2026 BPO, AI & CX Intelliview:

From Labor Arbitrage to Managed AI Operations

opusresearch

November 2026

Ian Jacobs, VP & Lead Analyst, Opus Research

Derek Top, Research Director & Principal Analyst, Opus Research

Amy Stapleton, Senior Analyst, Opus Research

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

www.opusresearch.net

Published November 2026 © Opus Research, Inc. All rights reserved.

opusresearch | Report »

2026 BPO, AI & CX Intelliview:

From Labor Arbitrage to Managed AI Operations

Generative and agentic AI are rewriting the economics of customer experience outsourcing. Opus Research's 2026 BPO, AI & CX Intelliview evaluates how six CX BPOs (Concentrix, Firstsource, Foundever, iQor, TELUS Digital, and TP) apply AI inside their own delivery operations and in the services they sell to clients. Using a two-axis BPO AI Maturity Model, a six-dimension AI operating model profile, and the Intelliview Map, the report finds that no provider has yet reached AI-native delivery, that BPOs sell AI better than they run it autonomously, and that proof of value and accountability for failure remain the weak points.

November 2026

Ian Jacobs, VP & Lead Analyst, Opus Research

Derek Top, Research Director & Principal Analyst, Opus Research

Amy Stapleton, Senior Analyst, Opus Research

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

www.opusresearch.net

Published November 2026 © Opus Research, Inc. All rights reserved.

2026 BPO, AI & CX Intelliview:
From Labor Arbitrage to Managed AI Operations

Table of Contents

Key Highlights 4
Introduction: The BPO Dilemma 4
The BPO AI Maturity Model 6
The AI Operating Model Profile 8
- Six Dimensions of the AI Operating Model 9
- Reading the Results 10

The 2026 BPO, AI & CX Intelliview Map 11
- Leaders 12
- Standouts 12
- Sidebar: Crescendo 12

The Proof Problem 13
Where Does BPO AI Differentiation Come From? 13
Who Owns the AI, and Who Owns the Failure? 14
How the Frontline Job Is Changing 15
Questions to Ask Your BPO 16
Methodology 16
Appendix A: Company Dossiers 18
- **Concentrix** 19
- **Firstsource** 22
- **Foundever** 25
- **iQor** 28
- **TELUS Digital** 31
- **TP** 34
- Reference Profile: **Crescendo** 37

Table of Figures

Figure 1: Primary Strategies by Provider 6
Figure 2: The BPO AI Maturity Model 7
Figure 3: Maturity-Stage Definitions by Axis 8
Figure 4: AI Operating Model Profile Scores 10
Figure 5: 2026 BPO, AI & CX Intelliview Map 11
Figure 6: Dividing AI Control Plane Responsibilities 15

Key Highlights

- **Leaders and Standouts.** TP, iQor, Foundever, TELUS Digital, Firstsource, and Concentrix* are leading solution providers highlighted in the 2026 BPO, AI & CX Intelliview Map, each with distinctive strengths buyers should weigh against specific needs.

- **BPOs sell AI better than they run it autonomously.** Client-facing AI services is the cohort's strongest dimension, with autonomous customer service the weakest.

- **Proof is the market's weak point.** Among the RFI respondents, reported containment ranged from roughly 35% to 75%, two reported no containment figure at all, and definitions differ from provider to provider. Buyers should insist on baselines, attribution, and evidence that results repeat across accounts.

- **Workforce change trails the rhetoric.** Every provider says frontline work is shifting toward exceptions and judgment calls, but transforming the workforce to maximize AI value is still in progress.

- **Accountability for AI failure has to be designed in.** When the client, the BPO, and the technology vendor each control part of the AI system, ownership of errors, handoffs, and remediation must be assigned before launch, not after the first failure.

Introduction: The BPO Dilemma

Over the past few decades, CX BPOs built an enormous business around a straightforward proposition. They could recruit, train, and manage people to handle customer interactions more efficiently than their clients could. Labor arbitrage supplied much of the economic engine. Operational discipline and excellence kept it running.

Modern AI puts pressure on both parts of that equation. Agent assistance can reduce the time and expertise required to complete work. Meanwhile, autonomous agents can take over portions of the work itself. For a provider paid by the hour or the seat, helping a client need fewer people creates an awkward sales conversation. Success can shrink the customer invoice.

The remaining human work also gets harder. As automation absorbs routine requests, agents inherit more exceptions, emotionally charged conversations, and problems that cross systems or departments. Those interactions demand different skills, more support, and performance measures that recognize complexity. Yesterday's staffing assumptions won't necessarily survive tomorrow's contact mix.

Clients also have more options. Software vendors, implementation partners, and AI specialists can compete for pieces of the customer service budget. A BPO's ability to supply thousands of trained agents becomes less persuasive when the buyer wants to know how many interactions can be resolved without one.

Yet deploying AI creates substantial new work. Someone has to connect systems, maintain knowledge, define authority, test behavior, manage escalations, and establish whether the results justify the cost. BPOs already understand many of the operational problems involved. However, turning that familiarity into a service clients will pay for takes investment and demonstrable expertise.

They also face a timing problem. Providers must fund new capabilities while supporting businesses that still depend on people. Move too slowly and competitors capture the opportunity. Move quickly without changing the commercial model and efficiency gains can outrun replacement revenue.

BPOs have several ways to respond. Most will combine approaches, although each makes different demands on their talent, investment priorities, and appetite for risk. They can:

- **Build AI products.** The desire among BPOs to build a real software business is nothing new. But today, the idea is to develop proprietary agents, copilots, analytics, or orchestration software. These products can generate revenue beyond staffing contracts and give clients a reason to stay. They also require sustained engineering, support, and product management. A branded interface alone won't create much protection from competitors.
- **Become the deployment and integration expert.** Here, providers move into systems integrator territory and help clients select, connect, configure, and operate third-party AI. The opportunity extends well beyond installation into evaluation, knowledge maintenance, and workflow improvement. Providers need to show why their operating experience produces better results than another implementation partner's technical skills.
- **Go deeper into specializations.** Some BPOs are doubling down on operational expertise, but through a narrower aperture. They combine AI with expertise in healthcare, financial services, collections, multilingual delivery, or other demanding work. Domain knowledge can shape better workflows and tighter controls. Credibility depends on how that expertise changes delivery, especially when an interaction gets complicated.
- **Make internal operations the advantage.** Many BPOs are turning AI on themselves. They use AI to improve recruitment, training, coaching, quality management, scheduling, and agent support while clients retain control of consumer-facing technology. This approach can strengthen delivery and margins, although contract negotiations will determine how much of the productivity benefit the provider keeps.
- **Manage the combined human and AI operation.** Playing up the strengths they have in managing human labor, many BPOs are taking responsibility for performance across automated service, human agents, and the handoffs between them. That creates opportunities in shared forecasting, quality standards, escalation design, and governance. It also gives the provider responsibility for failures that cross organizational and technology boundaries.
- **Change what clients buy.** BPOs are at the heart of the pricing debate around AI. How to charge for AI, let alone AI plus humans, remains an unsettled question. Some BPOs are looking to charge for managed capabilities, completed work, or agreed outcomes alongside traditional staffing. Providers can also use interaction insights to help clients remove recurring service problems. Both approaches require clear measurement and contract terms.

Figure 1: Primary Strategies by Provider

Provider	Primary strategies	How it shows up
Concentrix*	Build AI products; become the deployment expert	iX Hello and iX Hero licensed separately from services; Agentic Operating Framework with Forward Deployed Engineering
Firstsource	Become the deployment expert; build AI products	Transform–Implement–Operate methodology; Kairos AI operating system
Foundever	Manage the combined operation; change what clients buy	Operator Pods and Agentic Pods run human and AI workforces together; Bots-as-a-Service consumption pricing
iQor	Make internal operations the advantage	Valdi, Predictive NPS, and Coach iQ built on its own operating data; partner technology for conversational AI
TELUS Digital	Build AI products; become the deployment expert	Fuel iX tooling in every new engagement; Forward Deployed Engineers embedded with delivery teams
TP	Build AI products; manage the combined operation	TP.ai orchestration and a Trust & Control Layer spanning in-house and partner AI

The BPO AI Maturity Model

Opus Research created a BPO AI Maturity Model that assesses how deeply AI has changed a provider's delivery operations and the services it offers clients. We intend this to help buyers distinguish between adding AI tools to established practices and building an operating model in which people and AI share responsibility for delivering customer and business outcomes.

The vertical axis measures internal operations maturity. It examines, for example, how the BPO uses AI within its own service delivery, including agent assistance, quality management, coaching, training, workforce planning, and operational improvement. Higher maturity requires more than widespread tool adoption. It involves connected workflows, redesigned roles, shared performance measures, and clear accountability for managing human and AI work together.

The horizontal axis measures customer-facing offering maturity. It examines what the BPO brings to clients, including conversational automation, AI implementation, journey redesign, and ongoing management of AI-enabled services. "Customer-facing" refers to the provider's client offering, which can include both employee-facing tools and AI that interacts directly with consumers. Higher maturity here reflects the provider's ability to integrate these capabilities into repeatable services with demonstrable outcomes.

Both axes use four stages. AI-assisted delivery improves individual tasks within existing processes. AI-integrated operations connects capabilities across workflows and embeds them in routine delivery. AI-led transformation

changes processes, workforce structures, and service models around AI's capabilities. AI-native delivery treats human and AI work as parts of a holistic and unified system, with coordinated governance, measurement, and continuous improvement. AI-native describes how delivery operates, regardless of when the company was founded.

Readers should examine the relationship between the two placements. A provider may have advanced internal practices while offering relatively conventional services to clients. Another may market sophisticated AI services that aren't yet deeply embedded across its own operations. The differences and gaps can spark useful questions about implementation experience, consistency, and the transfer of capabilities across accounts.

In this cohort, the gaps run in both directions. Foundever and iQor are further along inside their own operations than in what they sell to clients. Concentrix, Firstsource, and TELUS Digital show the reverse, with client offerings ahead of their internal redesign. TP is the only provider at stage 3 on both axes. No provider reaches stage 4 on either axis, and none sits at stage 1. The market remains focused on integration and early-stage transformation rather than fully AI-native delivery.

Figure 2: The BPO AI Maturity Model

Internal operations maturity	Customer-facing offering maturity: 1 AI-assisted delivery	2 AI-integrated operations	3 AI-led transformation	4 AI-native delivery
4 AI-native delivery				
3 AI-led transformation		Foundever iQor	TP	
2 AI-integrated operations			*Concentrix** Firstsource TELUS Digital	
1 AI-assisted delivery				

(Names are alphabetical within cells. Evaluations are based on RFI responses and briefings. *Concentrix is assessed from the public record. Crescendo is assessed separately.)

Placing providers on the matrix. The BPO AI Maturity Model places each provider on two axes, internal operations and customer-facing offerings, across four stages. Stage judgments reflect demonstrated operating practice and client delivery, with uneven rollout and unresolved claims taken into account. Broad tool deployment alone doesn't establish stage 3, which requires redesigned roles, metrics, or commercial models. Stage 4 requires unified governance, measurement, and improvement across human and AI work. Partnering with an AI-native technology provider doesn't by itself establish stage 4.

Figure 3: Maturity-Stage Definitions by Axis

Stage	Internal operations	Customer-facing offerings	Evidence we looked for
1 AI-assisted delivery	Point tools support agents, QA, training, or workforce management.	AI features supplement conventional outsourced services.	Adoption, production coverage, and measured task-level gains
2 AI-integrated operations	Multiple functions share data, automation, and operating processes.	The BPO implements and manages AI across several customer journeys.	Integration depth, scale, operating ownership, and repeatability
3 AI-led transformation	Roles, staffing, metrics, and management practices are redesigned.	The BPO leads transformation and accepts broader outcome accountability.	A new operating model, attributable outcomes, and outcome-based engagements
4 AI-native delivery	AI performs or manages substantial work under unified governance and optimization.	AI-led services are the primary delivery model, not an added capability.	End-to-end autonomy, unified workforce controls, and durable economics

We believe the matrix should inform both a shortlist and a briefing agenda. Buyers should match maturity to their intended scope, risk tolerance, and capacity for change. A higher placement doesn't automatically indicate the best partner for every engagement. For the providers highlighted in this Intelliview, readers should use the accompanying provider dossiers to test deployment scale, evidence quality, governance, and client outcomes, and to determine how broadly the assessed maturity applies. Beyond this cohort, the model gives buyers a consistent way to compare any provider they are considering.

The AI Operating Model Profile

The maturity matrix answers one question: how far along is each provider? It doesn't answer the question buyers usually ask next: what kind of AI BPO is this, and is it strong where I need it? Two providers can share a cell on the matrix and still bring very different capabilities to an engagement.

To capture that difference, we built an AI operating model profile. It scores each provider on six dimensions where AI creates value in an outsourced CX engagement. These cover the support AI gives agents and supervisors, the interactions AI completes on its own, how the workforce is organized, how operational data drives improvement, the AI services sold to clients, and the governance and controls that keep it all safe. The profile's shape, not its total, defines the operating model. A provider with a 5 in analytics and a 2 in autonomous service is a different partner from one with the reverse, even if their totals match.

We used this lens for three reasons. First, it reflects how buyers truly source. Few engagements need equal strength everywhere, and a buyer moving collections to voice AI needs something different from one rebuilding agent coaching. Second, it separates what a provider sells from what it runs. Strength in client-facing AI services doesn't guarantee depth in autonomous customer service, and this cohort shows exactly that gap. Third, it rewards operation, not ownership. Partner technology earns credit when the BPO shows it deploys and operates that technology effectively, so a provider that runs a third-party platform well isn't penalized for not building its own.

Six Dimensions of the AI Operating Model

- **Frontline augmentation.** AI that supports human agents and supervisors in live work: real-time assistance, knowledge retrieval, summarization, simulation-based training, and automated coaching. We looked at adoption across the agent base, how deeply the tools are built into the workflow, and measured effects on quality, proficiency, and productivity.
- **Autonomous customer service.** AI that completes customer interactions and takes action without a person in the loop, across voice and digital channels. We looked at production volume, task complexity, multistep completion, the quality of handoffs to people, and outcomes.
- **Workforce transformation.** Changes to roles, skills, staffing, and management driven by AI. We looked for new roles, revised training and career paths, changed hiring profiles, and evidence that human and AI work are managed together.
- **Analytics and optimization.** Using interaction and operational data to improve customer experience and delivery. We looked at coverage, whether insight reliably turns into action, attribution rigor, review cadence, and repeatability.
- **Client-facing AI services.** The AI capabilities a BPO brings to clients as a service: advisory, implementation, proprietary assets, and ongoing management of AI operations. We looked at breadth, proprietary intellectual property, the provider's responsibility for deployment and operation, scale, and repeatability.
- **Governance and control.** The controls that span AI systems and customer journeys: inventory of models and agents, policies, approvals, evaluation before deployment, monitoring, auditability, and the ability to intervene.

Each dimension is scored from 1 to 5:

1. Limited capability or evidence
2. Developing capability or limited production evidence
3. Established and repeatable production delivery
4. Strong integration and demonstrated outcomes
5. Exceptional depth, adoption, and repeatability

Figure 4: AI Operating Model Profile Scores

AI value dimension	Concentrix*	Firstsource	Foundever	iQor	TELUS Digital	TP
Frontline augmentation	4	3	4	4	3	4
Autonomous customer service	3	2	3	2	2	3
Workforce transformation	2	2	4	4	3	3
Analytics and optimization	4	4	3	5	3	4
Client-facing AI services	4	4	4	4	5	5
Governance and control	3	3	3	3	3	5

Scores run from 1 (limited) to 5 (leading). Providers in alphabetical order. *Concentrix is scored from the public record only. Crescendo is assessed separately and isn't scored.

Reading the Results

BPOs sell AI better than they run it autonomously. Client-facing AI services is the cohort's strongest dimension, averaging above 4, with TELUS Digital and TP at 5. Autonomous customer service is the weakest. BPOs have built credible practices for advising on, implementing, and managing AI, but few can yet point to complex customer journeys that AI resolves end to end at scale. That pattern also reflects the economic transition described earlier in this report. BPOs need new sources of revenue as AI reduces the amount of human labor required in some programs, so it makes sense that client-facing AI services are developing faster than fully autonomous delivery. In other words, the commercial incentive to sell AI is arriving faster than the operational ability to let AI run large portions of the service on its own.

Governance separates the leader. TP is the only provider to score 5 on governance and control, on the strength of a control layer that covers agent inventory, approval gates, change control, and observability across vendors. As BPOs take responsibility for AI they didn't build, control across vendors becomes a prerequisite, not a differentiator.

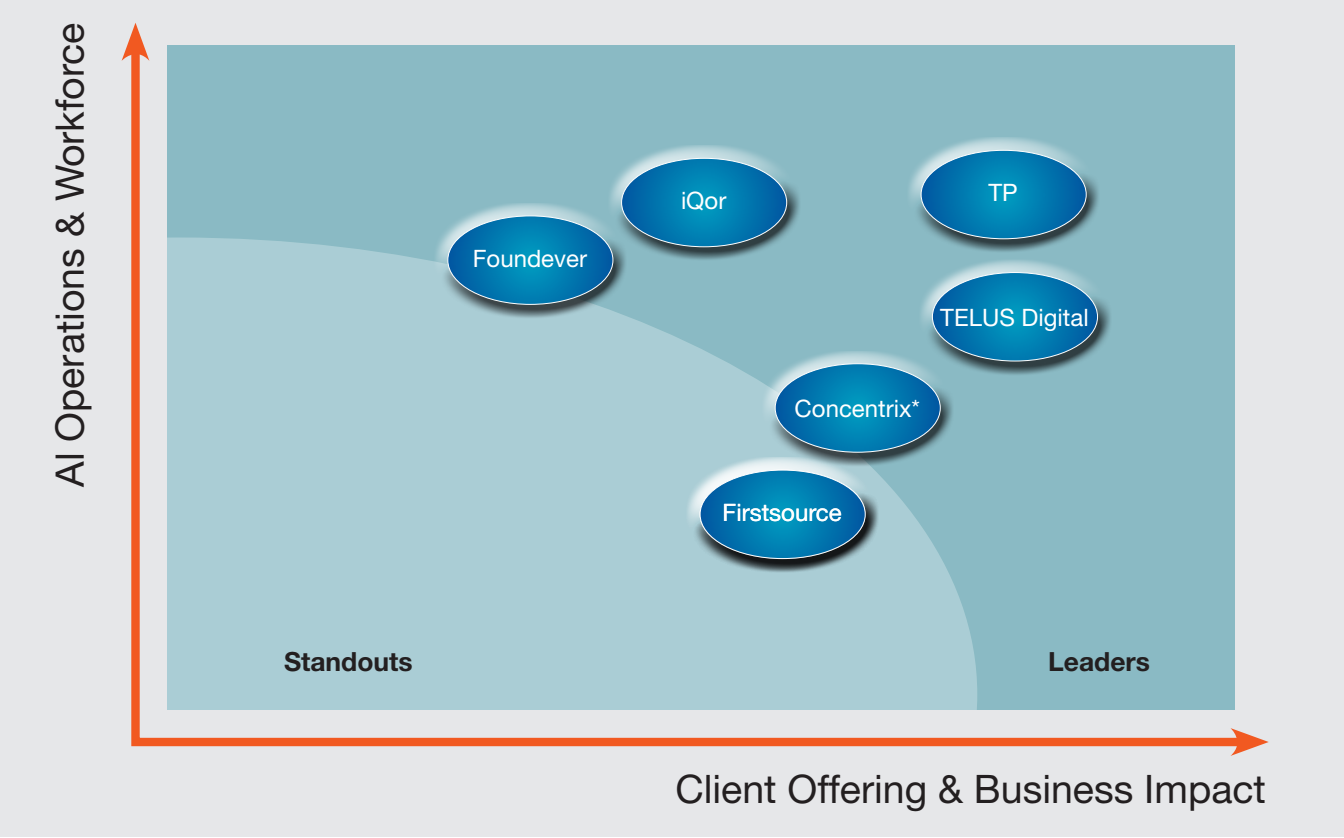
Peaks define distinct partners. iQor's 5 in analytics and optimization recognizes an analytics platform that has run in production for more than a decade. Firstsource also peaks in analytics and client-facing services. Concentrix scores 4 on frontline augmentation, analytics and optimization, and client-facing services, built largely on software it owns and licenses separately. Foundever has the most even profile among the Leaders, scoring 4 on frontline augmentation, workforce transformation, and client-facing services.

The 2026 BPO, AI & CX Intelliview Map

As in previous Opus Research Intelliviews, an Intelliview Map summarizes each provider's overall position (Figure 5 below). It combines the two lenses above. The vertical axis (AI Operations & Workforce) reflects what a provider runs inside its own delivery: its internal operations stage on the maturity model and profile scores for frontline augmentation, autonomous customer service, workforce transformation, analytics and optimization, and governance and control. The horizontal axis (Client Offering & Business Impact) reflects what it brings to clients: its customer-facing stage on the maturity model and its client-facing AI services score.

Providers fall into two groups. Leaders combine advanced operations with mature client offerings, or lead clearly on one axis while staying strong on the other. Standouts bring a distinctive strength that buyers should weigh against a specific need. Position within a group isn't a ranking, and several providers sit close to the boundary between groups.

Figure 5: 2026 BPO, AI & CX Intelliview Map



*Concentrix is assessed from the public record. Crescendo is assessed separately and isn't placed on the map.
Alorica was invited to participate in this Intelliview and provided an RFI response and briefing. Prior to publication, Alorica declined to participate in the final edition and is therefore not included in this assessment.

- **TP.** TP is the only provider at stage 3 on both maturity axes and the only one to score 5 on governance and control. Its TP.ai framework orchestrates people, AI agents, data, and workflows, and its Trust & Control Layer adds agent inventory, approval gates, audit trails, and cross-vendor observability.
- **iQor.** iQor is a clear specialist in analytics and optimization. Valdi, its operating-intelligence platform, feeds Coach iQ, Predictive NPS, and Insights iQ, and iQor reports about 35% of clients on outcome-based pricing. Its customer-facing automation runs largely on partner platforms inside clients' own stacks.
- **Foundever.** Foundever has the most even profile among the Leaders. Its dual-workforce model pairs AI-augmented agents with agentic "digital workers," managed through Operator Pods and Agentic Pods, and it is restructuring frontline work by complexity tier. It reports 57% chatbot containment for a US retailer and 74%+ voice containment for a financial services firm.
- **TELUS Digital.** For TELUS Digital, every new CX engagement includes a CX Strategy Assessment, Fuel iX tooling, and a Forward Deployed Engineer, and Fuel iX Fortify supplies automated red-teaming. Internal transformation trails its client offerings, although briefings showed progress, and customer-facing automation is largely partner-delivered.
- **Firstsource.** Firstsource stands out for its transformation posture. Its "UnBPO" strategy, Kairos AI operating system, and Transform–Implement–Operate methodology position it as an AI-first transformation partner, with high scores on analytics and optimization and client-facing AI services.
- **Concentrix*.** Concentrix, assessed from the public record, is distinguished by owning and licensing its AI. iX Hello and iX Hero are sold separately from services, with software ARR targeted at about $120 million by the end of FY2026, and its Agentic Operating Framework wraps them in a design-build-run model.

Crescendo

Crescendo responded to the RFI, but it isn't a traditional BPO. Publicly launched in January 2024, it sells an AI-native customer experience platform, reports hundreds of production AI deployments, and doesn't supply or price frontline labor. In October 2024, the company acquired PartnerHero, a mid-sized BPO. Running that operation gave Crescendo real-world proof of its technology and helped guide product development. But the company has since changed its approach and has begun transferring PartnerHero's agents to Alorica, a process that started after research for this Intelliview was underway. Accordingly, we are still including Crescendo in some parts of this report as a reference, but we are not including it in the scoring or the Intelliview Map.

For its technology business, Crescendo charges a recurring platform fee, tiered on volume, that includes the AI Deployment Engineer (AIDE) and AI Product Manager managed service, plus an AI usage fee per conversation, with no FTE or seat pricing. Because its model differs from the cohort's, we use it as a reference point for autonomous customer service rather than placing it on the matrix or scoring its profile.

The Crescendo-Alorica relationship also shapes how Crescendo goes to market. In June 2026, Alorica took a license to resell and deploy Crescendo's platform, agreed to co-develop an AI-native center of excellence, and gave Crescendo customers access to its human workforce. Taken together, these changes mean Crescendo customers can now contract human support through Alorica.

The Proof Problem

AI has made it easier for BPOs to make impressive claims and harder for buyers to determine what those claims actually mean. A provider may point to dozens of AI capabilities, hundreds (or even thousands) of deployments, or large improvements in efficiency. Those numbers can be useful, but they often leave important questions unanswered. Was the technology used in production or only in pilots? Was it deployed across a meaningful share of the provider's client base or concentrated in a handful of showcase accounts? Did the BPO build the capability, operate it, or simply resell technology from a partner?

The outcome claims require just as much scrutiny. A 20% reduction in handle time sounds impressive until you know whether the improvement came from AI, process redesign, staffing changes, or some combination of all three. The same problem applies to containment, CSAT, conversion, and cost savings. AI rarely arrives in isolation, which makes attribution difficult.

As Opus Research pointed out in The Proof-of-Value Scoreboard: Measuring What Matters for BPO AI (April 2026), credible AI measurement depends on more than headline metrics. Buyers need clear baselines, sensible attribution rules, and enough discipline to separate real improvement from noise. Buyers should therefore look for evidence at several levels.

First is deployment evidence. How many clients are using the capability in production? At what scale? For how long? Second is operational evidence. Can the provider show how the AI is managed, evaluated, updated, and governed after launch? Third is outcome evidence. Are results based on clear baselines and sustained performance, rather than a short pilot or a single favorable metric?

The strongest proof also tends to be specific. A provider that can explain the workflow, the role of AI, the role of humans, the operating model, and the measured result offers a much more credible story than one that presents a percentage improvement without context.

Another useful test is repeatability. A compelling case study can show that a provider solved one client's problem. It doesn't necessarily show that the capability has become part of the BPO's operating model. Buyers have to ask whether the provider can apply the same methods, governance, evaluation practices, and delivery assets across multiple accounts without rebuilding everything from scratch.

This cohort illustrates the problem. Among the five providers that responded to the RFI, reported containment ranged from roughly 35% to 75%, and two reported no containment figure at all. One provider estimated its deployment count by applying a multiplier to its client count. These figures aren't comparable across providers, and buyers shouldn't treat them as if they were.

Where Does BPO AI Differentiation Come From?

AI has created an awkward differentiation problem for CX BPOs. Most providers have access to many of the same foundation models, cloud platforms, CCaaS products, and conversational AI tools. Buying powerful technology is relatively easy. Turning it into something clients can't get just as easily from another provider is harder, and that's where the proverbial rubber meets the road.

That means differentiation is shifting away from the underlying AI itself. A BPO may build some proprietary technology, but few providers are likely to win simply because their model, copilot, or AI agent is dramatically

better than everyone else's. More durable advantages come from what surrounds the technology and how effectively the provider puts it to work.

One such source is operational data and experience. BPOs live inside millions of customer interactions and the workflows surrounding them. They can see where customers struggle, where agents get stuck, which handoffs fail, and which processes repeatedly create avoidable demand. Used well, that experience can help providers identify better automation opportunities and design AI around how service actually operates rather than how a workflow looks mapped out on paper.

Domain expertise can matter just as much. Customer service in healthcare, financial services, travel, collections, or utilities isn't interchangeable. The policies, risks, regulations, terminology, systems, and exceptions are different. A provider that understands those details can build better workflows, determine where autonomy is appropriate, and recognize when an AI agent should stop and hand work to a person. So while the technology may be broadly available, knowing how to apply it safely to a difficult process isn't.

BPOs can also differentiate through reusable intellectual property. That might include proprietary agents or copilots, but it can also be the less glamorous and more useful stuff. Prebuilt integrations, evaluation libraries, workflow templates, governance controls, industry-specific knowledge structures, and benchmarking data can reduce deployment time and improve consistency. The key question is whether the provider has created assets it can reuse across clients rather than starting over with each client.

AI operations are another critical source of advantage. Launching an AI system into production is the starting line, not the finish. Someone has to keep it performing as customers, knowledge, models, and business rules change. BPOs have spent decades running customer service operations continuously. The providers that can extend that discipline to AI have a credible advantage over firms whose involvement largely ends after implementation.

Finally, differentiation increasingly depends on how well the provider manages humans and AI as one "hybrid workforce" (to use the cliché of the day). The maturity model earlier in this report reflects that shift. Leading providers won't simply deploy automation alongside human agents. They'll coordinate work across both, apply shared quality and performance standards, and improve the operation based on what happens across the full customer journey.

Who Owns the AI, and Who Owns the Failure?

As BPOs take on more responsibility for AI-enabled service, an obvious but sometimes uncomfortable question follows. When something goes wrong, who owns it?

The answer gets complicated quickly. A brand may control the policies and knowledge. The BPO may design the workflow and operate the service. A technology provider may supply the AI agent, model, or orchestration layer. If the agent gives a customer the wrong answer, violates policy, or takes an action it shouldn't, responsibility can cross all three. That makes ownership more than some arcane contractual issue. It needs to be designed into the operating model.

Buyers should understand who controls the major parts of the AI system. Who approves prompts and policy changes? Who maintains the knowledge sources? Who sets escalation rules and autonomy limits? Who evaluates behavior before deployment and monitors it afterward? Who can pause or override the AI when something starts going wrong?

Figure 6 offers a starting point for dividing those responsibilities across the five layers of an AI Agent Control Plane. The right split varies by engagement and by client strategy.

Figure 6: Dividing AI Control Plane Responsibilities

Control layer	Client	BPO	Technology vendor
Journey and intent state	Defines journeys and intents in scope and sets autonomy limits	Maintains context across AI and human steps, including at handoff	Carries state and context through the platform
Identity and consent	Sets authentication and consent policy; owns customer data	Enforces verification in workflows and at escalation	Provides identity, consent, and data-handling controls
Policy and guardrails	Approves policy, prompts, and the actions AI may take	Turns policy into workflows and escalation rules	Enforces guardrails at runtime
Knowledge governance	Owns source content and approves changes	Maintains, versions, and tests the knowledge AI and agents use	Supplies retrieval and versioning infrastructure
Evaluation, testing, and audit	Sets acceptance thresholds and reviews results	Runs pre-launch evaluation and ongoing monitoring and QA across AI and human work	Provides logs, evaluation tooling, and notice of model changes

The same questions apply when work moves between AI and people. A failed handoff can be caused by bad routing, incomplete context, a poorly designed escalation rule, or a human agent who receives too little information to recover the interaction. Assigning all that failure to either “the AI” or “the agent” hides how the operation actually works.

More mature BPO relationships therefore need to define accountability across the full service chain. That includes the client, the provider, and any technology partners involved. Clear ownership matters for governance, but it also matters for improvement. Teams can’t fix recurring problems if nobody is clearly responsible for investigating them and making changes.

As BPOs move toward managing hybrid human and AI operations, this becomes part of the value proposition. Providers that want responsibility for the outcome need to show they can also take responsibility for the controls, handoffs, monitoring, and remediation that make those outcomes dependable.

How the Frontline Job Is Changing

Every provider in this report says the frontline job is shifting toward exceptions, judgment calls, and emotionally charged contacts. Yet workforce transformation is the second-weakest dimension in the cohort’s operating model profiles, and evidence of structural change is uneven.

The strongest responses describe structural change, not just training. Foundever, for example, is reorganizing

work by complexity tier, with AI handling Tier 0 and 1 contacts and people handling Tier 2 and 3, and it reports realigning KPIs and compensation to match. iQor is moving supervisors from monitoring lagging KPIs to behavior-based coaching and shifting QA from sampling to calibration and root-cause analysis.

New roles are appearing, most of them around AI operations. Foundever, Firstsource, and TP describe AI trainers, conversation designers, AI quality specialists, and workflow or orchestration specialists. TELUS Digital embeds Forward Deployed Engineers in every CX engagement, and Concentrix describes similar teams in its public materials. AI simulation has become standard in training. TELUS Digital reports that its Agent Trainer is used across about 90% of its CX clients and cuts time to proficiency by 20–50%. iQor reports up to 50% faster proficiency.

Measurement is the gap. When AI removes the easy contacts, handle time and first contact resolution can worsen even when the program is working. Buyers should ask how scorecards, staffing ratios, and pay change as the remaining work gets harder, and who absorbs the cost of the more skilled workforce that remains.

Questions to Ask Your BPO

The questions below gather the tests raised throughout this report. Use them in briefings and RFPs, and expect specific answers rather than capability lists.

- **Deployment.** How many clients use this capability in production, at what scale, and for how long?
- **Build or partner.** Which capabilities do you own, which do you operate on a partner's platform, and which do you resell?
- **Outcomes.** What baseline, control group, and attribution method support each result you cite?
- **Repeatability.** Which methods, controls, and delivery assets carry over to my account without rebuilding?
- **Operations.** Who evaluates, updates, and monitors the AI after launch, and how often?
- **Control.** Who approves prompt and policy changes, sets autonomy limits, and can pause or override the AI?
- **Handoffs.** What context do human agents have access to when an interaction is escalated, and who investigates failed handoffs?
- **Workforce.** How do recruiting, hiring, roles, staffing ratios, scorecards, and pay change as AI absorbs more routine work?
- **Commercials.** How does pricing share productivity gains, and what happens to price as automation rises?
- **Exit.** Who owns prompts, evaluation sets, knowledge, workflows, and interaction data if we part ways?

Methodology

This Intelliview evaluates six CX BPOs: Concentrix, Firstsource, Foundever, iQor, TELUS Digital, and TP. Five responded to an Opus Research request for information (RFI) issued in July 2026 and took part in follow-up briefings. Most also demonstrated products and supplied proof points. Concentrix didn't respond. Given its scale and its position in the market, we assessed it from the public record instead. Crescendo also responded. Because it operates as an AI-native CX platform rather than a BPO, we treat it as a reference point for autonomous customer service, not as a member of the cohort.

Evidence base. Assessments draw on four sources. The RFI response is the primary source. It covers AI capabilities and service delivery, people and workforce models, track record and business impact, and future vision and roadmap, plus a quantitative section on adoption and deployment. Analyst briefings, demonstrations, and follow-up clarifications come next, followed by supplied proof points such as case studies and client results. We use public material to verify or challenge a provider's claim, never to supply one the provider didn't make.

Evidence, not reputation. We credit what a provider documents or demonstrates. Size, revenue, and brand recognition aren't evidence of AI maturity, and a large delivery footprint doesn't by itself earn a higher placement. Partner technology earns credit when the BPO shows it deploys and operates that technology effectively. A thin response can produce a lower placement than a provider's actual capability would merit. Where we believe that is the case, the provider dossier says so.

Public-record assessment. Concentrix is assessed on a different basis from the rest of the cohort. Its placement and scores draw only on public sources available as of October 2, 2026: SEC filings, earnings calls, investor presentations, press releases, and company web content. Concentrix didn't supply, review, or approve any of it. Investor- and marketing-facing material tends to report wins without denominators, so we applied the same criteria and credited only capabilities with specific public evidence of production use. Where the record is silent, as it is on job redesign, measurement method, and containment, the score reflects the absence of public evidence, not a finding that the capability is missing. Concentrix is marked with an asterisk in the figures and in provider lists.

Profiling operating models. The AI operating model profile scores each provider from 1 (limited) to 5 (leading) on six dimensions, against written criteria for each. It answers a different question from the matrix (what kind of AI BPO a provider is), so we read it by shape rather than total and don't rank providers on profile totals.

Building the Intelliview Map. A provider's vertical position averages two equally weighted inputs: its internal operations stage and the mean of its profile scores for frontline augmentation, autonomous customer service, workforce transformation, analytics and optimization, and governance and control. Its horizontal position averages its customer-facing stage and its client-facing AI services score. Each input is rescaled to a common range before averaging. The boundary between Leaders and Standouts reflects analyst judgment about where the evidence separates the groups. The vertical position also reflects analyst judgment from in-person discussions with provider teams, which may show operational progress beyond written responses.

Self-reported metrics. Quantitative responses are self-reported, and definitions vary from provider to provider; one firm's containment rate is another's deflection rate. We don't score these figures directly. We use them to corroborate or challenge qualitative claims, and we report them only in aggregate, as ranges without attribution.

What this assessment is not. This Intelliview evaluates documented evidence gathered between July and early October 2026. It isn't an audit, a market-share analysis, or a financial rating, and matrix placement isn't a ranking. Providers within a cell are listed alphabetically. Apart from Concentrix, assessed from the public record as described above, providers that didn't participate weren't assessed, and their absence implies no judgment about their capabilities.

Appendix A: Company Dossiers

The following dossiers summarize how each of the six BPOs evaluated in the 2026 BPO, AI & CX Intelliview describes its AI strategy and delivery model. Providers are listed alphabetically. Crescendo, an AI-native CX platform that also responded to the RFI, follows as a separate reference profile because it is not assessed as part of the BPO cohort.

Each dossier follows the four pillars of the Opus Research RFI: AI Capabilities & Service Delivery; People & Workforce Models; Track Record & Business Impact; and Future Vision & Roadmap. Content is drawn from each provider's RFI response and supporting materials it submitted. The exception is Concentrix, which did not respond to the RFI; its dossier draws on public sources only, as noted at the start of its profile. Outcomes and figures are as reported by the provider and have not been independently audited.

Each BPO dossier opens with its AI operating model profile: the six dimensions the analysts scored, each from 1 (limited) to 5 (leading). The profile shows what kind of AI BPO a provider is, so read it by shape. Scores are not summed or ranked.

Concentrix

AI Operating Model	Score (1–5)
Frontline augmentation	●●●●○ 4
Autonomous customer service	●●●○○ 3
Workforce transformation	●●○○○ 2
Analytics and optimization	●●●●○ 4
Client-facing AI services	●●●●○ 4
Governance and control	●●●○○ 3

Concentrix did not respond to the RFI or take part in a briefing. Given its scale, it is assessed from the public record (SEC filings, earnings calls, investor materials, press releases, and company website).

Concentrix Corporation (NASDAQ: CNXC) is a CX services and technology provider headquartered in Newark, California. Its origins date to 1983. It was spun off from SYNNEX as an independent public company in 2020 and merged with Webhelp in 2023. It reports about 455,000 employees across roughly 483 locations in 74 countries, and more than 2,000 clients, including 160+ of the Fortune Global 500. FY2025 revenue was $9.8 billion. Technology and consumer electronics is its largest vertical (27%), followed by retail, travel, and e-commerce (25%), media and communications (16%), banking, financial services, and insurance (16%), and healthcare (7%). Its offerings span strategy and design, data and analytics, enterprise technology, and digital operations. Concentrix positions itself as an “intelligent transformation partner” and describes its approach as “blended AI and services.” Its AI portfolio is organized under the proprietary iX Suite.

AI Capabilities & Service Delivery

Concentrix builds and licenses its own AI applications, sold separately from services, and integrates them with partner platforms. The Agentic Operating Framework (AOF), launched in September 2025, packages readiness consulting, brand-aligned language models, agentic engineering, data management, and monitoring into a design-build-run model, supported by Forward Deployed Engineering teams.

Interaction and Journey Automation

- **iX Hello (proprietary):** No-code generative and agentic AI self-service across web chat, SMS, voice, smart devices, and mobile apps, in 90+ languages. An LLM gateway connects to multiple model providers. Concentrix states that iX Hello recognizes “the moment a case needs human judgment” and hands off with full context.
- **Pre-built agents:** Since December 2025, four “emotionally aware” agents for product support, order status, appointment scheduling, and collections and payments. They adapt to tone and to cultural and linguistic nuance.

Employee Augmentation

iX Hero (proprietary), launched in May 2025, serves advisors, supervisors, QA analysts, and trainers:

- **Insights, Knowledge, Summarize, and Announcements:** Performance and sentiment analytics, real-time answers, automated transcripts and summaries, and updates pushed to advisors.

- **Harmony and Clarity:** Speech refinement and noise suppression, optimized for delivery in the Philippines, India, and Latin America. Concentrix reports a 33.6% rise in communication scores and a 22.3% NPS gain.
- **Roleplay:** AI voice and chat simulation for training, in 50+ languages.

Concentrix reports AI deployed on 400,000+ desktops across 1,000+ clients, and more than 30,000 advisors added to the iX Suite in the third quarter of FY2026. It has not defined what counts as an AI desktop.

Revenue Services

Revenue work runs through Concentrix's B2B sales and marketing operations, expanded by its 2022 acquisition of ServiceSource. Collections and payments is one of its pre-built agents. Published results include sales conversion rising from 2% to 7% with iX Hero, a 53% increase in sales conversion from generative AI advisor assistance, and a 56% increase in revenue per chat from AI-driven sales training.

AI Safety, Governance, and Controls

In December 2025, the iX Suite was certified to ISO/IEC 42001:2023 (AI management systems), ISO 31700:2023 (privacy by design), and HITRUST AI requirements. Enterprise certifications include ISO 27001:2022, SOC 2 Type II, PCI DSS 4.0.1, HITRUST CSF, and ISO 22301.

People & Workforce Models

- **Development:** Concentrix University (about 74,000 courses) and the iRise mentoring program. Concentrix reports a 79.6% positive engagement rating.
- **AI in training:** iX Hero Roleplay, which Concentrix reports raised production readiness 47% at an automotive client and brought proficiency two weeks sooner at a financial services client.
- **New roles:** Forward Deployed Engineering teams for agentic deployments, and a data and analytics practice that includes data annotation.
- **Change management:** The AOF names skills and trust as a pillar. FY2026 restructuring, including severance and facility consolidation, totaled $100.4 million in the first half.

Service portfolio: Concentrix attributes about $3 billion of revenue to clients that have gone through heavy transformation or where AI influences the revenue. It attributes about $1.3 billion to traditional services "compressed" onto the iX platform, and about $700 million to new high-value services such as risk and compliance.

Track Record & Business Impact

Concentrix reports 1,000+ clients running AI in production and says about 42% of its new business wins include its technology.

Client / program	Reported outcome
Energy client (iX Hero pilot)	CSAT from 72% to 81.8%
Energy sector client	16.5% FCR improvement; 6% lower AHT
Retail and e-commerce client	175,000 calls transcribed; 85% accuracy on sales insights; 4% higher close rate
Airline (Agentic Operating Framework)	$150M in potential new revenue; $45.8M cost-saving ROI

Measurement: Concentrix frames proof of value through Agentic Value Maps within the AOF. It reports a deployment success rate of "almost 80%" for AI implementations. Public sources do not describe baselines, control groups, or a review cadence.

Commercial models: Core services run on multi-year contracts priced on prevailing labor costs, with performance bonuses and penalties under 1% of revenue. iX Hello is priced by consumption; management says it runs at a negative margin early, then scales to software-like margins. iX Hero is licensed per seat. Concentrix reports about 350 basis points higher margins on clients using the iX Suite.

Future Vision & Roadmap

Concentrix's stated priorities center on scaling the iX Suite as a software business:

- **iX software growth:** ARR from about $60 million at the end of FY2025 to about $120 million by the end of FY2026, with iX Hero expected to reach "SaaS-type margins" at critical mass.
- **Agentic AI:** More pre-built agents and autonomous capabilities, delivered through the AOF and Forward Deployed Engineering.
- **Agentic commerce:** Supporting clients' infrastructure for AI-driven shopping. Concentrix says it has built similar technology for its own use.
- **Platform partnerships:** Deeper integration with agentic platforms such as NiCE Cognigy alongside its own products.

Concentrix builds the application layer itself (iX Hello, iX Hero, the LLM gateway, and AOF tooling) and sources models, cloud, and contact center platforms from partners such as AWS, Microsoft, Google, OpenAI, and NiCE. Management describes value moving from labor toward platform, counting "compressed" services as a revenue category and pricing automation by consumption.

Firstsource

AI Operating Model	Score (1–5)
Frontline augmentation	●●●○○ 3
Autonomous customer service	●●○○○ 2
Workforce transformation	●●○○○ 2
Analytics and optimization	●●●●○ 4
Client-facing AI services	●●●●○ 4
Governance and control	●●●○○ 3

Firstsource Solutions is a business process services provider headquartered in Mumbai, India, founded in 2000. It reports 39,937 employees as of October 2026, about 29,000 of them in CX roles, with delivery in India, the Philippines, the UK, the US, Mexico, Australia, Romania, South Africa, Trinidad, and Turkey. It supports roughly 20 languages and serves telecom and media, banking and financial services, insurance, retail and digital marketplaces, energy and utilities, technology, and healthcare payers and providers. Customer care (41%) and back office (34%) make up most of its CX work, followed by collections, sales, retention, and technical support. Firstsource frames its strategy as "UnBPO," a shift from labor-based outsourcing to an AI-first transformation partner. In 2026 it introduced "Intelligence That Operates" and Kairos, its AI operating system.

AI Capabilities & Service Delivery

Firstsource describes "Intelligence That Operates" as domain expertise encoded into AI that operates work end to end, with accountability, rather than only advising or automating tasks. Kairos, its operating system for AI-native operations, is how it delivers that approach. Firstsource extends Kairos through partnerships with conversational AI, contact center, and workflow automation vendors. These capabilities serve customers, agents, and supervisors, and also run within Firstsource's own operations.

Interaction and Journey Automation

Kairos orchestrates workflows and applies domain intelligence across customer service and back-office operations, so AI can take on complex work within defined governance. Within Kairos, the Intelligent Context Framework (ICF) gives AI agents the memory of the business: account history, how past exceptions were handled, and which decisions were approved. Agents act on the same context human teams rely on rather than on isolated prompts.

Front-office interactions across voice, chat, messaging, and email draw on this foundation for self-service, intent recognition, authentication, guided resolution, knowledge retrieval, summarization, and case management. A Domain Harness sets the rules for when AI resolves an interaction and when it hands off to a live agent, weighing intent, complexity, confidence thresholds, compliance requirements, and customer preference. On handoff, the full conversation context passes to the agent so customers do not have to repeat themselves.

Employee Augmentation

Agent Assist provides real-time guidance, next-best-action, contextual knowledge retrieval, and compliance prompts. Generative AI produces interaction summaries and disposition notes. Supervisors and QA analysts use

interaction analytics for coaching, sentiment monitoring, risk detection, and automated evaluation. Workforce planners use AI-driven forecasting for staffing and scheduling.

Operational Augmentation
Firstsource follows a "skills-first" talent strategy supported by IntelliPrism, a skills-intelligence and workforce-planning platform. Two knowledge platforms, OneFirstsource and Orange Pages, support AI-assisted knowledge discovery and collaborative learning. UnBound, its internal learning platform, delivers personalized, continuous learning; Firstsource reports using it to reskill thousands of employees from transactional roles to judgment-driven positions.

Revenue Services
Predictive models identify customers likely to churn, accept an offer, or respond to collections outreach. Agent Assist supplies next-best-action and offer guidance. AI-led decisioning supports segmentation, prioritization, and channel and timing selection across sales, retention, collections, and cross-sell programs.

AI Safety, Governance, and Controls
Firstsource applies Responsible AI principles with human-in-the-loop oversight for high-impact decisions. It reports ISO 27001, ISO 27701, ISO 22301, ISO 9001, PCI DSS, and SOC 2 (where applicable), along with compliance with GDPR, HIPAA, and CCPA. Solutions support regional data residency according to client and regulatory requirements.

People & Workforce Models
Firstsource describes the frontline role shifting from transaction handling to "AI-enabled customer experience specialist," focused on exceptions, problem resolution, and advisory-led engagement.

- **Hiring:** Recruitment now weighs digital fluency, analytical thinking, learning agility, problem solving, and customer empathy, alongside the ability to interpret AI recommendations.
- **New roles:** Firstsource is developing AI trainers, conversation designers, AI quality specialists, and workflow orchestration experts. Emerging responsibilities include AI governance, knowledge management, and prompt engineering.
- **Training:** AI identifies individual skill gaps, recommends role-based learning paths, and personalizes coaching using QA and interaction analytics.
- **Change management:** Firstsource describes a people-first approach built on continuous learning, AI literacy, and leadership-driven adoption.

Service portfolio: Firstsource has expanded into Consulting & AI Advisory, CX Transformation, Applied AI, Agentic AI, and Decision Intelligence. It delivers these through a Transform–Implement–Operate methodology that spans advisory, implementation, and managed AI operations. The portfolio also includes AI governance, model lifecycle management, and Responsible AI services. Firstsource sees the strongest fit for AI-led delivery in high-volume, repeatable, data-intensive workflows. Complex healthcare cases, financial advice, fraud investigations, and complaint resolution remain human-led with AI support.

Track Record & Business Impact
Common deployments include conversational AI and self-service, intelligent routing, Agent Assist, summarization, automated quality management, demand forecasting, back-office automation, and revenue

programs using next-best-action and predictive insights. Firstsource reports the following outcome ranges across programs:

- 12–18% AHT reduction
- 10–12% FCR improvement
- 15–20% higher customer satisfaction scores
- 15–20% higher conversion rates
- 10–15% improved retention rates

Firstsource offered to share case studies during the analyst briefing rather than in the written response. Several recent case studies, including one for KFC, are published on its website.

Measurement: Engagements begin with a baseline across AHT, after-call work, FCR, CSAT, NPS, productivity, quality, compliance, and cost to serve. Firstsource uses phased pilots, control groups where feasible, and normalization for case mix, complexity, seasonality, and volume. Governance runs through weekly or bi-weekly operational reviews, monthly value-realization reviews, and quarterly executive reviews.

Commercial models: Input-based (FTE, hour, minute), output-based (incident, fixed fee), and outcome-based pricing tied to CSAT, NPS, total cost of operations, FCR, or revenue. Proprietary AI platforms are licensed on a platform or subscription basis. Transformation engagements increasingly use hybrid models with performance incentives or gainsharing.

Future Vision & Roadmap

Over the next 12–24 months, Firstsource plans to invest in agentic AI, AI-native workflow orchestration, Applied AI, predictive analytics, and its proprietary platform Kairos.

- **Kairos** is being built out as an AI-native operating system that combines domain intelligence, orchestration, and governed AI agents across customer and back-office operations.
- **Decision intelligence** investments target predictive insights, next-best actions, and proactive customer engagement.

Firstsource will build proprietary platforms and industry-specific accelerators in-house. It will source foundation models, cloud infrastructure, CCaaS platforms, and enterprise applications from partners, and expects closer collaboration with hyperscalers. It expects the contact center to shift toward AI handling routine interactions while people focus on complex problem solving, AI supervision, and exception handling. In its view, the BPO model moves from headcount-based delivery to outcome-based, platform-enabled managed services. Risks it identifies include inaccurate or biased outputs, data privacy, regulatory compliance, workforce disruption, and over-reliance on automation for sensitive interactions.

Foundever

AI Operating Model	Score (1–5)
Frontline augmentation	●●●●○ 4
Autonomous customer service	●●●○○ 3
Workforce transformation	●●●●○ 4
Analytics and optimization	●●●○○ 3
Client-facing AI services	●●●●○ 4
Governance and control	●●●○○ 3

Foundever is a privately owned CX services provider headquartered in Luxembourg. It was founded in 1996 and rebranded as Foundever in 2023. The company reports about 130,000 employees across 150+ delivery centers in more than 45 countries, supporting 60+ languages and 800+ brands. It reports handling about 3.3 billion customer conversations a year. Verticals include financial services, retail and e-commerce, technology, telecommunications, travel and hospitality, automotive, healthcare, insurance, government, and media. Foundever positions itself as a "managed AI operations partner and transformation partner" built on a dual-workforce model: AI-augmented human agents and agentic AI "digital workers" that it designs, deploys, and manages within client technology stacks.

AI Capabilities & Service Delivery

Foundever describes itself as tool- and platform-agnostic. Its proprietary layers are its orchestration capability, the PeakOS operating framework, the Ever-branded employee tools, and pod-based delivery teams. Foundation models, NLU engines, and channel platforms are selected to fit each client.

Interaction and Journey Automation

Foundever's Agentic AI Workforce covers voice, chat, messaging, email, and social channels. Use cases include chatbots and voicebots for account lookups, status checks, scheduling, and payments; natural-language IVR modernization; and end-to-end journeys in KYC onboarding, collections, insurance first notice of loss (FNOL), technical triage, and sales qualification. AI agents assess complexity, sentiment, and confidence in real time and escalate with full context when thresholds are crossed.

Agentic Pods, teams of conversation designers, AI engineers, and CX analysts, manage each agent through design, deployment, and ongoing tuning. Reported results include 57% containment with an 18% CSAT improvement for a US retailer's chatbot and 74%+ containment for a financial services firm's voice agent.

Employee Augmentation

EverAssist: Real-time copilot for knowledge retrieval, next-best-action, compliance prompts, and response suggestions.

- **EverCoach:** AI-generated coaching recommendations based on individual performance data.
- **EverTrainer:** AI simulation and role-play with personalized learning paths.
- **GenAI QA:** Evaluation of up to 100% of interactions, versus 2–5% manual sampling.
- Automated wrap-up summaries and real-time sentiment and intent analysis.

Operational Augmentation

Smart Hiring (proprietary) uses predictive analytics to match candidates to role requirements. PeakOS embeds AI in forecasting, scheduling, and real-time adherence. Predictive models flag attrition risk from engagement, performance, and behavioral signals.

Revenue Services

Foundever runs Sales as a Service (SDR/BDR prospecting, telesales, upsell, renewals, win-back) and Collections as a Service (forbearance, skip tracing, omnichannel recovery). AI supports propensity scoring, next-best-offer, AI-driven segmentation, contact-timing optimization, and churn prediction for save desks. It reports a 20% reduction in non-value transactions with a CSAT increase for a European car finance client.

AI Safety, Governance, and Controls

Controls include retrieval-augmented generation (RAG) grounded in approved content, confidence thresholds that trigger human escalation, and bias testing and auditing. Decision logic for routing, next-best-action, and QA scoring is documented and auditable. Data residency is managed to meet GDPR and regional sovereignty rules. All deployments keep human-in-the-loop governance, and supervisors retain override authority.

People & Workforce Models

Foundever is restructuring frontline work by complexity tier: AI handles Tier 0/1 interactions, while human agents focus on Tier 2/3 work requiring judgment, negotiation, compliance expertise, or emotional sensitivity. It reports realigning role expectations, KPIs, and compensation to match.

- **Operator Pods:** Cross-functional teams of operations leads, data analysts, process engineers, and change managers embedded in client operations. They identify automation opportunities, redesign workflows, and upskill agents. Foundever reports these programs deliver a 15–25% reduction in cost per contact, 8–12 point CSAT improvement, and a 10–15% FCR increase.
- **PeakOS lifecycle:** Governs the associate lifecycle of Hire, Train, Plan, Engage, Grow, and Retain, with AI at each stage.
- **New roles:** AI trainers, conversation designers, prompt engineers, AI quality auditors, and orchestration specialists within Agentic Pods.
- **Change management:** Emphasizes role evolution over elimination, structured reskilling, and engagement programs. Foundever cites 700,000 work hours saved through process automation at a UK bank, with associated staff redeployment.

Service portfolio: AI is extending Foundever beyond customer care into Financial Crimes Compliance and KYC, claims and insurance, Collections as a Service, trust and safety, Tier 2/3 technical support, Sales as a Service, and CX consulting through Operator Pods. Foundever engages clients along a three-stage maturity path: building CX foundations, transforming CX operations, and optimizing the agentic workforce.

Track Record & Business Impact

Common deployments span customer-facing automation, agent and supervisor augmentation, internal operations (Smart Hiring, EverTrainer, attrition modeling), and revenue programs. Reported outcomes include:

Program	Reported outcome
Operator Pod programs	15–25% lower cost per contact; 8–12 pt CSAT gain; 10–15% FCR uplift
Deployed programs	25–30% contact deflection
US retailer	$10M annual cost-to-serve reduction; 57% chatbot containment with +18% CSAT
Global financial services	74%+ AI voice containment
UK bank	500K+ applications processed; 700K work hours saved
Digital travel brand	80%+ multilingual agent retention; NPS +60
Global retail bank	93% risk management index vs. 80% target

Measurement: Baselines are set during an Operator Pod "Diagnose & Plan" phase. Foundever uses control groups where feasible, normalized case-mix analysis, and attribution modeling that separates AI, agent augmentation, and process changes. Reviews run weekly or bi-weekly (operational), monthly (executive), quarterly (strategic), and annually. Its KPI framework includes AI-specific measures such as hallucination rate and human override frequency.

Commercial models: FTE-based, per-transaction, outcome-based, gainshare, platform or managed-service fees for Agentic Pod support, and hybrid models. Foundever also offers Bots-as-a-Service, a consumption-based model in which clients pay for AI outcomes such as contained interactions or completed processes.

Future Vision & Roadmap

Foundever's 12–24 month priorities are:

- **Agentic AI expansion:** Scaling from 25+ live AI agents to multi-step, multi-system workflows such as end-to-end claims adjudication, collections negotiation, and diagnostic troubleshooting.
- **Voice AI:** Improved natural language understanding, real-time translation, and domain-specific fluency.
- **Multimodal AI:** Visual, document, and screen-sharing capabilities, such as document verification for KYC.
- **Orchestration:** Dynamic routing among AI agents, AI-augmented humans, and specialists based on complexity, sentiment, value, and compliance.
- **Analytics and workforce intelligence:** Predictive and prescriptive insights, and models for balancing the human workforce against growing AI capacity.

Foundever will build orchestration, workforce tools, and operating frameworks in-house and partner for LLMs, NLU, cloud, and channel platforms. It expects the contact center to move from "human-primary, AI-assisted" to "AI-primary, human-elevated" within two to three years. In that model, it says providers must manage human and AI workforces as one operation with shared governance and outcome accountability. Risks it highlights include hallucination, regulatory exposure, workforce displacement anxiety, vendor lock-in, and commoditization of basic automation.

iQor

AI Operating Model	Score (1–5)
Frontline augmentation	●●●●○ 4
Autonomous customer service	●●○○○ 2
Workforce transformation	●●●●○ 4
Analytics and optimization	●●●●● 5
Client-facing AI services	●●●●○ 4
Governance and control	●●●○○ 3

iQor is a CX BPO headquartered in Fort Lauderdale, Florida. It was founded in 1998 as IntelliRisk Management Corporation and renamed iQor in 2007. The company reports 45,000 employees, including about 40,000 contact center staff, across 40+ sites in 10 countries. About 83% of headcount is offshore, 9% nearshore, and 7% onshore. It supports 20+ languages. Telecom is its largest vertical (39%), followed by banking and financial services, utilities, legal and government, retail, healthcare and insurance, travel, and high tech. Service lines are omnichannel care (65%), revenue recovery (18%), sales and retention (10%), and technical support (7%). iQor is a "BPO-first" managed AI operations partner, explicitly not a software company. It reports that 64% of its revenue is AI-enabled.

AI Capabilities & Service Delivery

iQor's AI-powered CX platform is branded infinityAiQ. Its foundation is proprietary data and analytics, combined with partner technology for conversational AI, agent assist, CCaaS, and language services. An AI Center of Excellence (AI CoE) provides governance and moves proven use cases from pilot to production.

Interaction and Journey Automation

- **Valdi (proprietary):** iQor's internal operating-intelligence and analytics platform for operations leaders, supervisors, and QA teams.
- **Insights iQ (proprietary):** A client-facing interaction-intelligence layer built on Valdi and integrated with OpenAI and Anthropic models. It surfaces intent, sentiment, friction, repeat-contact drivers, compliance risk, revenue opportunities, and automation candidates.
- **Analyst GPT (proprietary):** Natural-language exploration of Insights iQ intelligence to uncover recurring interaction patterns, investigate root causes, and quantify opportunities for improvement and automation.
- **Automation Discovery:** Uses that intelligence to decide where automation should be applied.

Automation use cases include self-service optimization, predictive intent and resolution, intelligent routing, payments and account servicing, order and service status, scheduling, collections, claims and voucher processing, and authentication support. Handoff logic weighs intent, confidence, sentiment, journey complexity, authentication status, compliance, and revenue or retention risk, and passes full context to the agent.

Employee Augmentation

- **Agents:** Agent Assist, Real-Time Agent Assist, Knowledge Assist, Predictive Intent and Resolution, next-best-action, and Auto Call Summarization.

- **Supervisors:** Coach iQ, which iQor reports enables about four times more coaching sessions per week, plus Predictive NPS and Analyst GPT.
- **QA and compliance:** Auto QA evaluates interactions against client-specific criteria across 60+ parameters. Predictive NPS forecasts satisfaction with more than 80% accuracy, per iQor.
- **Trainers and planners:** AI simulated training, Learning as a Service, WFM as a Service, STAMP workflow methodology, and Attrition Predictor with Check-Ins.

Operational Augmentation

In 2025, 63% of iQor's new hires came through partner-enabled Recruiting AI, with candidates able to reach ready-to-hire status in under 24 hours. iQor reports up to 500% year-over-year seasonal staffing flexibility and 100% peak staffing compliance for a major retailer.

Revenue Services

Revenue work is an established business line: revenue recovery is about 18% of the business, and sales and retention about 10%. iQor Grow (Growth as a Service) covers digital marketing, lead generation, inbound and outbound sales, and account management. QNEX, its proprietary recovery platform, supports collections with account prioritization, payment-intent detection, and compliance monitoring. Insights iQ and Predictive NPS identify churn risk and buying signals.

AI Safety, Governance, and Controls

The AI CoE evaluates each use case before production on business purpose, data sensitivity, risk, regulatory obligations, and required human oversight. iQor reports PCI DSS, SOC 1 Type 2, SOC 2 Type 2, HITRUST, ISO 27001, HIPAA, and GDPR compliance. The degree of AI autonomy depends on the use case, confidence level, and risk. People remain accountable for QA calibration, employee performance decisions, and compliance-sensitive actions. iQor describes its approach as "bounded, accountable AI."

People & Workforce Models

iQor reports that about 70% of its work is already complex, specialized, or regulated, including fraud prevention, content moderation, loan servicing, legal services, credit assessments, and revenue operations. It expects AI to raise the value of skilled human roles rather than simply remove transactional work.

- **Role redesign:** Frontline roles are becoming more skill-based and specialized. Performance measures are shifting toward resolution, quality, compliance, sentiment, revenue impact, and effective use of AI tools.
- **Team structures:** Supervisors move from monitoring lagging KPIs to behavior-based coaching. QA shifts from sampling to calibration and root-cause analysis. Workforce teams align staffing to work complexity and value.
- **New capabilities:** AI training and validation, conversation and workflow design, knowledge optimization, automation discovery, prompt and use-case design, and human-plus-AI orchestration, developed across iQor Labs, operations, training, QA, and transformation teams.
- **Training:** AI simulation and scenario-based learning, which iQor reports have improved speed to proficiency by up to 50%.
- **Retention and development:** Programs include iQor University, iLead, sQholars, LevelUp, and Coach-the-Coach. iQor reports 89% employee satisfaction and 94% of management promoted from within.

Service portfolio: AI is moving iQor toward "insights-led, AI-enabled managed operations," including CX and gaming analytics, journey diagnostics, automation design, QA transformation, workforce optimization, and revenue operations.

Track Record & Business Impact

iQor groups its deployments into customer-facing automation, agent and supervisor augmentation, internal operational automation, and revenue optimization. Reported results include:

Area	Client / program	Reported outcome
Automation	Airline voucher processing	29% lower handle time; ~1,965 hours saved annually
Automation	Order-status use case	~83% lower cost per interaction
Augmentation	Predictive NPS deployments	38% FCR improvement; 87% NPS gain among lower-quartile agents
Training	Home services	50% faster proficiency; 32% higher sales conversion
Sales	Social media company	$90M new SMB revenue; 275% of quota
Retention	Home warranty	30x increase in saves within four weeks
Collections	Credit card issuer	50% more card payments processed via voice-to-chat and SMS
Compliance	Regulated energy provider	About $3M in regulatory penalty exposure avoided

Measurement: iQor sets baselines from historical performance and interaction intelligence before deployment. It normalizes for volume, seasonality, channel, segment, tenure, and complexity, and uses phased rollouts, control groups, or A/B tests where practical.

Commercial models: FTE, hourly, seat, and capacity models for complex human-led work. Transaction- and volume-based pricing for defined workflows. Platform, usage, or managed-service fees for analytics and automation. Gainsharing or outcome-based components where attribution is clear. iQor reports about 35% of clients already operate under outcome-based pricing.

Future Vision & Roadmap

iQor's 12–24 month roadmap aims to connect individual AI use cases into one operating model.

- **Knowledge as a Service:** Dynamic, context-aware knowledge for agents and AI systems.
- **Agentic workflow orchestration:** AI that completes defined workflow steps and escalates to humans, within governed boundaries.
- **Predictive CX:** Extending Predictive NPS and interaction intelligence to anticipate dissatisfaction, churn, and revenue opportunities earlier, enabling proactive outreach.
- **Workforce enablement:** Coach iQ, AI micro-training, and simulation that deliver continuous, personalized performance support.

iQor will build where its operating data and workflow knowledge differentiate, including Valdi, Predictive NPS, Coach iQ, and orchestration. It will partner for conversational and agentic AI, CCaaS, LLMs, real-time assist, and language technologies. Named partners include NiCE Cognigy, OpenAI, Anthropic, AWS, Sanas, Balto, Second Nature, Zenerate, Capacity.com, and Replit. iQor expects BPO value to tie less to FTE counts and more to orchestration, customer intelligence, and outcomes, with hybrid commercial models bridging the transition. Risks it highlights include over-automation, model drift and bias, workforce disruption, and an industry tendency to equate AI adoption with labor removal.

TELUS Digital

AI Operating Model	Score (1–5)
Frontline augmentation	●●●○○ 3
Autonomous customer service	●●○○○ 2
Workforce transformation	●●●○○ 3
Analytics and optimization	●●●○○ 3
Client-facing AI services	●●●●● 5
Governance and control	●●●○○ 3

TELUS Digital is a digital CX and IT services company headquartered in Vancouver, Canada, founded in 2005. It was taken private in 2025 and operates within the TELUS family of companies. The company reports 80,000+ employees, including about 65,000 frontline CX staff and 4,000+ digital and AI practitioners, across 70+ delivery centers and digital studios in 35+ countries. It supports 50+ CX languages and 500+ languages and dialects for data annotation. Service lines include CX management, trust and safety, financial crime and compliance, AI data solutions, AI transformation, and digital solutions such as web, mobile, and enterprise technology modernization. TELUS Digital positions itself as a transformation partner and states that every CX engagement is "AI-first by design."

AI Capabilities & Service Delivery

Every new CX engagement includes three elements by default:

1. A 4–6 week **CX Strategy Assessment** covering digital journeys, contact center operations, and the technology stack, producing an AI transformation roadmap.
2. **Core tooling:** Fuel iX Agent Trainer and two-way noise filtering.
3. **Forward Deployed Architects/Engineers,** embedded on site with delivery teams, who own the client's AI roadmap.

Interaction and Journey Automation

Automation is largely partner-delivered. CCaaS and CRM automation draws on Five9, Google Cloud, Salesforce, Verint, Zendesk, Amazon Connect, Genesys, Automation Anywhere, and eGain. Virtual agents for service and sales are deployed with Gemini Enterprise for CX, ElevenLabs, Cresta, and Parloa. TELUS Digital also builds mobile and web applications intended to remove the need for contact. It developed TELUS Expert Messaging, a bot-triaged asynchronous messaging channel, for TELUS Communications and reports 87% first contact resolution, 40% lower agent attrition, and a 17% CSAT increase.

Employee Augmentation and Operations

The proprietary Fuel iX suite includes:

- **Agent Trainer:** AI voice and chat simulations with configurable scenarios, automated feedback, and machine translation for multilingual programs. Planned enhancements include workflow and agent desktop simulation.

- **Agent Quality** (beta; launch planned for Q1 2027): AI review of 100% of calls, chats, and emails against quality rubrics.
- **Agent Interviewer** (beta; launch planned for Q4 2026): Structured 15–25 minute AI voice interviews scored against job rubrics. Hiring managers retain decision authority.

Partner tools include Sanas for accent softening and noise cancellation and Verint for workforce management and employee self-service scheduling.

Revenue Services

TELUS Digital uses a proprietary method to codify sales "gold standards." Human annotators and LLM analysis of transcripts isolate what top sellers do differently, and the findings feed agent guidance. In a proof of concept now in production at TELUS Communications, it reports a 22% increase in sales conversion and a 17% increase in retention. It also builds bespoke upsell engines for clients and deploys virtual sales agents with Salesforce, ElevenLabs, Cresta, Parloa, and Gemini Enterprise for CX.

AI Safety, Governance, and Controls

Fuel iX Fortify is a proprietary automated red-teaming application. It simulates adversarial attacks before and after deployment, classifies vulnerabilities against OWASP, MITRE ATLAS, and NIST AI-RMF, and generates audit-ready evidence for the EU AI Act, GDPR, HIPAA, and ISO 42001. TELUS Digital also offers RLHF, adversarial red-teaming, and LLM output-monitoring services. For data residency, it offers TELUS's Sovereign AI Factory (SAIF), which keeps training, processing, and storage within Canadian borders.

People & Workforce Models

As AI absorbs routine queries, TELUS Digital is reorienting agents toward complex, judgment-based interactions. It does this through hiring for relevant skills, AI-simulated training, agent tooling, and automated quality management that targets coaching.

- **Forward Deployed Engineers:** Deployed on every CX engagement and recruited in-region. They serve as the catalyst for upskilling and reskilling.
- **Specialized academies:** A Sales Excellence Academy for revenue programs, and a path toward ACAMS certification for financial crime and compliance leadership.
- **AI trainers:** Experts Engine, a proprietary sourcing platform, matches subject-matter experts to AI data-creation, annotation, and validation tasks.
- **AI in training:** TELUS Digital reports that Fuel iX Agent Trainer is used across about 90% of its CX client base and cuts time to proficiency by 20–50%. ELSA Speak, a voice-enabled generative AI tutor, helps non-native English speakers build language proficiency.
- **Change management:** Framed as "retraining, not reduction," supported by certification, simulation, wellness, and upskilling programs. TELUS Digital reports a 78% team member engagement score.

Service portfolio: About 60% of CX accounts now include services beyond standard operations, such as Fuel iX deployments, CX Strategy Assessments, agentic automation projects, trust and safety, and AI data operations. TELUS Digital sees high-volume, low-ambiguity, digitally native workflows as best suited to AI-led delivery. It notes that AI performance is often limited more by knowledge quality than by model capability.

Track Record & Business Impact

The most common AI use cases across TELUS Digital clients are AI-simulated training, assistive copilots, AI-driven talent acquisition, AI-led workforce management, task automation, automated quality management, accent softening and translation, and AI-driven coaching. TELUS Digital measures outcomes in five categories: efficiency, workforce, revenue, quality and risk, and customer outcomes. It notes that specific client figures are governed by confidentiality terms. Public case studies cover conversational AI in financial services, ElevenLabs-based customer onboarding, and Agent Trainer proficiency gains.

Measurement: TELUS Digital proves value through controlled pilots. Baselines draw on up to two years of account history, with outliers removed, on segments untouched by other AI initiatives. Parallel control groups run alongside, and results are tracked weekly for at least four weeks after launch before moving to dashboard monitoring. Results feed existing monthly and quarterly business reviews. TELUS Digital flags that mature KPIs such as AHT and FCR can worsen even when AI is working, because automating tier-one contacts shifts the remaining mix toward harder cases.

Commercial models: Input-based (worked or productive hour, FTE, worked minute), output-based (transaction, contact, handled minute), fixed-rate, managed services, and outcome-based pricing. AI setup, platform, and licensing are modeled separately so clients can see bundled and unbundled costs. Core CX AI tooling is included in every new engagement, with part of the deployment cost absorbed in program margin. Outcome-based pricing is being extended as a standard option where results are measurable, attributable, and within TELUS Digital's control.

Future Vision & Roadmap

TELUS Digital's 12–24 month roadmap includes:

- **Conversational Intelligence (proprietary):** Real-time, natural-language analysis of interaction data, including calls and transcripts, to identify contact drivers, emerging issues, and behavioral trends. In production at TELUS, with a product launch planned for 2027.
- **Fuel iX Agent Quality (proprietary):** Full-coverage automated QA with coaching dashboards and compliance-risk flagging.
- **Virtual agent implementation (partner-driven):** With partners including Cresta, ElevenLabs, Salesforce, Gemini Enterprise for CX, and Parloa.

It expects four shifts in the contact center. Routine work continues moving to automation. Commercial models move from headcount to outcomes. The human role concentrates on judgment, escalation, and automation failures. Governance grows in importance. Risks it identifies include automating broken processes on fragmented knowledge sources, the proliferation of AI tools in client stacks, and bias in automated decisions. Its stated mitigation is to assess the data and knowledge foundation, through the CX Strategy Assessment, before automation goes live.

TP

AI Operating Model	Score (1–5)
Frontline augmentation	●●●●○ 4
Autonomous customer service	●●●○○ 3
Workforce transformation	●●●○○ 3
Analytics and optimization	●●●●○ 4
Client-facing AI services	●●●●● 5
Governance and control	●●●●● 5

TP (formerly Teleperformance) is a digital business services provider headquartered in Paris, France, founded in 1978. It reports about 490,000 employees across 600+ facilities in nearly 100 countries, with support for 400+ languages and dialects. Care makes up about 63% of its business, followed by trust and safety (9%), technical support (8%), sales (8%), back office (5%), consulting, analytics, and technology products (5%), and data services for AI (2%). It serves a broad set of verticals, including financial services and insurance, technology, telecom, healthcare, retail, travel, media and gaming, public sector, energy, and automotive. TP positions itself as a transformation partner. Its AI program is organized around the TP.ai framework, which it describes as orchestrating people, AI agents, data, and workflows under one operating model.

AI Capabilities & Service Delivery

TP combines its own orchestration and governance layers with partner technologies through an LLM-agnostic, API-first architecture.

Interaction and Journey Automation

TP-developed capabilities:

- **TP.ai:** AI orchestration, workflow management, journey design, observability, and AI lifecycle management.
- **Agentic CX Orchestration:** Workflow redesign, AI agent coordination, and human-in-the-loop operations.
- **Trust & Control Layer:** Governance, policy enforcement, auditability, approval controls, and escalation management.

Partner-enabled capabilities include voice bots and agentic workflows (in-house and with Parloa), chatbots (in-house and with Kore.ai and AWS), agentic workflows with Ema, and IVR through CCaaS partners. Client-facing use cases include account servicing, claims, order management, scheduling, and onboarding. Routing and escalation are governed by a decisioning engine that weighs intent, confidence, risk, compliance, and complexity. Routine requests can run fully autonomously; regulated or emotionally sensitive interactions go to human advisors; and hybrid workflows let AI agents and specialists collaborate within a single interaction.

Employee Augmentation

TP.ai Assist is embedded in agent workflows. It provides next-best-action, intent and topic recognition, suggested responses, summarization, wrap-up notes, and RAG-based knowledge retrieval. TP also uses AI-generated simulations for onboarding and training. AI augmentation extends to mid- and back-office functions such as finance and accounting.

Operational Augmentation

AI-driven analytics, simulation models, and demand-supply optimization support workforce planning and resource allocation. The TP.ai platform applies security-by-design, auditability, explainability, and industry-specific compliance controls.

Revenue Services

- **TP.ai Growth:** A "Revenue as a Service" offering covering lead sourcing and enrichment, outreach automation, proposal generation, and sales analytics.
- **TP.ai Collect:** Predictive analytics for collections, including account segmentation and strategy-engine routing.

AI Safety, Governance, and Controls

TP.ai Connect provides governance, auditability, explainability, and risk management, with human experts kept in complex or sensitive cases. TP holds ISO/IEC 42001:2023 certification for its AI governance framework, integrated with ISO 27001 and ISO 27701. It also reports PCI and HIPAA-aligned data handling.

People & Workforce Models

TP describes its operating model as three interdependent layers:

- **Empathy Layer:** AI delivers speed and consistency, while people own moments that require trust, judgment, and emotional intelligence, with handoffs designed to be invisible to the customer.
- **Trust in Both Directions:** Agents must trust AI recommendations without depending on them, and customers must trust that AI-enabled service stays accountable.
- **Feedback Loop:** Human escalations, interventions, and exception handling generate data that improves models and workflows.

Frontline agents are becoming "AI-augmented specialists," with teams organized around intelligent task orchestration. New roles include AI trainers, conversation designers, workflow orchestrators, model optimization specialists, and human oversight experts. Employees are reskilled through real-world AI projects, simulation-based development, and feedback-driven learning.

Service portfolio: TP is expanding into AI consulting and strategic advisory, AI-powered managed services, agentic workflow orchestration, data services for AI (model training, annotation, data refinement), and industry-specific offerings. These include financial crime operations, claims, underwriting, medical coding, utilization management, and revenue cycle management. For clients still forming an AI strategy, TP provides operational assessments, AI simulation models, and transformation roadmaps.

Track Record & Business Impact

Reported client outcomes include:

Client	AI program	Reported outcome
Global airline	Demand-supply optimization, AI simulation, roadmap	$10M annual cost savings; 50% of centers consolidated
US retailer	TP.ai Collect	11.2% better debt liquidation; strong increase in debt recovery over nine months
Household appliance manufacturer	TP.ai Assist	24% more volume at flat cost; 16% higher productivity; 10% lower AHT
US healthcare insurer	TP.ai Assist with GenAI knowledge base	8% lower AHT; 7% higher new-hire productivity over two months
Major US bank	Copilots, simulations, incentive campaigns	20%+ efficiency gain; 15% shorter interactions; +6 NPS
UK digital bank	AI-enabled financial crime operation	2x fraud detection

Across anonymized clients, TP reports TP.ai Assist AHT reductions of 20–30% (public sector, retail, automotive). It also reports onboarding acceleration of 13–19%, with more than 20% lower attrition in one financial services program.

Measurement: TP evaluates proof of value on business outcomes rather than AI adoption metrics. Its KPI categories span customer experience, operational performance, financial impact, employee proficiency, compliance and risk, and client-specific outcomes. It uses analytics, simulation models, and workflow analysis to set baselines and track impact.

Commercial models: TP supports FTE/seat, interaction, consumption, platform-fee, managed-service, outcome-based, gainsharing, performance-based, and hybrid models. For AI-enabled programs it increasingly uses transaction-based and outcome-oriented structures that scale with a client's AI maturity.

Future Vision & Roadmap

TP's investment is centered on agentic AI, orchestration, and industry-specific solutions delivered through TP.ai. Current priorities include:

- Intelligent task routing, hybrid human-AI workflows, and an agent control plane.
- Real-time learning and continuous feedback loops.
- Industry AI blueprints for banking and financial services, insurance, healthcare, and retail.
- Interaction analytics, AI simulation models, model training, and data refinement.

The evolving Trust & Control Layer adds agent inventory management, ownership controls, approval gates, audit trails, cost visibility, human override, and cross-vendor observability. TP will build orchestration logic, industry blueprints, and analytics in-house, and partner with hyperscalers and vertical specialists through its LLM-agnostic architecture, with acquisitions also an option. It expects the BPO model to shift from labor-based delivery toward outcome-based services measured by productivity, experience, and revenue impact. Risks it identifies include failing to reach measurable impact, generic AI that lacks enterprise context, organizational silos, workforce resistance to untrusted tools, legacy integration, and data fragmentation.

Reference Profile: **Crescendo**

Crescendo responded to the RFI but operates as an AI-native CX platform rather than a BPO. It is profiled here as a reference point for autonomous customer service and is not placed on the BPO AI Maturity Model.

Crescendo (CrescendoAI Inc.) is headquartered in San Francisco and was founded in 2024. It describes itself as "the AI-native customer experience platform." Its product, the Crescendo Customer Experience Platform (CXP), launched in August 2026 and is built around an operating model in which AI agents do the customer and operator work and people govern it. Crescendo reports hundreds of production AI deployments worldwide across retail and e-commerce, consumer social and subscription apps, connected devices, health and wellness, and technology and SaaS. Customers keep their own human teams or work with partners such as Alorica.

AI Capabilities & Service Delivery

Crescendo's view is that launching an AI agent is a project, but running one well is a permanent job. CXP assigns that job to agents. CX agents serve customers and support human teams. An Agentic Foundation of Optimization, Simulation, Knowledge, and Integration agents deploys, tests, grounds, and improves them, and the customer's team approves every change. Crescendo delivers CXP as managed AI software, with a deployed team of AI Deployment Engineers (AIDEs) and AI Product Managers.

Concierge AI

Concierge AI is the customer-facing layer: a set of AI agents for each business, each with its own skills, knowledge, and authority, working across voice, chat, email, SMS, and WhatsApp. Typical agents include:

- **Retail and e-commerce:** Refunds, returns, cancellations, order status, and a shopping assistant for product discovery, sizing and fit, and purchase guidance.
- **Consumer social and subscription apps:** Plan advisory, account and billing support, subscription changes, and support for marketing campaigns and events.
- **Connected devices:** Setup, troubleshooting, and diagnostics.
- **Technology and SaaS:** Product and account support at scale.
- **Revenue across verticals:** The agents also guide purchases, run retention and save motions, and make upsell and cross-sell offers using authorized customer context. Influence, launched in March 2026, adds personalization and memory across chat, voice, email, and messaging. Crescendo reports an average chat-to-order conversion of 24%, with peaks of 58%, and cites an industry benchmark of 5–15%. Discounts, credits, and other consequential commitments stay within explicit authority and escalation rules.

Escalation is configured around authority, risk, confidence, customer request, emotion, and policy, and the human team receives the full conversation context.

Agentic Foundation and the Improvement Loop

The Agentic Foundation runs the work that surrounds a production agent:

- **Optimization** configures, deploys, and tunes agents after go-live, updating prompts, instructions, and guardrails.
- **Simulation** runs regression testing on every change before release. Evals test instructions, SOPs, skills, tools, and guardrails against simulated conversations, comparing old and new behavior.

- **Knowledge** keeps reference content current, finding and filling gaps from live conversations.
- **Integration** connects agents to production systems, including custom and homegrown ones. Crescendo reports integrations in under 30 minutes via MCP, and backend integrations cut from about 30 days to as few as 3.

According to Crescendo, Quality scores 100% of AI- and human-handled conversations, and Applied Insights turns those scores into themes and root causes. Together they run a five-step loop:

1. Quality scores every conversation.
2. Applied Insights spots a pattern, such as a rising handoff rate, and identifies the root cause.
3. Optimization proposes a fix, working with the Knowledge and Integration agents.
4. Simulation runs a regression test, rerunning the full eval set against the change.
5. The customer's team reviews and approves. Every change is versioned and logged.

Crescendo reports that across its customers, the customer dissatisfaction rate fell from 5.8% to 0.68% in under a year.

Customer Profiles, Memory, and Skills

A shared platform layer carries customer profiles (identity, history, and journey), agentic memory, and reusable skills. A skill built once works across agents, channels, and Agent Assist, and context carries from one channel to the next so customers don't repeat themselves. The same data reaches beyond support: Crescendo cites RealVNC using Applied Insights to send account-friction patterns to its product team, and NIMA using support conversations to drive product and shipping improvements.

Agent Experience

For human agents, Agent Assist provides conversation context, customer history, knowledge, drafted replies, and next best actions. Quality turns conversation scoring into targeted coaching.

AI Deployment Engineers and AI Product Managers

Every deployment includes AIDEs and AI Product Managers. They turn a customer's top contact drivers into custom AI use cases, build custom insights and analytics for continuous improvement, integrate agents into production systems, and keep the platform calibrated to the business. Crescendo reports that deployments go live in about four weeks.

AI Safety, Governance, and Controls

Governance rests on explicit decision rights, role-based access, least-privilege skills, answers grounded in approved knowledge, separation of test and production, pre-launch simulation, versioning, rollback, and a full audit trail. Crescendo states that no change reaches customers without human approval; its principle is that "self-improving doesn't mean self-authorizing." Crescendo reports SOC 2 Type II compliance and describes itself as HIPAA-capable. Its data processing agreement addresses GDPR and UK GDPR, and data location is validated per client and use case.

People & Workforce Models

Crescendo does not employ frontline agents. Human teams handle the interactions that need judgment and approve every change to the AI.

Track Record & Business Impact

Crescendo reports hundreds of production AI deployments. It also reports that across its customers, AI-handled volume grew 600% and the number of AI customers grew 200% in the year to August 2026. Public customer examples it cites include:

Customer	Reported outcome
Lovepop	18-second email replies; 94% overall CSAT; shopping assistant converts 25% of chats into orders
Stewart Golf	75% ticket automation; 1-minute response times
RealVNC	60% AI resolution; 90% backlog reduction; chat time-to-agent cut to 21 seconds from 100+
Rachio	1M+ customers supported; 65% instant resolutions; 95–99% accuracy
EVPassport	70% instant AI resolutions; human response in under 30 seconds
Good Eggs	CX stack costs cut by more than 60%

Measurement: Each program is baselined against the prior provider and measured on verified resolution, with 100% of conversations scored.

Commercial model: Managed AI software. A recurring platform fee, tiered on volume, includes the AIDE and AI Product Manager managed service, plus an AI usage fee per conversation. Human support is contracted separately. Crescendo does not offer FTE or seat pricing.

Future Vision & Roadmap

Crescendo's investment centers on having the improvement loop take on more of the work while people govern it:

- **Optimization:** Moving up a level of abstraction, from tuning prompts to working from business requirements to managing against KPIs, with humans approving every change.
- **Conversation data beyond CX:** Product, trust and safety, legal, and marketing teams get direct use of the same insights, such as evidence packets for review decisions and product bugs filed straight from support conversations.
- **Revenue:** Deeper personalization and attribution for Concierge's shopping, retention, and upsell agents.
- **Skills and integrations:** A growing library of reusable skills and deeper integration with enterprise systems.
- **Voice and cross-channel continuity:** Identity, memory, and context carry across every channel.

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. **www.opusresearch.net**

For sales inquiries please e-mail info@opusresearch.net or call +1(415) 505-2511

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 inadequacies 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 November 2026 © Opus Research, Inc. All rights reserved.