

Agent assist across 1,200 BPO seats — a co-sell success

Case study · BPO · Agent assist

A global BPO wanted to introduce agent assist across three of their largest client programs without disturbing the existing SLA structure. pronix.ai delivered the platform, the change program and the client co-sell motion — 22% AHT reduction, 14% QA lift, and a repeatable playbook the BPO now sells into their own book.

CLIENT

Global BPO, 1,200 seats across three programs

INDUSTRY

BPO

PLATFORM

 NICE CXONE

 KORE.AI

 AZURE AI FOUNDRY

-22%

AHT

+14%

QA scores

3

New client wins attributed to the capability

6 wk

Time to first live program

01 The situation

A global BPO wanted to introduce agent assist across three of their largest client programs without disturbing the existing SLA structure. pronix.ai delivered the platform, the change program and the client co-sell motion — 22% AHT reduction, 14% QA lift, and a repeatable playbook the BPO now sells into their own book.

Three programs on NICE CXone, three different clients, three different QA rubrics — and a contractual expectation that any AI rollout would not shift the SLA calculus. The BPO also wanted to keep the resulting IP as a differentiator for future client pitches.

-22%

AHT

+14%

QA scores

3

New client wins attributed to the capability

6 wk

Time to first live program

02 How we delivered it

STEP 01

One platform, three tuning tracks

A single Kore.ai and Azure AI Foundry deployment with per-client prompt libraries, knowledge sources and QA models. Shared infrastructure, isolated content.

STEP 02

SLA-safe rollout

Shadow-mode for 21 days per program, then A/B by team for 21 more. No agent went to production suggestions until each program hit the pre-agreed guardrails.

STEP 03

QA lift, not QA replacement

Automated QA scored 100% of interactions and flagged the bottom decile to human reviewers — QA team capacity became a lever, not a constraint.

STEP 04

Co-sell packaging

We packaged the platform, the change model and the results into a two-page co-sell brief the BPO now uses in every new-logo pitch.

03 Reference stack and assumptions

LAYER	COMPONENT	DESIGN ASSUMPTION
Contact center	NICE CXone	Voice + digital already live across all three programs; agent desktop is CXone-native.
Conversational AI	Kore.ai XO	One tenant, per-program bots and knowledge sources; shared intent taxonomy for reporting.
LLM / evaluation	Azure AI Foundry (Azure OpenAI)	GPT-4o-class model for suggestions; cheaper 4o-mini for classification; prompt caching enabled.
Knowledge	Per-program SharePoint + curated FAQ	Grounding sources versioned per client; retrieval scoped by tenant.
QA & eval	Azure AI Foundry evaluations + human calibration	100% automated QA; 5% analyst dual-score for drift monitoring.
Data & attribution	Snowflake + BPO's existing BI	Per-program AHT, QA, containment and cost surfaced weekly to client ops reviews.

Written for

BPO COO · VP Delivery · Client Partner — and the delivery leaders who have to make the same call on their own estate.

THE BOTTOM LINE

Every engagement ends in production, under one SLA. The team that builds it runs it — which is why these numbers are measured against a baseline captured before anything shipped, not estimated afterwards.

Run the same play on your estate. Book a working session with the practitioners who delivered this program, or model the business case first with our ROI calculators.

pronix.ai/book
pronix.ai/roi

Pronix.ai is a specialized AI & CX Systems Integrator. Leading with AI. Innovating with Digital. Representative outcome; results vary by client, scope and platform configuration. Figures are measured against a pre-engagement baseline. Third-party names and marks belong to their owners.