

Dear KODIF community,

2025 was a year in which many long-held assumptions across e-commerce and customer support were tested in production.

- AI in CX became louder.
- New "agents" launched weekly.
- Demos became increasingly polished.
- Claims grew broader and more confident.

As AI moved from pilots to production, its impact extended far beyond internal CX teams. BPOs began redefining their role as AI absorbed more volume. Software vendors faced pressure as human-centric workflows started to break down. Systems that once sat at the center of operations increasingly became peripheral, while new intelligence layers began to emerge.

This shift is not simply about replacing humans with AI. It was about who owns the work, who controls outcomes, and where strategic value now lives.

In this annual letter, we share our point of view on how the customer support ecosystem is evolving - across brands, BPOs, platforms, and technology partners - and why we believe the next phase of customer support will be defined less by tools and headcount, and more by intelligence, orchestration, and operational ownership.

What follows is not theory. It is what we saw when AI met real customers, real volume, and real accountability.

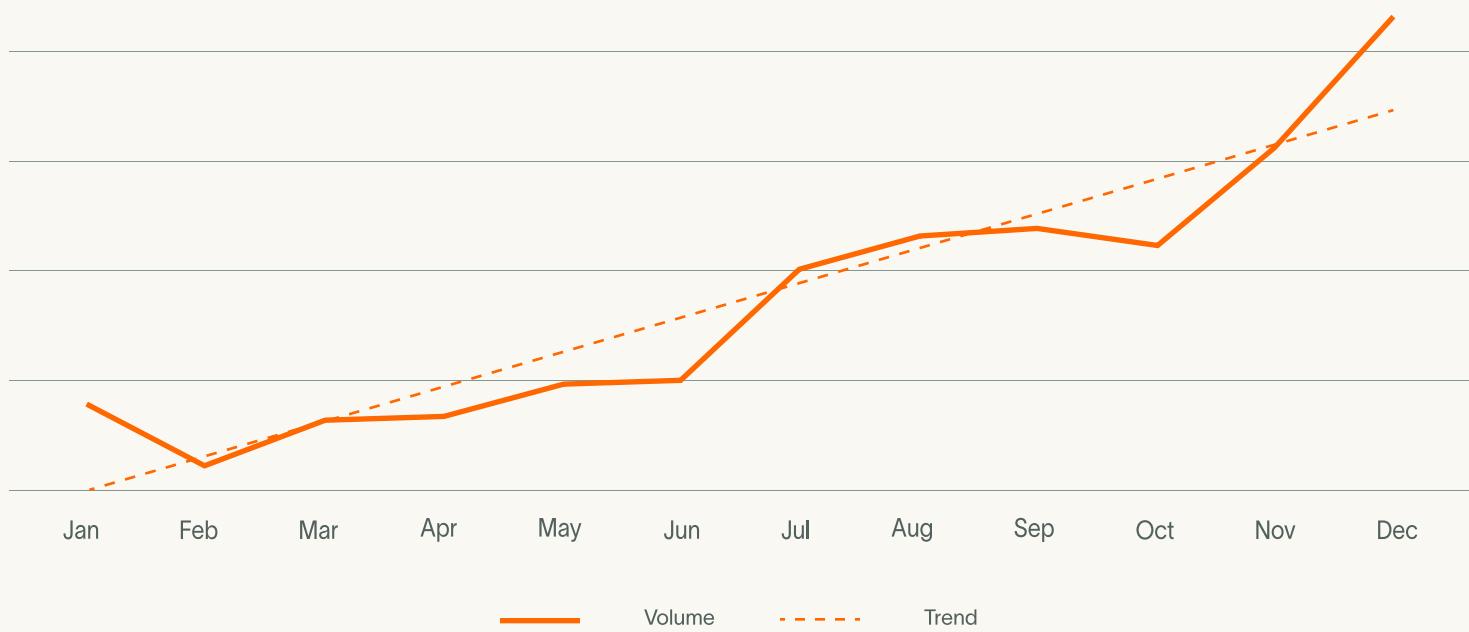
2025 - A Year of Focus and Depth

In 2025, KODIF handled more than **6 million customer conversations end-to-end**, representing significant year-over-year growth. Over the same period, KODIF **tripled ARR** and began approaching profitability.

2025 Growth:

- Overall Conversation Growth - 2.8x
- SMS Automation Growth - 8x
- Email Automation Growth - 5.6x
- 88 third party integrations built - 6x
- 2000+ Agentic tools - 7.5x

2025 Monthly Resolved Conversation Volume



While growth mattered, it was not the most important business outcome for KODIF this year. Early in 2025, we chose **focus over optionality**. We made e-commerce our core domain.

Not because other industries are unimportant, but because e-commerce is uniquely demanding. It is fast-moving, margin-sensitive, operationally brittle, and deeply interconnected. The customer journey in this environment is not theoretical - it is relentlessly practical.

What became clear over time was the **compounding effect of vertical focus**.

Earlier, we tried to sell broadly. The intent was reasonable: maximize surface area and preserve flexibility. In practice, this diluted momentum and slowed learning. By narrowing our focus, we unlocked a different growth dynamic. Reputation began to compound within a single ecosystem. KODIF became known not as a general-purpose AI vendor, but as a system purpose-built for e-commerce support. Reference calls became more credible. Recommendations began to travel laterally - from operator to operator, brand to brand - within the same vertical.

Focus also sharpened internal execution. Hiring sales and solutions leaders with deep e-commerce experience shortened sales cycles and improved qualification. Conversations moved faster because shared context already existed. Objections became more precise - and more solvable.

Over time, this created a flywheel:

- clearer positioning
- stronger brand recognition within the vertical
- higher-quality inbound from similar operators
- denser reference networks
- more disciplined product decisions

The result was not just faster growth, but **more efficient growth**. This reinforced a broader lesson from 2025: in operationally complex markets, **focus is not a constraint - it is a force multiplier**. We found that depth compound and complexity compound. In the e-commerce both of these compound faster than anywhere else.

Why Agentic Platforms Are Becoming the Strategic Core

In 2025, agentic systems have moved from experimentation to production - and are now operating at scale. This marks a fundamental shift in how work gets done.

For decades, **the atomic unit of knowledge work belonged to humans**. Software existed to enable people to do that work better. Ticketing systems routed tasks. Workforce management scheduled humans. QA tools evaluated human output. Knowledge bases helped humans find SOPs. VoC platforms analyzed what humans produced. For the first time, this relationship is flipped.

AI agents are now doing the work - and humans are enabling the software.

Nowhere is this more visible than in customer support. Historically, BPOs supplied labor, and software existed to maximize human productivity. But modern agentic platforms can already perform the full support lifecycle end-to-end: execution, routing, quality control, knowledge creation, and continuous optimization. This shift is not incremental - **it is structural**.

Pre-AI Support Stack

Enhance human productivity



AI-Native Support Stack

Deliver completed work

AI Platform

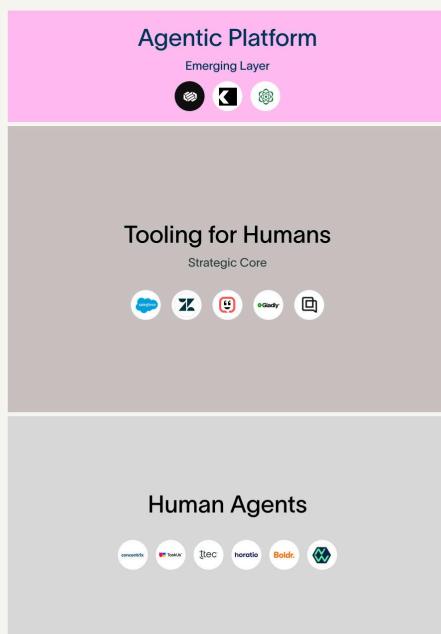
- AI Quality assurance
- AI + Human workforce management
- AI Voice of the Customer
- AI Agents
- AI Knowledge management

 Tier 3 Support
(With AI Handoff)

As a result, strategic value is moving away from labor providers and human-centric software toward agentic platforms. BPOs and traditional ticketing systems increasingly become commodities, while agentic platforms become the strategic core of operations.

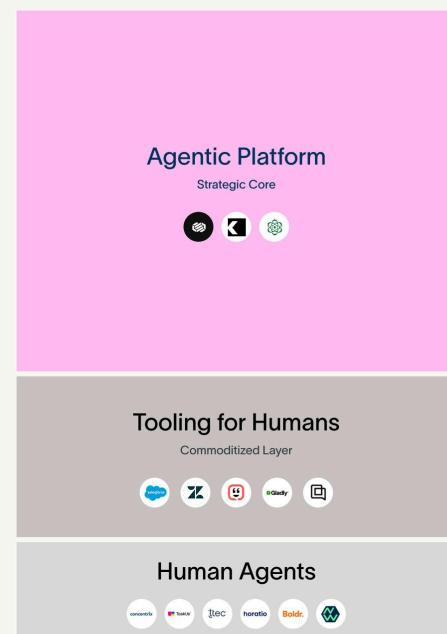
Today

Tooling for Humans Dominates



Tomorrow

Agentic Platforms Dominate



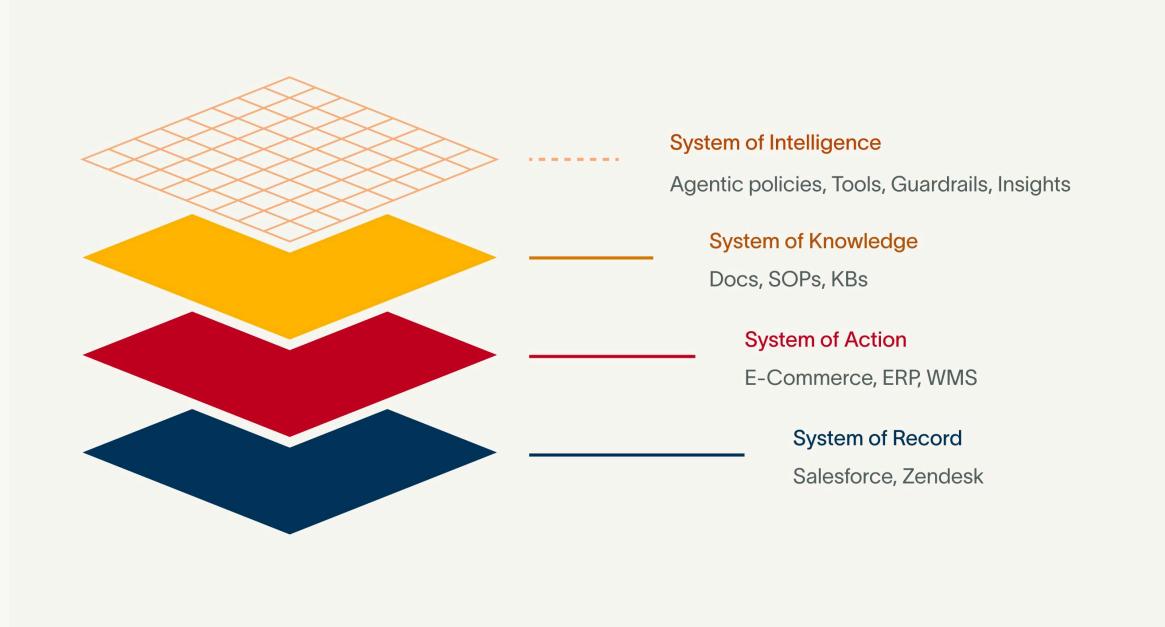
In a typical support stack today, work flows across three layers: System of Record (e.g., Salesforce, Zendesk); System of Action (e.g., Shopify, ERP, WMS); System of Knowledge (docs, SOPs, KBs)

We believe a fourth layer is emerging - and will become the most important. KODIF is the System of Intelligence. It is the intelligence layer that:

- Holds and operationalizes institutional knowledge
- Continuously learns from every interaction
- Improves autonomously over time
- Orchestrates actions across systems
- Becomes smarter as volume and complexity increase

Agentic platforms will increasingly subsume and demote legacy systems, owning the most strategic layer of the stack: intelligence itself.

This is our long-term vision for KODIF - and the future we are building toward.



The State of Support in 2025: Complexity Over Volume

In 2025, support isn't defined by volume, it's defined by complexity.

Customer Support is evolving from a high-volume ticket-handling function into a sophisticated automation organization. Tickets increasingly involve **policy interpretation rather than factual questions, partial or conflicting system states, emotionally charged escalations tied to delays or errors, and cascading issues triggered by small upstream changes**. The work has gotten harder and simply hiring more headcount is no longer a sustainable solution, both economically and operationally.

With the rise of large language models and agentic systems, straightforward question-answering has become table stakes. And in many organizations, multi-step automations are no longer experimental: they're running in production, handling real customer interactions end-to-end. As the industry clears that bar, a new set of challenges is coming into focus: operational coordination, system constraints, edge cases, and the day-to-day realities of running support at scale.

KODIF's view is that the next wave of progress won't come from replacing humans with a single "super-bot," but from emulating how high-performing support teams actually operate. That's why in 2025 we introduced the concept of [AI teammates](#), agents designed to mirror real support roles and responsibilities, enabling companies to map automation to familiar operational language: handoffs, escalation, ownership, quality control, and collaboration.

[Integration depth](#) is also becoming a primary differentiator. Effective automation increasingly depends on what the system can reliably do, not just what it can say. The teams that win will be the ones that are vertically and deeply integrated into the support stack, and able to take precise actions across the systems of record, enforce policy, and navigate real workflows with minimal friction.

As workflows become more capable, [observability and explainability](#) become non-negotiable. Organizations need to understand what happened, why it happened, and what the system will do next, both to build internal trust and to operate safely under real business constraints.

Finally, adoption at scale requires that automation be [self-service](#). If authoring and iterating on workflows is gated behind a vendor or specialized teams, companies will struggle to keep pace with change. The future belongs to platforms that allow support teams to own their automations quickly, safely, and with confidence.

What Broke: The Limits of Answer-First AI

A consistent pattern emerged across the industry: AI systems designed primarily for [answer generation](#) performed well in controlled environments and degraded quickly in production.

They struggled to reconcile data across multiple systems or knowledge sources, apply nuanced or evolving business rules, take actions accurately and ensure data consistency, escalate with context and intent and adjust to changing customer tone or need.

In many cases, teams concluded that "AI is not fully ready for CX." This is the wrong framing. As we worked closely with many teams, one pattern kept surfacing: CX is not a single job. It is a collection of distinct roles with different workflows, risk profiles, and success metrics.

- Shopping assistants guide customers toward the right products and drive conversion.
- Support agents resolve issues and execute operational workflows across systems.
- Analysts surface patterns, friction, and opportunities across conversations.
- Managers continuously maintain knowledge bases, refine policies and processes, and optimize operations over time.

AI is already strong in some of these roles and weaker in others. The mistake was treating CX as a monolithic problem rather than a portfolio of jobs.

AI Teammates: from “Bots” to Systems of Work

This observation led to a conceptual shift in 2025: High-functioning support organizations already operate as systems that handle **frontline resolution, escalation paths, quality assurance, insight generation, continuous improvement loops**.

AI systems that attempted to collapse these roles into a single interface consistently underperformed. As a result, CX teams tend to buy separate solutions for individual systems.

This year, KODIF started framing AI, not as a single interface, but as a team: **AI teammates** designed around distinct responsibilities. These teammates have different prompts, goals, personas, and data, all of which should be measured and acted upon differently.

Effective AI in CX needs to resemble a **coordinated system of teammates and tools**, including human and AI Agents using their unique strengths to provide a quality customer experience.

Depth Over Breadth

In practice, support outcomes are highly operational: containment and deflection, resolution time, backlog stability, refund and replacement accuracy, CSAT and long-term trust preservation.

Improving these metrics requires systems that can read from multiple sources of truth, interpret policies consistently, execute workflows end-to-end and fail predictably and recover cleanly.

By the end of 2025, a clearer distinction emerged across the market: AI that can only answer customer questions is **assistive**. AI that can deliver meaningful outcomes is **operational**. Support work depends on coordinated actions across CRMs, helpdesks, OMS, subscriptions, fulfillment providers and many others.

The challenge is not the number of integrations – it is their **semantic depth**. Natural language must reliably map to real business actions. Without deep integration, automation remained brittle and stalled.

That's why we invested effort into creating integrations to 88 new platforms (and growing) used within the common e-commerce tech stacks so that KODIF allows more than 2000 automation actions which can be combined to operate against any backend set up, improving self-service, accuracy, and reducing manual human agent effort.

E-commerce vertical surfaced these failures earlier than other domains for structural reasons. E-commerce stacks are fragmented for a variety of reasons. Order management, subscriptions, fulfillment, loyalty, payments, marketing platforms, and CRMs often each hold a partial state and are owned by disparate teams. Policies change frequently, promotions introduce ambiguity and small misconfigurations can trigger thousands of downstream conversations.

As a result, shallow integrations fail quickly, generic automation degrades under pressure and exception handling dominates human workload.

This reinforced a broader pattern visible throughout 2025: **Horizontal AI systems struggle under vertical operational pressure**. Depth matters more than breadth.

If You Can't Explain It, You Can't Trust It

As AI adoption expanded through 2025, skepticism expanded with it. Early wins in automation created momentum but also raised the bar for operational accountability. Support leaders began asking questions that weren't philosophical. They were practical, day-to-day, and unavoidable:

- Why did this response happen?
- What data was used?
- Which rule or policy applied?
- What should we change to correct this behavior and how do we verify it?

In production support environments, "because the model said so" is not an acceptable answer. Trust isn't rebuilt through generic explanations or polished reasoning. It's rebuilt through **observable, explainable, testable, and correctable behavior**.

By 2025, one pattern became clear across the market: **opaque AI systems are incompatible with real support operations**. If a brand can't trace what the system saw, what it decided, and what it did, the automation can't be safely operated, especially when outcomes affect refunds, replacements, delivery expectations, and customer trust.

Operational trust requires more than transparency as a principle. It requires transparency as a workflow:

- **Testing against real conversations** to validate behavior under realistic edge cases
- **Clear failure classification** so teams know whether an issue is data, policy, tool execution, or model behavior
- **Fast feedback loops** that shorten the time from "this was wrong" to "this is fixed"
- **Explicit correction paths** that make it clear how to retrain, retest, or redesign the decision entirely

The standard for production AI is simple: every response, action, and decision should be **human-readable, traceable, and actionable**, with clear direction on how to change the underlying behavior.

This is the core shift of 2025. **AI transparency isn't optional: It's a prerequisite** for confidence, internal buy-in, and safe scale.

At KODIF, we prioritized features that make automations transparent in practice, not just in theory, so teams can understand what happened, adjust behavior quickly when needed, and build confidence that grows with usage rather than collapsing under scrutiny.

Self-Service Policy Authoring as the Adoption Lever

Adoption at scale requires that automation be **self-service**. If authoring and iterating on workflows is gated behind a vendor, forward-deployed engineers, or a specialized internal team, organizations can't keep pace with the rate of change. New promotions, policy updates, fulfillment exceptions, and tooling changes quickly outstrip centralized capacity. The future belongs to platforms that let support teams **own their automation**: update behavior quickly, safely, and with confidence.

This need is especially visible in fast-moving e-commerce brands, where a handful of recurring workflows dominate support volume: **order fulfillment and modifications, delivery issues, returns and warranties, subscriptions, and appointments**. Because these workflows change often, the operational requirement isn't just "automation," but configurable automation: teams must be able to operate, configure, and evolve workflows themselves without waiting on engineering cycles. Where organizations achieved this, automation rates increased materially. See the case studies for [JustFoodForDogs](#) and [Dollar Shave Club](#). Crucially, these outcomes were driven less by prompt tuning and more by **deep integration and operational self-service**, the ability to translate policy and intent into reliable, editable workflows that map to real business actions.

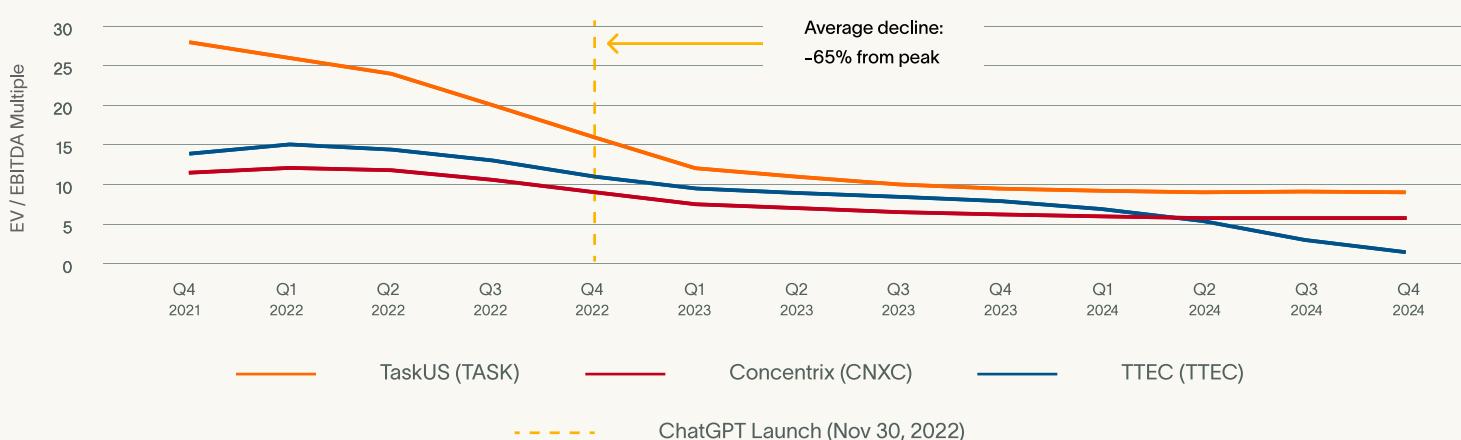
BPOs in the Age of AI Agents: From Headcount to Solution Providers

While many expected AI to significantly displace traditional outsourcing models, BPOs remained structurally resilient through 2025. But their role evolved. As AI absorbed more volume, human operators increasingly focused on exception handling, quality assurance, policy interpretation and training and feedback.

In this model, **BPOs function less as labor arbitrage** and more **as operators of hybrid systems**. The divide is no longer AI vs. humans - it is between organizations that can run AI-augmented operations and those that cannot.

The economics of the BPO industry are changing fast, and not subtly. BPOs are seeing the public markets price in the AI disruption that is almost certain in the short term. Since ChatGPT launched in 2022, public BPOs have seen an average EBITDA multiple drop by roughly 65%!

BPO Valuation Multiple Compression Since ChatGPT Launch



Agentic systems may be new - but disruption is not. History repeatedly shows what happens when a new operating model replaces an old one.

Netflix dismantled the video rental market, and Blockbuster - a behemoth with over 9,000 stores - went bankrupt. Uber redefined transportation economics: taxi medallion prices peaked at \$1.05M in 2014, then collapsed by more than 86%, reaching roughly \$150K by 2025.

The same structural forces are now at work in the BPO market. As work shifts from humans to AI agents, the value of labor-centric models compresses. While incumbents are adapting, for example, Wipro COO Sanjeev Jain recently noted a 140% increase in AI adoption across existing projects - the direction is clear. Adoption is accelerating, but the center of gravity is moving. As with past disruptions, efficiency gains are not enough to preserve long-term value when the underlying unit of work fundamentally changes.

There are two main reasons why markets don't believe that BPOs can adapt and transform their businesses:

1. **BPOs and Agentic businesses are built on fundamentally different economics.** BPOs sell people with a 20-30% markup, while agentic companies sell software that replaces people. If a BPO fully embraced an AI-first, product-led model, it would compress margins and cannibalize its core revenue—something that's nearly impossible for a public company to justify to the market.
2. **Talent is the second hard constraint.** Staying competitive in AI requires dedicated, Agentic teams that can continuously absorb fast-moving research, evaluate new models and tools, and translate them into production-grade systems. This depth of technical and product expertise is rare, highly specialized, and structurally misaligned with how BPOs hire, operate, and build capabilities.

This creates a situation where companies are able to grow revenue, while enterprise value drops. The market is making a clear statement that valuations built on headcount growth are not coming back. In an agentic world, BPOs cannot simply grow their way back to old multiples using the same labor-based playbook. The only way for BPOs to stay relevant and competitive is to transform from Headcount to Solutions Providers.

However, this transformation requires a deliberate strategy to move from service providers to full solutions providers. There are three things BPOs can do, and many have started to do, to begin this transformation.

1. **Change the measure of success:** Adding headcount no longer equals success. Occupancy, utilization, and seat growth for BPO clients are no longer the goal. The agentic model means that success requires resolving more issues with less human labor.
2. **Quality and Confidence:** Change the goal to handling volume and measuring the quality score of your agents to understanding why a ticket was created in the first place. Build confidence that you can help prevent issues, understand policy or knowledge gaps faster, and automate the gaps quickly and easily. This will enable higher quality products and services for the BPO clients, as well as strengthen confidence in your ability to provide end-to-end CX Solutions.
3. **Enable Operations:** While the C-Suite must set the tone and direction for this transformation, the very nature of agentic solutions is that they are bespoke to the individual client's needs. There has to be a recognition that operational teams understand, recognize, and are empowered to find technology solutions given their direct oversight of the work. They need to know when AI Agents can safely take over and when human experts are required and then be given options to add or augment the technology solutions in real time.

Last but not the least, BPOs now must decide their approach to transform to Solutions Providers. While building or buying the technology is an option, these often require large scale cultural, staffing, and investment shifts which is a significant, time consuming effort which will fundamentally alter the organization. Partnership offers a faster start and more flexibility and stability. As our CTO, Norm Usenkanov puts it "For a BPO, the most attractive thing isn't just "cool tech"; it's a platform that can enable your talented agents to be more productive where it counts, empathetic reasoning, dealing with ambiguity, or trusted motivation and advocacy" while making the remaining headcount 2x more productive." The right technology partner will work with you on providing the technology the market is demanding, training your team, and enabling new revenue streams for your business.

And the good news is that today's BPOs already operate across much of what AI agents now do end-to-end. They execute the work by responding to tickets, manage and optimize processes, maintain and update knowledge bases, and run quality assurance to ensure consistent outcomes. An agentic platform doesn't replace this model - it amplifies it. By embedding AI agents into their operations, BPOs can scale execution, standardize quality, and continuously improve performance with far less manual effort. That's why we believe partnership offers the fastest path to deploying agentic platforms at scale today. Over time, as intelligence becomes core infrastructure, we expect leading organizations to tightly integrate or internalize these capabilities to maintain strategic control and differentiation.

Looking Ahead to 2026

There are many important problems to work on, but we're continually reminded how fortunate we are to play even a small role in helping teams serve their customers better. The companies we work with are operating in increasingly complex environments, and every day we see how thoughtful use of AI can remove friction, unlock focus, and create better experiences for both customers and operators. To the teams who have trusted KODIF with real workflows and real decisions - **thank you**. Your ambition, feedback, and willingness to build with us are what make this work meaningful.

If this letter sparked your interest in how AI is reshaping customer operations, we look forward to continuing the conversation throughout the year - through customer conversations, partner collaborations, and live discussions with operators building the future of CX. We'll be sharing more on where we're headed, what we're learning, and how AI teammates can take on more real work, end to end.

Until next year,
Chyngyz, Norm, and Mike