

32% forecast-accuracy gain for a top-10 BPO — AI-driven workforce management

Case study · BPO · AI Ops

A top-10 BPO's staffing model was rebuilt every Monday in spreadsheets and always wrong by Wednesday. pronix.ai shipped an AI forecasting engine with intraday re-optimization across 60 sites — 32% forecast-accuracy improvement, 9% shrinkage reduction and \$11M in avoided overstaffing.

CLIENT

Top-10 global BPO, 60+ contact center sites

INDUSTRY

BPO

PLATFORM

NICE CXONE

Snowflake SNOWFLAKE

DATABRICKS

+32%

Forecast accuracy (MAPE)

-9%

Shrinkage

\$11M

Annual staffing cost avoided

3d → 4h

Weekly WFM cycle time

01 The situation

A top-10 BPO's staffing model was rebuilt every Monday in spreadsheets and always wrong by Wednesday. pronix.ai shipped an AI forecasting engine with intraday re-optimization across 60 sites — 32% forecast-accuracy improvement, 9% shrinkage reduction and \$11M in avoided overstaffing.

Forecast MAPE ran at 22-28% by program, and the WFM team spent three days a week rebuilding models by hand. Overstaffing cost was hidden inside client billing and no one could isolate it.

+32%

Forecast accuracy (MAPE)

-9%

Shrinkage

\$11M

Annual staffing cost avoided

3d → 4h

Weekly WFM cycle time

02 How we delivered it

STEP 01

Program-level feature store

Consolidated seven years of interaction, roster and outcome data in Snowflake with a shared feature store — one place to compute drivers.

STEP 02

Hierarchical forecasting

Trained per-program models with cross-program regularization — better accuracy on low-volume queues without overfitting.

STEP 03

Intraday re-optimizer

Every 30 minutes the model re-forecast the next 4 hours and proposed shift swaps and voluntary time-off offers to supervisors.

STEP 04

Shrinkage attribution

Made unplanned shrinkage visible per site, per shift, per reason — closing the loop with operations.

STEP 05

Client-billable transparency

Every staffing decision tagged to a client, a program and a driver — clean cost attribution for pass-through billing.

03 Reference stack and assumptions

LAYER	COMPONENT	DESIGN ASSUMPTION
Contact center	NICE CXone (WFM module retained)	Schedules and adherence stay in CXone; AI forecasts flow in via WFM Integration Hub.
Data platform	Snowflake	Seven years of interaction, roster, outcome and shrinkage data unified in one feature store.
ML platform	Databricks (MLflow + Feature Store)	Per-program hierarchical models trained weekly; model registry gates promotion to production.
Intraday orchestration	Custom scheduler on AWS ECS	30-minute re-forecast horizon; proposes shift swaps and VTO offers to supervisors via CXone.
BI / attribution	Snowflake + Sigma Computing	Client-billable staffing tagged per program and driver; pass-through billing reconciled monthly.
Change / adoption	Supervisor cockpit in Microsoft Teams	Supervisors accept/decline AI proposals in Teams; adoption instrumented as a leading metric.

Written for

BPO COO · VP Workforce Management · CFO — 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.