ensure an honest assessment, you should evaluate your current maturity level based on your weakest dimension, not your strongest.

Most organizations claim Phase 2 because they have one impressive bot — but their agent-side tooling is still a tab sprawl, their channels operate independently, and their use cases are still defined at call-type level. In practical terms, that is Phase 1, with a Phase 2 demo.


The path from Phase 1 to Phase 4 is not linear and not optional. Every stage requires both technology and operating-model change. Skipping the operating-model work is the most common reason transformations stall.

In the following sections, we will examine a few of the operational and behavioral changes it takes to progress along this journey. We will pay particular attention to achieving the major jump required between early-stage and scaled maturity levels.

The Dual-Track Operating Model

The **Zero Contact Center** is not a chatbot strategy. It is a fundamental redesign of how your organization relates to customer contacts. Most contact center transformation initiatives pursue a single track: Automate the contacts that come in

The leaders who reach Zero Contact faster run two parallel tracks:



Track 1 - Automate

Every incoming contact is met first by AI. Voice, chat, email, and messaging all flow through an AI layer before reaching a human agent. The objective is end-to-end resolution, not deflection.

Where a human does pick up, an AI copilot sits beside them, shortening handle time and improving accuracy. The three-layer funnel within this track is described below.



Track 2 - Eliminate

This is the more powerful and more overlooked track. Instead of resolving a contact after it arrives, ask **“Why is the customer contacting us at all?”**

It could be the wrong price displayed on a website, a poorly explained policy, a confirmation email that never arrives — each of these generate thousands of contacts that should never exist. By working with upstream teams in reservations, operations, IT, and policy to fix the root cause, you can eliminate those contacts permanently. No AI is needed, just cross-functional accountability.



Figure 3. A dual-track model anchored by Coforge Signal diagnostics as the starting point, running Automate and Eliminate in parallel, converging on a contact center that is a strategic insight hub, not a cost center.

Both tracks compound over time. As AI handles more contacts, the data exhaust gets richer, making it easier to identify elimination candidates. As elimination reduces volume, the AI handles a higher proportion of remaining contacts more accurately. The feedback loop between the two tracks is what makes the economics accelerate after year two.

The Three-Layer Funnel

Within the **Automate** track, three layers compound. The first is the AI front door where every contact starts. The second is the agent copilot: where AI assists every human interaction. The third is the human core, reserved for complex cases, VIPs, and moments where human judgment genuinely changes the outcome.

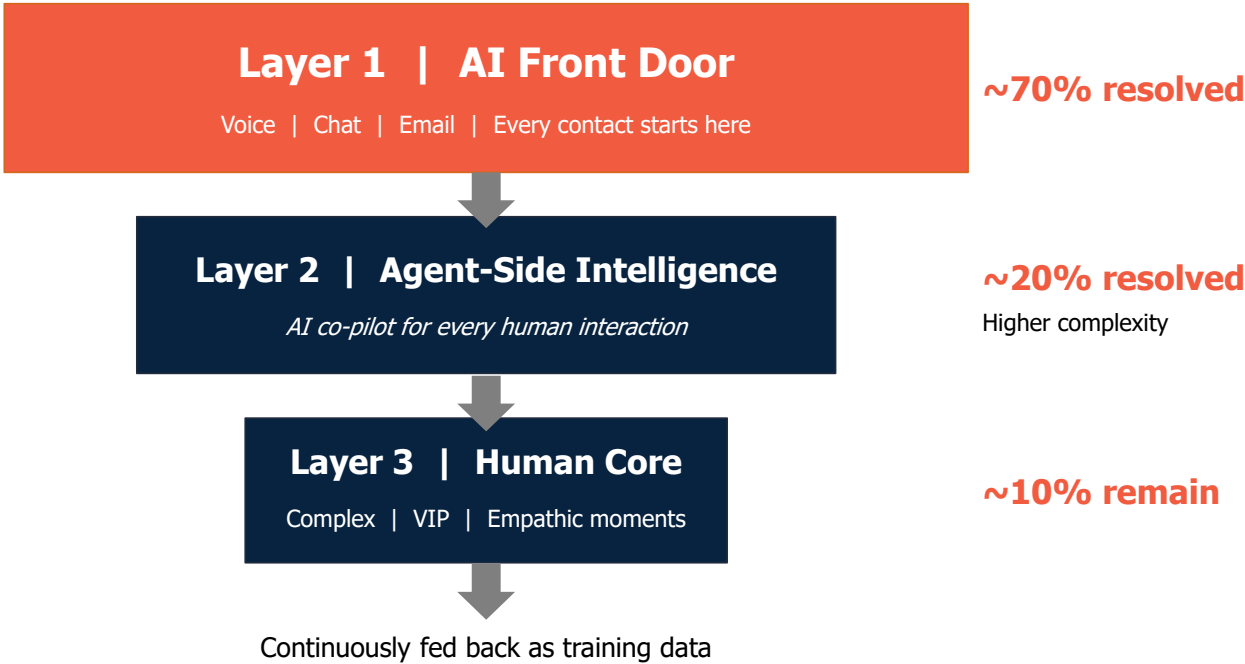


Figure 4. The Zero Contact Center funnel. Volumes shown are target end-state within 24-36 months of mature implementation.

- 01 Layer 1: The AI Front Door**

Every incoming contact, regardless of channel, is met first by AI. Voice calls hit a conversational AI before any human; chat and email are answered by an LLM-powered agent with full system access. The AI is not a deflection layer, it is the primary handler. It books the change, processes the refund, updates the policy, sends the confirmation. It only escalates when it genuinely cannot resolve, or when policy mandates a human (high-value transactions, sensitive situations).
- 02 Layer 2: Agent-Side Intelligence**

When a contact does reach a human, the human is no longer working alone. A second AI layer sits beside them in real time: surfacing context from prior interactions, drafting responses, checking policy, auto-filling system fields, summarizing the call for the wrap-up note. The agent's role shifts from data-entry-and-recall to judgment-and-empathy. AHT drops 25-30%, error rates fall, and the training curve for new agents (human) shortens dramatically.
- 03 Layer 3: The Human Core**

A small, highly skilled tier handles what neither AI nor co-piloted humans can: ambiguous investigations, sensitive emotional moments, high-stakes decisions where accountability matters, and VIP relationships where the human relationship itself is the value. This tier is smaller in headcount but higher in skill, compensation, and CSAT contribution.

The critical question:

Before you decide which contacts to automate and which to eliminate, you need to know exactly why customers are contacting you today — at the level of specific intent, not broad call type. This is the role of Coforge Signal — the diagnostic foundation for the transformation.

Key Enabler: Coforge Signal

Making sense of the ocean of contacts

Every Zero Contact Center journey begins with the same question: What is happening in your contact center, at the level of granularity that drives design decisions?

The answer is rarely what leaders think. Boards and dashboards show a breakdown by call types: Booking change: 22%; Baggage: 12%; Refund: 10%. **These numbers are accurate but useless for transformation.** They cannot tell you which sub-intents are automatable, which are driven by upstream failures that can be eliminated, which require human judgment, or what the repeat-contact rate looks like by reason. You cannot design Zero Contact Center programs on call-type data.

Coforge Signal is a proprietary generative AI analytics solution built specifically for contact centers. It is industry agnostic and can be fully customized to meet the specific needs of your industry and your customers. Coforge Signal ingests large samples of live contacts (from voice, chat, email, and messaging), transcribes them at scale, and applies NLP and LLM-based analysis to produce a definitive intent taxonomy with quantified automation and elimination potential for each cluster.

This is not a consulting exercise — it is a data-driven diagnostic. It is a data-driven diagnostic that takes 4-6 weeks and produces a roadmap with specificity that no workshop can match.




Figure 5. The four steps of Coforge Signal. The output is a complexity vs. value matrix and intent taxonomy that drives every subsequent design decision. This diagnostic becomes the starting point for both the Automate and Eliminate tracks.



What Coforge Signal Produces

The intent map generated by Coforge Signal does not just rank contact reasons by volume. For each intent cluster, it calculates the share of contacts that:

- Are automatable today
- Will become automatable with system access or process change
- Have an identifiable upstream fix (elimination candidates)

It also details the repeat contact rate (indicating customer frustration and operational waste), and the relative complexity and business value of addressing each cluster.

Left: contact volume vs automation potential (bubble size = share of total volume). Right: recommended track by contact reason.




Figure 6. Illustrative Coforge Signal output.

Left: Contact volume vs automation potential. Bubble size = share of total volume. Green bubbles are elimination candidates.

Right: Recommended track by contact reason.

Why This is the Right Starting Point

Organizations that go straight to automation without mapping and understanding customer intent typically achieve 10-15% deflection and plateau.

Why? Because they have automated the wrong things, left the elimination opportunities untouched, and built a roadmap on call-type data that obscures more than it reveals.

Phase 1 costs a small fraction of the implementation investment. The organizations that run it first consistently produce programs that are two to three times more effective at reducing cost per contact.

A note on approach:

Coforge Signal is deployed as the intelligence layer in a Zero Contact Center engagement — the diagnostic that powers both the **Automate** and **Eliminate** tracks. The output belongs entirely to your organization, and every subsequent design decision is grounded in that insight.

From Call Types to Task-Level Granularity

This is the part most contact center leaders get wrong, and it is why so many AI projects plateau at 10-15% deflection.

Most contact centers describe their work in terms of **call types** such as booking changes, billing disputes, or technical support. These are useful for reporting, but useless for AI design.

Key insight:

A “call type” is an aggregation of dozens of distinct tasks. Until you decompose it, you cannot tell which parts AI can handle today, which parts it will handle in 12 months, and which parts will always need a human.

For example, take a **booking change** — a top-line item that might represent 12% of your call volume in travel, hospitality, or any reservation-driven business. Inside it sit at least six distinct task families: date changes, route changes, passenger additions, cabin upgrades, refund splits, and fee waivers.

Each of those decomposes further into atomic units like identity verification, inventory check, payment processing, and others.




Figure 7. A single call type decomposes into multiple task families, each of which decomposes further into atomic, automatable units.

When you map your operation at this level, three things happen. You discover that perhaps 60% of the atomic units can already be fully automated. Another 25% can be automated with a modest investment in system access or process simplification. You also discover where your real human-judgment moments actually sit — and they are often a much smaller set than you assumed.

This decomposition exercise is the single highest-leverage piece of work in the contact center journey. It is also the one most organizations skip, because it is unglamorous, slow, and requires process expertise that **no AI vendor can supply for you.**

The Economics

A five-year cost-per-contact view

Strategy without numbers is opinion. Here is the cost-per-contact trajectory across a transformation period comparing a traditional operation against a serious Zero Contact Center program.

5-year cost-per-contact: Traditional vs. Zero Contact Center




Figure 8. Indexed cost-per-contact over five years. Traditional ops reflect typical wage inflation and modest productivity gains. ZCC path reflects compounding AI deflection plus agent productivity gains.



40-50%+

Reduction in cost per contact

within 24-36 months



30-40%

Reduction in average handle time

with agent co-pilot



2-3x

CSAT increase

on AI-handled simple queries

Where the Savings Actually Come From

It is worth being precise: The cost reduction is not a single lever. Roughly half comes from absolute volume reduction at Layer 1. The AI front door simply resolves contacts that previously required a human.

The other half comes from productivity at Layer 2: agents handling more contacts per shift, with shorter handle times, lower error rates, and shorter ramp-up periods for new hires. A smaller (but real) contribution comes from the structural shift in workforce mix as the operation moves toward a smaller, higher-skilled human core.

These numbers assume execution. Organizations that buy AI tools without redesigning the operating model around them typically capture less than half of the potential value. This is the difference between deploying AI and operating as a Zero Contact Center.

Three Mindset Shifts

Technology is the easy part. The harder shift is in how leaders, middle managers, and front-line supervisors think about the work. Our experience has shown that there are three key changes that are non-negotiable.

Shift 1: From projects to operating model

Stop running AI as a portfolio of pilots. Start running it as a redesign of how your contact center operates. That means a single multi-year program, a single accountable executive, a single integrated roadmap.

Pilots that do not lead to an operating model change are noise.

Shift 2: From AI as tool to expertise as multiplier

AI is the enabler, not the answer. The organizations achieving real results have deep process expertise (how does this booking change actually work?), deep systems expertise (which APIs, which data, which legacy constraints?), and apply AI as the layer on top.

Buying a vendor platform without that underlying expertise produces demos, not transformation.

Shift 3: From today's metrics to tomorrow's mix

Your current dashboards are designed for an operation where humans handle 100% of contacts. They will mislead you for an operation where AI handles 70%.


To avoid optimizing the wrong things, build the metric set for your end-state now. It should include: AI containment quality, AI-to-human handoff effectiveness, human-tier CSAT on complex cases, and cost-per-resolved-contact across the full funnel.

Key decision point:

A useful stress test for any leadership team is to ask the following question:

"If AI handled 70% of your contacts tomorrow, would your org structure, vendor contracts, and training programs still make sense?"

If the answer is "no" — and it almost always is — start there.



A 90-Day Starting Roadmap

You do not need a full long-term plan to start — you need a focused 90-day entry point. You need a credible 90 days that turns into a plan. Below we demonstrate what that looks like.

Phase	Days	What you do	What you get
Diagnose	0-30	• Deploy Coforge Signal to identify contact drivers at task-level granularity and decompose them into atomic tasks. • Map current AI footprint and gaps. • Baseline cost-per-contact and AHT by task family.	• Honest current-state map • ZCC maturity rating • Sized opportunity by task family
Design	30-60	• Define your 24-month target operating model across all three layers. • Pick 2-3 task families for the first deep automation. • Lock metrics and governance.	• Operating-model blueprint • Phased roadmap • Investment case for executive approval
Deploy	60-90	• Stand up Layer 1 AI for one channel and one task family. • Stand up Layer 2 co-pilot for one queue. • Instrument the new metrics.	• First production proof points; • Live data on deflection, AHT, CSAT • Foundation for scale

Two Things to Avoid

First, do not start by picking a vendor. Pick a target operating model and fit your vendors into it. Picking the vendor first is the single most common cause of stalled transformations.

Second, do not run this as an IT program. The contact center P&L owner must own it, with IT and AI as core partners — but the strategic and operational accountability must sit with operations.

Ready to get started?

Our experts work with senior CX, COO, and transformation leaders to design and execute Zero Contact Center programs. If you would like to compare your current state against the maturity model, or pressure-test your existing AI roadmap, contact us at information@coforge.com or **visit us online**.

About the author:

Ajay Singh specializes in customer experience (CX) transformation, contact center operations, and AI-led enterprise solutions. He works closely with senior CX, COO, and transformation leaders to design and implement next-generation operating models, including Zero Contact Center frameworks and AI-driven automation strategies.

With deep expertise across business process outsourcing (BPO), conversational AI, back-office operations, and digital transformation, Ajay has led large-scale initiatives focused on improving operational efficiency, enhancing customer experience, and driving measurable business outcomes.

He brings strong experience in strategic account management, solution development, and complex deal structuring, partnering with global organizations to modernize service delivery through innovation, analytics, and process re-engineering.

Ajay Singh

Vice President, CX Transformation & Solutions, Travel, Transportation & Hospitality



About Coforge

Coforge is an AI-native engineering services leader, where AI is the very foundation of how we design, build, and deliver intelligent solutions for our clients. We use AI and hyperspecialized industry expertise to engineer autonomous enterprises. We combine AI agents with our AI-enabled workforce, including specialized FDEs in hybrid pod-based delivery units. With a deep focus on trusted AI, our solutions are secure, governed, and enterprise-grade. We are outcome-led by design. Moving beyond AI experimentation, we deliver measurable business outcomes – lower operating costs, faster cycle times, higher conversion rates, and sustained margin growth.

Learn more: www.coforge.com

For more information, contact information@coforge.com

Coforge