PITON Global

EXECUTIVE WHITE PAPER · WP-01

CALL CENTER OUTSOURCING · PHILIPPINES

Holding the line: the executive economics of call center outsourcing to the Philippines.

An analysis of market structure, fully-loaded seat economics, benchmark performance, and vendor-selection discipline in the world's largest voice-services market.

AUTHORS

John Maczynski — Chief Executive Officer
Ralf Ellspermann — Chief Strategy Officer

JULY 2026
MANILA · ESTABLISHED 2001

WP-01 · CALL CENTER OUTSOURCING

Contents

ABOUT THIS SERIES

The PITON-Global Executive White Paper Series is a companion library to the firm's advisory practice. Each volume examines one outsourcing service line or industry vertical in depth — its economics, benchmarks, vendor landscape, and governance — and substantiates the performance figures published on piton-global.com with the operational and financial analysis behind them.

THE PITON-GLOBAL LEADERSHIP TEAM

About the authors

John Maczynski

Chief Executive Officer

Over a thirty-plus-year career, John has directed contact-center programs from the buyer's chair — commissioning, scaling, and turning around offshore operations for some of America's largest corporations — before dedicating himself to guiding companies into the Philippine market. Every PITON-Global engagement runs under his personal direction, from the first requirements workshop through RFP, commercial negotiation, and go-live.

j.maczynski@piton-global.com
+1 402-598-8740

Ralf Ellspermann

Chief Strategy Officer

Few Western executives can match Ralf's tenure in Philippine outsourcing: he arrived in Manila when the sector was still measured in thousands of seats, and has spent the decades since running and advising delivery operations spanning financial services, insurtech, health services, travel, and online retail. The firm's benchmarking models, provider intelligence, and vetting methodology carry his signature.

r.ellspermann@piton-global.com
Manila, Philippines

PITON Global

ABOUT PITON-GLOBAL

Operating from Manila since 2001, PITON-Global has advised on Philippine outsourcing longer than any comparable practice and has collected numerous industry honors along the way. Its work is strictly buyer-side: mapping a provider landscape of more than a thousand operations, qualifying the small fraction that truly fit a given program, and staying at the client's side as an independent advisor until the engagement is live and stable. Companies that have sourced teams through the firm include Tripadvisor, TSYS, Garmin, HealthPartners, Yetipay, Workrise, and Norman Disney & Young.

SECTION 01

Executive summary

Voice remains the highest-stakes customer channel: it carries the interactions that chat deflection cannot absorb — the anxious, the angry, the high-value, and the ready-to-buy. It is also the channel where operational failure is most visible. A center that misses its service level does not merely inconvenience callers; it abandons revenue in the queue, in measurable dollars, every interval of every day.

The Philippines answers more customer calls for Western brands than any other country, and has since it overtook India in 2010. Yet the market's scale is precisely the buyer's problem: more than 1,000 BPO operations compete for every program, and the spread between the top decile and the market average — in service level, abandonment, first-call resolution, and attrition — is wider than the spread in price. Picking on price alone is how buyers convert a 60-percent cost saving into a customer-experience liability.

This paper sets out the economics an executive needs before signing a voice contract: the fully-loaded cost stack behind the per-hour quote, the benchmark performance the top 1% of Manila centers sustain, the selection discipline that reliably finds them, and a reconstructed engagement showing how an 80-seat blended program produced a 6.4× first-year return. Five findings frame the analysis:

1 **The arbitrage is intact — and it is not the point.** A fully-loaded Manila voice seat runs 60–70% below its US equivalent. The durable value, however, comes from what disciplined operators do with that headroom: tenure, training, and QA density an onshore P&L cannot fund.

2 **AI has re-priced the easy call, not the hard one.** Tier-0 automation removes password resets and order status from the queue. What remains is longer, harder, and more valuable per contact — raising, not lowering, the premium on agent quality.

3 **Service level is an economic variable, not an ops metric.** Moving a queue from 72% to 88% (80/20) at constant volume recovers abandoned-call revenue and conversion that typically exceeds the entire cost of the offshore program's governance overlay.

4 **Vendor selection determines 70% of program outcome before go-live.** Across PITON-Global's engagement base, programs matched on industry, scale-fit, and compliance posture reach steady-state roughly twice as fast as programs sourced on rate cards.

5 **Right-sizing beats brand-sizing.** For 10-to-100-seat programs, a right-sized specialist with global-grade compliance consistently outperforms a global BPO on governance access, flexibility, and executive attention — at a comparable or lower rate.

The intended reader is a CEO, COO, CFO, or CX leader weighing a first Philippine voice program or repairing an underperforming one. Sections 2–4 supply the market and cost evidence; sections 5–7 supply the method.

SECTION 02

The 2026 Philippine voice market: scale, maturity, and the AI reset

The Philippine IT-BPM industry closed 2025 at roughly $40 billion in revenue and enters 2026 on a trajectory toward $42 billion, employing about 1.9 million people. Voice-centric customer experience work remains its largest single component. The industry's quarter-century compounding — from roughly $1 billion when PITON-Global was established in 2001 — is the single most important fact about it: this is not an emerging destination but a mature industrial ecosystem, with three generations of team leaders, workforce planners, and QA professionals promoted from the phones.

FIG. 1 — PHILIPPINE IT-BPM REVENUE, SELECTED YEARS (US$B)

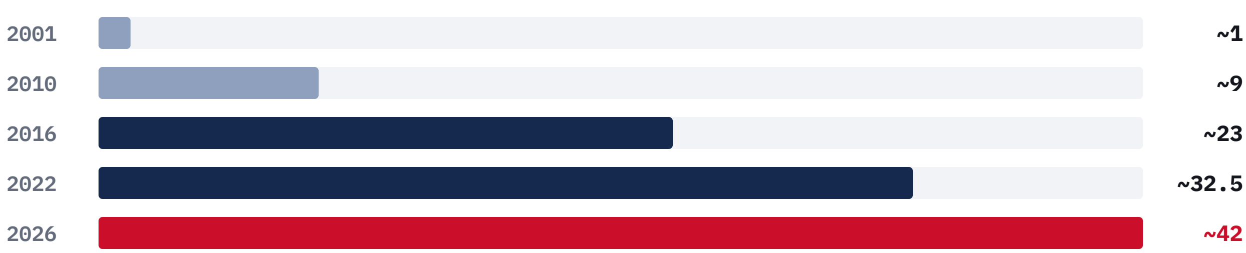

Source: IBPAP industry reporting and PITON-Global market analysis, July 2026. 2026 = projection.

The supply side is deep but radically uneven. PITON-Global's provider mapping counts more than 1,000 BPO operations nationwide — roughly 200–300 captive centers and 700–800 third-party providers — spread across Manila, Cebu, and eight further delivery hubs including Davao, Clark, Iloilo, and Bacolod. They range from 25,000-seat global majors to 40-seat boutiques. For any specific program, PITON-Global's vetting work shows the field collapsing to a handful — typically six to ten operators, under 1% of the market — once industry track record, scale fit, compliance posture, and commercial terms are applied as filters.

The AI reset of 2024–26 has restructured demand rather than reduced it. Voice AI and tier-0 self-service now absorb the shortest, most scriptable contacts. The calls that remain — disputed claims, complex troubleshooting, retention saves, high-value sales — carry longer handle times, higher emotional stakes, and more revenue per contact. Buyers in 2026 are therefore purchasing judgment, empathy, and conversion skill, delivered by tenured agents inside AI-augmented workflows. That shifts even more of the outcome onto which provider is selected: commodity centers staffed for yesterday's call mix are being repriced out of the market, while the top tier is investing in agent-assist tooling, real-time QA, and analytics that most buyers could not justify building in-house.

SECTION 03

The true cost stack: what a voice seat actually costs

Most outsourcing business cases fail at the same point: they compare an offshore hourly rate against an onshore wage, and both numbers are wrong. The economically honest comparison is between fully-loaded costs per productive hour — wages plus benefits, facilities, technology, management overhead, and the perennially omitted line item, attrition and backfill. Table 1 reconstructs that stack for a US metro center and a PITON-Global-vetted Manila operation.

FIG. 2 — THE FULLY-LOADED COST STACK PER PRODUCTIVE AGENT HOUR (US$, 2026)

Basis: PITON-Global engagement pricing data, 2025–26; US metro Tier-2 city baseline; productive hour = paid hour net of shrinkage at 30% (US) / 28% (PH).

Three observations matter more than the headline 64%. First, the largest onshore penalty after wages is attrition: a US center running 45–60% annual turnover pays for recruitment, training weeks, and the productivity ramp of every replacement — nearly five times the Manila figure, where vetted centers hold voice attrition near 25–35% and tenure routinely exceeds two years. Second, technology is the smallest gap in the stack: top Philippine centers run the same Genesys, Five9, NICE CXone, and Twilio Flex platforms under SOC 2 Type II and PCI-DSS controls, so buyers concede little on the stack itself. Third, the spread funds quality. A buyer who banks the entire 64% has bought a cheaper version of their current problem; the disciplined move is to reinvest 5–8 points of the saving into higher agent-to-QA ratios, retention bonuses, and speech analytics — expenditures that determine whether Section 4's benchmarks are achievable.

Beware the two commercial traps at contract stage: per-minute pricing that quietly shifts occupancy risk to the buyer, and "blended rate" quotes that average senior and trainee agents while the roster skews to trainees. Both are detectable in the RFP — if the RFP asks.

SECTION 04

Performance benchmarks: what separates the top 1%

Averages flatter the market. Table 2 contrasts the operating performance PITON-Global-vetted centers sustained across 2025–26 engagements with the blended onshore/commodity-offshore baseline a buyer inherits by default. The deltas compound: a queue that answers faster also abandons less, converts more, and generates fewer repeat calls.

TABLE 2 — VOICE OPERATING BENCHMARKS, VETTED VS. BASELINE (2025–26)

METRIC	VETTED TOP 1%	BASELINE	EXECUTIVE READING
Service level (80/20)	**88%**	~72%	Revenue reached vs. lost in queue
Average speed of answer	**18s**	~48s	Caller effort & first impression
Abandonment rate	**2.8%**	~7.5%	Each point ≈ lost sales + repeat demand
First-call resolution	**81%**	~68%	Deflects tomorrow's call volume
Agent occupancy	**85%**	~72%	Paid hours doing productive work
Annual voice attrition	**25–35%**	45–60%	Tenure is the quality flywheel

Source: PITON-Global call center operating data, 2025–26 engagements; baseline = onshore and generic-offshore voice averages.

What produces the left-hand column is unglamorous and highly replicable — which is exactly why it discriminates so well between providers. Four practices recur in every top-tier center PITON-Global vets: **interval-level Erlang-C staffing** against a forecast the client can audit, rather than flat headcount; **real-time queue management** with authority to flex breaks, training, and overflow inside the interval; **QA density of one evaluator per 12–15 agents** with calibration against the client's own rubric; and **a tenure economy** — career ladders and retention pay that keep the median agent past the 24-month mark, where handle time, FCR, and conversion all inflect.

None of these practices is visible on a rate card, and only some are visible on a site tour. They surface under structured diligence: workforce-management walkthroughs, QA calibration sessions, attrition cohort data, and reference calls with clients of comparable size and industry — the work Section 5 systematizes.

SECTION 05

Sourcing discipline: the 7-Step Vendor Vetting Framework, applied to voice

PITON-Global's vetting framework exists because voice programs fail at selection more often than at operation. Applied to a call center engagement, the seven steps translate as follows — each one eliminating a class of provider that would otherwise reach the shortlist on charm or price alone.

FIG. 2 — THE 7-STEP FRAMEWORK AS A VOICE-PROGRAM FUNNEL

	STEP	WHAT IT ELIMINATES IN A VOICE SEARCH	FIELD REMAINING
1	**Requirements definition**	Vague scopes that invite lowest-common-denominator bids; fixes call mix, hours, languages, SL targets, compliance set	1,000+
2	**Market mapping**	Providers outside the required hubs, scale band, or delivery model (onsite / hybrid / WFH mix)	~120
3	**Capability screening**	Centers without the industry track record, platform experience (Genesys, Five9, NICE, Twilio Flex), or language coverage	~40
4	**Compliance & security audit**	Anyone who cannot evidence SOC 2 Type II, PCI-DSS, TCPA/DNC process controls, BCP and site resilience	~20
5	**Operational diligence**	Weak WFM discipline, thin QA density, ugly attrition cohorts, references that don't check out	6–10
6	**Competitive RFP & negotiation**	Pricing games (per-minute traps, blended-rate rosters); finalists compete on transparent, comparable terms	2–3
7	**Transition & launch governance**	Launch drift: PITON-Global stays embedded on the client's side from day one to steady state, then hands over	1

Field-remaining counts are representative of a mid-market voice search (20–100 seats), PITON-Global engagements 2024–26.

Two design principles distinguish the framework from a broker's matchmaking. First, it is **vendor-neutral by construction**: PITON-Global holds no inventory and carries no placement quota, so no provider can buy its way onto a shortlist. Second, it is **right-sizing-led**: for a 10-to-100-seat program the framework deliberately surfaces mid-sized specialists with global-grade compliance — providers for whom the program is a flagship account commanding daily executive attention, not a rounding error routed to a B-team. The output is a shortlist of six to ten best-qualified centers delivered within 48–72 hours, competing for the business through a fully transparent, fair, comprehensive, and highly competitive RFP.

SECTION 06 · ANONYMIZED CASE STUDY

From 71% to 88% service level: an 80-

seat blended program, reconstructed

The engagement behind this paper's headline figures — designated CC-070 in PITON-Global's records — is a US consumer-services company with roughly $180 million in revenue, whose in-house center of 62 agents in a Tier-2 US city was missing service level in 9 of 12 months (71% average against an 80/20 target), abandoning 7.9% of calls, and losing a third of its agents annually. The board's mandate was blunt: fix the phones without growing the cost line. Identifying details are anonymized under NDA; the operating and financial mechanics below are typical of engagements at this scale.

SELECTION (WEEKS 0–8)

Requirements fixed a blended inbound-service / outbound-winback scope, 6am–10pm CT coverage, PCI-DSS handling for card payments, and TCPA/DNC discipline on the outbound side. From 1,000+ providers, mapping yielded 118 candidates; capability and compliance screens cut the field to 19; operational diligence produced a shortlist of 8 Manila and Cebu specialists in the 800–3,000-seat band. A three-round RFP settled on a Manila operator with deep consumer-services experience — at $12.40 per productive hour, 11% below the midpoint bid, with an SL-linked service credit regime.

TRANSITION (WEEKS 9–26)

An 80-seat build (62-seat like-for-like plus queue-depth headroom) ramped in three waves. Knowledge transfer ran against the client's own QA rubric from day one; Erlang-C staffing was rebuilt on 15-minute interval forecasts, exposing chronic Monday under-staffing the in-house WFM had never resolved. PITON-Global chaired weekly governance through launch, holding the provider to ramp gates before each wave released. Steady state was certified in month 6; only then did day-to-day governance hand over to the client.

TABLE 3 — CC-070 FIRST-YEAR VALUE BUILD-UP (US$)		
VALUE COMPONENT	YEAR 1	BASIS
Labor & overhead arbitrage	$2.71M	126k productive hrs × ($33.60 – $12.10)
Abandoned-call revenue recovered	$0.86M	7.9%→2.8% abandon on 640k calls; 14% sales-line share × $58 AOV × recovery factor
Outbound winback conversion lift	$0.54M	Winback conversion 9.2%→13.1% on tenured outbound pod
Repeat-call volume avoided (FCR 68%→81%)	$0.31M	~52k repeat calls avoided × $6.00 marginal cost
Gross value created	**$4.42M**	
Less: engagement & transition costs	($0.69M)	Transition, travel, parallel-run, governance overlay; advisory fee $0 (provider-paid model)
Net value ÷ cost = first-year ROI	6.4×	$4.42M ÷ $0.69M

Figures rounded; revenue effects use conservative recovery factors (not every abandoned call is a lost sale). Method detailed in Methodology & notes.

The instructive point is the shape of the return. Labor arbitrage supplied about 61% of first-year value; the remaining 39% — abandoned-call recovery, conversion lift, repeat-call avoidance — existed only because selection screened for operators capable of running 88/18/2.8 queues. A rate-card sourcing process would have captured the first number and forfeited the second, which is the difference between a cost project and a 6.4× value program.

SECTION 07

The executive playbook: governance, risk, and the first 180 days

Selection creates the possibility of the Section 6 result; governance realizes it. The playbook below condenses what PITON-Global installs, on the client's side of the table, in every voice engagement. It assumes the advisory posture the firm considers non-negotiable: the client calls every shot, and the advisor's job is to make sure nothing goes unexamined.

DAYS 0–30

Contract for outcomes

Tie 10–15% of monthly fees to a small SLA set — service level, abandonment, FCR, QA score — with service credits that bite and a cure path that is operational, not legal. Price per productive hour, never per minute. Lock the named-roster clause: the agents and leaders who won the RFP are the ones who staff the program.

DAYS 30–90

Ramp against gates, not dates

Release each hiring wave only when the prior wave clears QA and proficiency gates. Run weekly governance chaired by someone with authority to stop the ramp. Instrument the queue at 15-minute interval granularity from the first day of live traffic — averages hide the intervals where service level is lost.

DAYS 90–180

Convert stability into value

Once SL holds for 60 consecutive days, shift governance attention from delivery to yield: conversion coaching on sales queues, FCR root-cause work, AI agent-assist deployment where handle time concentrates. Certify steady state, then hand over day-to-day governance deliberately — a dated, criteria-based handover, not a fade-out.

THE FIVE FAILURE MODES WE ARE HIRED TO PREVENT

Sourcing on rate card instead of scale-fit · accepting blended-rate quotes against trainee-heavy rosters · contracting per-minute and donating occupancy risk · skipping reference diligence on attrition cohorts · letting the provider chair its own governance. Every one is avoidable before signature, and expensive after it.

The through-line of this paper is that call center outsourcing to the Philippines in 2026 is neither a gamble nor a commodity purchase — it is a solvable selection problem with quantifiable returns. The market offers genuine top-1% operators; the cost stack funds quality rather than merely discounting it; and a disciplined, vendor-neutral process reliably separates the handful of operators genuinely built for a given program from the thousand-strong field competing for it. That is the difference between guessing and knowing — and it is a difference available before signature, never after.

METHODOLOGY & NOTES

How the numbers were built

Operating benchmarks (service level, ASA, abandonment, FCR, occupancy, attrition) are drawn from PITON-Global-vetted engagements active during 2025–26, measured on provider ACD and QA systems and reviewed in client governance. Baselines combine client-reported pre-engagement performance with onshore and generic-offshore industry averages; they describe the position a typical buyer inherits, not any single center.

Cost figures are fully-loaded costs per productive hour: wages, benefits and statutory costs, facilities, telephony and security stack, management/WFM/QA overhead, and attrition-driven recruiting and ramp costs, net of shrinkage. US baseline reflects a Tier-2 metro center; Manila figures reflect 2025–26 PITON-Global engagement pricing. Exchange rate assumptions are held constant within the modeled year.

The case study (CC-070) is anonymized under NDA: industry descriptors are generalized and figures rounded. Revenue-side effects apply conservative recovery factors — abandoned-call recovery assumes only the sales-line share of abandoned calls at observed conversion, not all of them; winback lift is measured against the client's own trailing-12-month baseline. ROI = net first-year value ÷ all engagement-related costs borne by the client. PITON-Global's advisory fee is provider-paid and appears as $0 in the client's cost base; the firm's vendor-neutrality safeguards for that model are documented at piton-global.com.

Market data combine IBPAP industry reporting with PITON-Global's proprietary provider mapping (1,000+ operations, updated July 2026). Figures marked as estimates or projections are PITON-Global analysis and should be cited as such.

Which providers deserve a place on your shortlist?

Forty-five minutes with John will tell you. No cost to buyers, no vendor allegiance — a curated shortlist within two to three business days.

j.maczynski@piton-global.com
+1 402-598-8740