PITON Global

EXECUTIVE WHITE PAPER · WP-06

AI CALL CENTER OUTSOURCING · PHILIPPINES

The blended floor: the executive economics of AI call center outsourcing to the Philippines.

An analysis of how the human+AI floor actually divides its labor, why containment quality beats containment rate, the escalation contract between bot and agent, and vendor-selection discipline for the operating model every provider now claims to run.

AUTHORS

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

JUNE 2026
MANILA · ESTABLISHED 2001

Contents

ABOUT THIS SERIES — VOLUME 6

Every provider deck in 2026 has an AI slide. Almost none has an AI operating model. This volume is about the difference — the specific division of labor between machine and agent that separates floors where AI compounds the economics from floors where it merely decorates the proposal. The full series lives at piton-global.com.

THE PITON-GLOBAL LEADERSHIP TEAM

About the authors

John Maczynski

Chief Executive Officer

John has sat through three decades of contact-center technology cycles from the buyer's chair — IVR, CTI, chatbots, and now generative AI — and signed the budgets that separated the tools that worked from the demos that didn't. His filter has never changed: show me the operating model, not the roadmap. That filter, applied at PITON-Global, is what clients get on every AI-enabled sourcing decision, personally, without an account layer in between.

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

Ralf Ellspermann

Chief Strategy Officer

Ralf has spent twenty-five years running Philippine floors through every automation wave, and he judges them all by the same floor-level question: what happens in the ten seconds after the machine gives up? The escalation-contract framework in Section 3 codifies what he learned rebuilding operations where that handoff had been left to chance. He now reads AI-enabled providers the way he once ran them — from the escalation queue backward.

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

PITON | Global

ABOUT PITON-GLOBAL

PITON-Global is a buyer-side advisory practice, Manila-based since 2001, that has never operated a delivery floor or taken a placement commission. Its stock-in-trade is knowing the Philippine provider market at the operating level — 1,000+ mapped operations, re-verified continuously, now including each provider's actual AI maturity rather than its claimed one — and converting that knowledge into a shortlist of six to ten genuinely qualified partners per engagement. Buyers from Tripadvisor, Garmin, TSYS, HealthPartners, Yetipay, Workrise, and Norman Disney & Young have run their searches through the practice.

SECTION 01

Executive summary

-60%	28%	-31%	6.2×
Fully-loaded cost per productive hour vs. US onshore	Contact volume absorbed by the AI tier at stable CSAT, vetted floors	Handle-time reduction from agent-assist on escalated contacts	First-year ROI, reconstructed case (Section 6)

The AI call center is not a technology purchase; it is a division of labor. Somewhere in every blended floor a line is drawn: below it, the machine resolves; above it, a human does — with the machine now at their elbow. Where that line sits, how contacts cross it, and what condition they arrive in on the other side determine whether AI compounds the offshore cost advantage or quietly refunds it through repeat contacts and rage-dialed escalations.

The Philippines is where the blended model is being industrialized fastest, because the human tier there is deep enough to absorb what the machine hands back. Five findings:

1. **Containment rate is the industry's most gamed number.** Any bot can "contain" a customer who gives up. Vetted floors measure containment quality — resolved-and-not-recontacted within seven days — and run 24–32% true absorption. Weak floors quote 50%+ containment that dissolves under a recontact audit. Section 2 shows the gap.

2. **The handoff is the product.** CSAT on a blended floor is decided in the ten seconds after escalation: does the agent inherit the transcript, the account context, and the bot's failed attempts — or does the customer start over? The escalation contract of Section 3 is the single highest-leverage artifact in AI-floor diligence.

3. **Agent-assist beats agent-replace on current economics.** Real-time knowledge retrieval, draft responses, and auto-wrap cut escalated-contact handle time 25–35% on instrumented floors — value that arrives quarter one, carries no CSAT risk, and compounds the arbitrage rather than gambling it.

4 **The human tier gets harder, not smaller-skilled.** What escalates past a competent AI tier is the angry, the ambiguous, and the high-stakes. Floors that re-graded their hiring, coaching, and pay for that mix hold CSAT through the transition; floors that kept FAQ-tier staffing watch it fall exactly as their volume "improves."

5 **Pay for outcomes, not licenses.** The contracting failure of the AI era is paying a technology fee for tooling the provider never operationalizes. Vetted deals price the blended floor per resolved contact or per absorbed volume band, with containment-quality gates — making the AI's performance the provider's problem, which is where it belongs.

The intended reader runs CX, operations, or technology at a company being pitched "AI-powered" outsourcing weekly. Sections 2–4 separate the operating models from the slideware; sections 5–7, the method and the proof.

SECTION 02

The division of labor: what the AI tier actually absorbs, and what it hands back

Figure 1 follows 100 inbound contacts through a vetted blended floor — and, in the right column, through the same volume as a weak vendor reports it. The difference between the two columns is the difference between an operating model and a marketing claim.

FIG. 1 — 100 CONTACTS THROUGH THE BLENDED FLOOR: TRUE ABSORPTION VS. REPORTED CONTAINMENT

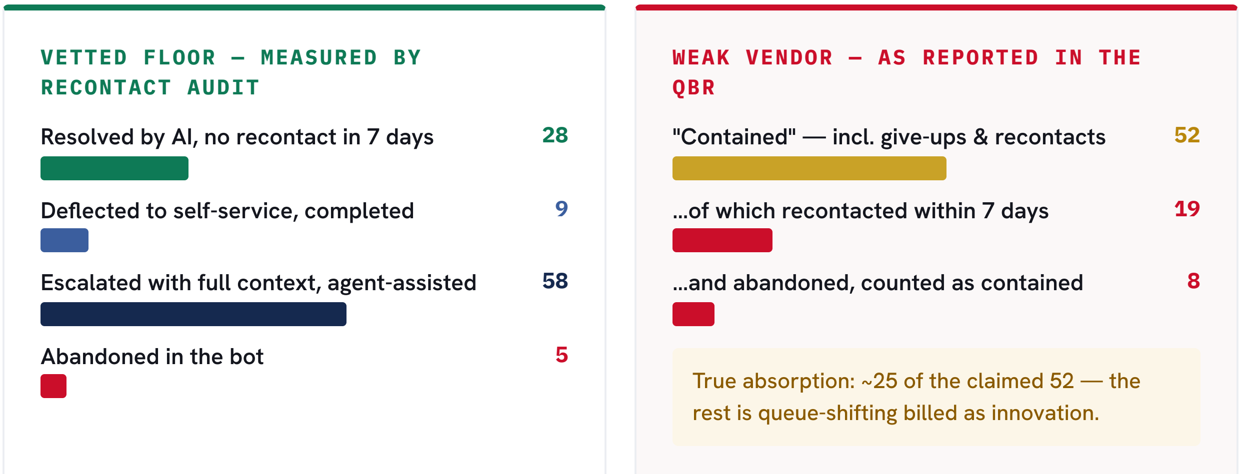

Left: PITON-Global engagement instrumentation, 2025–26, recontact-audited. Right: composite of pre-engagement vendor reporting reviewed in diligence.

The economics follow the left column. Twenty-eight truly absorbed contacts convert directly into staffed-hour reduction; nine completed deflections cost nearly nothing; and the fifty-eight escalations arrive with context, which is what lets agent-assist cut their handle time by a third. The right column produces the opposite: phantom savings booked in the QBR, real costs returning as repeat volume, and a CSAT chart management cannot explain. The recontact audit — seven lines of SQL any honest provider will run for you — is the cheapest due-diligence instrument in this paper.

SECTION 03

The escalation contract: containment quality, agent-assist, and the handoff that decides CSAT

Figure 2 lays out the three clauses of a working escalation contract — the machinery a floor read verifies before any AI line item is believed.

FIG. 2 — THE ESCALATION CONTRACT, AS OPERATED ON VETTED FLOORS

KNOW WHEN TO QUIT

The bot escalates on confidence, sentiment, and stake — two failed turns, detected frustration, or any account-risk intent — never after exhausting the customer. Escalation triggers are tuned monthly against recontact data.

HAND OVER THE FILE

The agent inherits transcript, account context, intent classification, and what the bot already tried — on screen before the greeting. The customer never repeats themselves. This single clause moves post-escalation CSAT 10+ points.

ASSIST, THEN LEARN

Agent-assist retrieves knowledge, drafts replies, and auto-wraps; every human resolution is tagged and fed back to expand what the AI tier can absorb next quarter. The boundary moves by evidence, not by roadmap.

↓ yields ↓

THE COMPOUNDING FLOOR

28% TRUE ABSORPTION | −31% AHT ON ESCALATIONS | CSAT HELD THROUGH TRANSITION

Absorption grows a few points a quarter as the feedback loop expands the bot's competence — while the human tier, re-graded for the harder residual mix, holds the experience. The floor gets cheaper and better simultaneously, which is the entire promise of the blended model.

Operating parameters from PITON-Global AI-enabled engagements, 2025–26.

In diligence, the contract is tested from the customer's side: call the provider's live client lines, fail the bot on purpose, and time what happens next. Did the agent know what you had already tried? Did you repeat your account number? Was the tone of the first human sentence an apology or a continuation? Perhaps one provider in twelve passes cold. Steps 3 and 5 of the framework exist to find that one before the contract does — because everything in Table 3 of Section 6 depends on it.

SECTION 04

Cost and performance: blended-floor

economics and top-tier benchmarks

TABLE 1 — COST PER 1,000 RESOLVED CONTACTS: THREE OPERATING MODELS (US$, 2026)		
OPERATING MODEL	COST / 1,000 RESOLVED	VS. US BASELINE
US onshore, human-only floor	$8,900	baseline
Vetted Manila floor, human-only	$3,560	–60%
Vetted Manila blended floor — 28% AI absorption + agent-assist	$2,340	–74%

Basis: PITON-Global engagement pricing 2025–26; resolved contact = closed without recontact in 7 days; AI platform fees included in blended figure; US comparator metro contact-center labor, fully loaded.

TABLE 2 — AI-ENABLED FLOOR BENCHMARKS, VETTED VS. BASELINE (2025–26)			
METRIC	VETTED TOP TIER	BASELINE	EXECUTIVE READING
True absorption (recontact-audited)	24–32%	10–18%	The only containment number worth paying for
Bot-abandonment rate	<6%	15–30%	Every abandonment is a customer punished for calling
Context-intact escalation	95%+	30–60%	"Hand over the file" — kept or broken
AHT on escalated contacts (agent-assist)	–25–35%	–0–10%	Where AI pays before it absorbs
CSAT through AI transition	held ±1 pt	–5–12 pts	The human tier re-graded, or not

Source: PITON-Global AI-enabled engagement data 2025–26; baseline = typical pre-engagement operations with vendor-led AI deployments.

The strategic read: the blended floor's –74% is not one decision but two, made in the right order. Selecting the right human floor delivers –60% with certainty; the AI layer adds a further 14 points only when the operating model of Sections 2–3 exists. Buyers who reverse the order — chasing the AI number first, from a weak floor — end up with the right column of Figure 1: a discount that recontacts its way back to full price.

SECTION 05

Sourcing discipline: the 7-Step Vendor Vetting Framework, applied to AI-enabled floors

AI diligence adds one hard requirement to every step: claims must survive instrumentation. The funnel below is representative of a 30–80-seat AI-enabled search.

FIG. 3 — THE 7-STEP FRAMEWORK AS AN AI-FLOOR FUNNEL

	STEP	WHAT IT ELIMINATES IN AN AI-ENABLED SEARCH	FIELD REMAINING
1	**Requirements definition**	Technology-first scoping: defines the absorption boundary by intent — what the AI tier may own, what it must never touch, and the containment-quality bar that gates every expansion	1,000+
2	**Market mapping**	Slideware AI: floors whose "AI practice" is a reseller agreement, and giants whose AI attention goes to nine-figure accounts	~45
3	**Capability screening**	Floors without the machinery: recontact-audited containment reporting, context-intact escalation, agent-assist in production (not pilot), a feedback loop that has actually moved the boundary	~12
4	**Compliance & security audit**	Weak AI governance: where transcripts go, what trains on client data, PII redaction before the model, human review of AI-issued account actions, SOC 2 scope covering the AI stack	~9
5	**Operational diligence**	The adversarial read: fail the bot on live client lines and time the handoff; pull recontact audits; watch agent-assist used mid-call; interview agents on what the AI gets wrong and what happens when it does	5–8
6	**Competitive RFP & negotiation**	License-fee pricing: finalists bid per resolved contact or absorption band, with containment-quality gates, abandonment ceilings, context-intact floors, and CSAT holds — all carrying credits	2–3
7	**Transition & launch governance**	Big-bang AI cutover: the human floor launches and stabilizes first; the AI boundary expands intent by intent behind the quality gate, with PITON-Global embedded buyer-side to steady state	1

Field-remaining counts represent a 30–80-seat AI-enabled search, PITON-Global engagements 2024–26.

The standing guarantees hold: **vendor neutrality** — PITON-Global operates no floors, resells no AI platforms, and takes no placement economics, so the shortlist reflects instrumentation, not inventory; **right-sizing** — providers for whom your program is a flagship account. Six to ten genuinely qualified providers, delivered within 48–72 hours, free to the buyer, competing through a fully transparent, fair, comprehensive, and highly competitive RFP.

SECTION 06 · ANONYMIZED CASE STUDY

Re-dividing the floor: a 50-seat human+AI rebuild, reconstructed

Engagement AC-025 is a US consumer subscription platform — roughly $420 million in revenue, 1.1 million support contacts a year — that had already been burned once: an incumbent vendor's "AI transformation" had booked a 47% containment rate while NPS fell nine points and repeat contacts rose a third. The board wanted the AI savings and refused the AI risk. Identifying details are anonymized under NDA.

SELECTION (WEEKS 0–8, INCL. ADVERSARIAL READS)

Requirements drew the absorption boundary by intent: password, billing-status and shipping intents eligible for AI ownership behind a containment-quality gate; cancellation, dispute, and account-risk intents human-only, permanently. Mapping produced 44 claimants; the machinery screen left 11; the AI-governance audit — including what trains on client transcripts — left 8. Adversarial reads (bot failed deliberately on live lines, handoffs timed, agents interviewed) shortlisted 6. A Cebu floor won at $10.90 blended human rate plus per-resolved-contact AI pricing, with a 7-day recontact audit written into the reporting spec.

TRANSITION (MONTHS 3–8, BOUNDARY MOVED BY EVIDENCE)

The human floor launched first and certified at month 4 — then the AI tier took its first intent, passwords, and held the quality gate for six weeks before earning the next. Agent-assist went live for all escalations from day one; the human tier's hiring profile and pay bands were re-graded for the harder residual mix. By month 8 the boundary covered nine intents at 27% true absorption, with post-escalation CSAT two points above the human-only baseline. PITON-Global chaired the boundary-review board to steady state.

TABLE 3 — AC-025 FIRST-YEAR VALUE BUILD-UP (US$)		
VALUE COMPONENT	YEAR 1	BASIS
Labor arbitrage on the 50-seat human tier	$1.61M	93.8k productive hrs × ($28.10 – $10.90)
AI absorption — 27% of volume at quality, net of platform fees	$0.54M	Absorbed contacts × blended human unit cost, recontact-audited, less AI pricing
Agent-assist AHT reduction on escalated contacts (–29%)	$0.27M	Avoided staffed hours at steady volume, net of ramp
Repeat-contact collapse vs. prior AI deployment (recontacts –38%)	$0.21M	Avoided contact volume vs. incumbent-era baseline, margin basis
Gross value created	**$2.63M**	
Less: engagement & transition costs	($0.42M)	Extended dual-running, boundary governance, re-graded hiring, integration; advisory fee $0 (provider-paid model)
Net value ÷ cost = first-year ROI	**6.2×**	$2.63M ÷ $0.42M

Figures rounded; absorption and recontact effects audit-verified. Method in Methodology & notes.

Sixty-one percent arbitrage, thirty-nine percent from the blended operating model — the highest AI-attributable share the series has recorded, earned precisely because the AI was sequenced second. Every line traces to a diligence step: the boundary design to Step 1, the honest containment reporting to Step 6's pricing, the intact handoffs to Step 5's adversarial read. The CX VP's line at the year-one review: "Last time, the AI hid our problems. This time it does the work."

SECTION 07

The executive playbook: governance for the first 180 days

An AI-enabled transition succeeds when the boundary moves by evidence and the humans are staffed for what remains — client in command throughout.

DAYS 0–30

Draw the boundary, contract the quality

Define AI-eligible and human-only intents in writing before launch. Contract per resolved contact or absorption band — never per license — with containment-quality gates, abandonment ceilings, and context-intact floors carrying credits. Stand up the 7-day recontact audit as the program's arbiter of truth.

DAYS 30–90

Stabilize the humans, assist before absorbing

Certify the human floor first; switch on agent-assist for every escalation from day one — it pays immediately and trains the feedback loop. Re-grade hiring and pay for the post-AI contact mix now, not after CSAT tells you. Publish the weekly dashboard: true absorption, abandonment, context-intact rate, CSAT by tier.

DAYS 90–180

Expand intent by intent, certify steady state

Let the AI tier earn each new intent behind the quality gate — six clean weeks before the next. Review the boundary monthly with recontact data on the table. Certify steady state on two clean months including one volume peak; hand over on documented criteria; keep the quarterly adversarial read standing.

THE FIVE FAILURE MODES WE ARE HIRED TO PREVENT

Paying for containment rate instead of containment quality · big-bang AI cutover onto an unstabilized floor · escalations that arrive without the file · keeping FAQ-tier staffing for a post-AI contact mix · licensing technology instead of contracting outcomes. All five are settled at selection and contract.

The AI call center argument, in one sentence: the blended floor is a division of labor, and divisions of labor are won by whoever assigns the work best — a vetted Philippine operation delivers –74% per resolved contact with CSAT held, for buyers who audit recontacts instead of believing containment slides and compete the finalists through a fully transparent, fair, comprehensive, and highly competitive RFP. The machine is ready for its share of the work. We make sure you buy the floor that knows which share that is.

METHODOLOGY & NOTES

How the numbers were built

Absorption and containment figures (Fig. 1, Table 2) derive from platform instrumentation in PITON-Global AI-enabled engagements, 2025–26. True absorption counts contacts resolved by the AI tier with no recontact on any channel within seven days; abandoned sessions are never counted as contained. The weak-vendor column is a composite of pre-engagement vendor reporting reviewed in diligence, shown for contrast.

Cost-per-resolved figures (Table 1) are fully loaded — wages, benefits and statutory charges, facilities, technology including AI platform fees, management and QA overhead, attrition-driven recruiting — net of shrinkage; US comparators reflect metro contact-center labor on the same basis. Agent-assist effects are measured as AHT deltas on escalated contacts at constant intent mix.

The case study (AC-025) is anonymized under NDA; volumes and dates are generalized, figures rounded. Absorption and recontact effects are audit-verified against the reporting spec; ROI = net first-year value ÷ all engagement-related client costs. PITON-Global's advisory fee is provider-paid and appears at $0 in the client cost base; neutrality safeguards — including that the firm resells no AI platforms — are documented at piton-global.com.

Market context draws on IBPAP reporting and PITON-Global's proprietary provider mapping (1,000+ operations; AI operating maturity re-verified May 2026). Estimates and projections are PITON-Global analysis and should be cited as such. The full series is available at piton-global.com.

Would your vendor's containment rate survive a recontact audit?

Thirty minutes with John will tell you. Free to buyers, strictly vendor-neutral — an instrumentation-verified shortlist within two to three business days.

j.maczynski@piton-global.com

+1 402-598-8740