

TECHNICAL SUPPORT OUTSOURCING • PHILIPPINES

Resolution rate: the executive economics of technical support outsourcing to the Philippines.

An analysis of tier-stack economics, cost per resolved ticket, FCR and MTTR benchmarks, and vendor-selection discipline for SaaS companies, device makers, and IT organizations sourcing in the Philippines.

AUTHORS

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

JULY 2026
MANILA • ESTABLISHED 2001

Contents

–	About the authors	03
01	Executive summary	04
02	The tier stack: where support cost lives, and the only metric that prices it honestly	05
03	Support is retention: the churn arithmetic executives skip	06
04	Cost and performance: role-band economics and top-tier benchmarks	07
05	Sourcing discipline: the 7-Step Vendor Vetting Framework, applied to technical support	08
06	Case study: a 55-seat SaaS support program, reconstructed	09
07	The executive playbook: governance for the first 180 days	10
–	Methodology & notes	11

ABOUT THIS SERIES

A good sourcing decision needs three things that vendor websites never supply: honest unit costs, benchmarks measured on real programs, and a sober count of how many providers can actually do the work. Each volume of this series delivers all three for one service line or industry, drawn from the engagement records of a quarter century of buyer-side advisory in the Philippine market.

About the authors



John Maczynski

Chief Executive Officer

Before founding an advisory practice, John spent three decades accountable for the outcomes this paper measures — running outsourced service and support operations for some of the largest companies in America, where a missed SLA had his name on it. That accountability is the firm's inheritance: he leads every engagement personally and holds every provider to the standard he once had to answer for.

j.maczynski@piton-global.com

+1 402-598-8740



Ralf Ellspermann

Chief Strategy Officer

Technology programs bookend Ralf's twenty-plus years in Manila: he ran some of the earliest offshore tech-support floors in the country and has since advised delivery organizations across software, devices, connectivity, and platforms. That history is why the diligence protocols in this volume ask the questions they ask — he has watched providers pass rate-card scrutiny and fail a severity-one incident. The firm's provider intelligence and benchmark models remain his charge.

r.ellspermann@piton-global.com

Manila, Philippines



ABOUT PITON-GLOBAL

PITON-Global exists because the Philippine outsourcing market is opaque to the people spending money in it. Since 2001 the Manila-based firm — many times honored by its industry — has closed that information gap for buyers exclusively: verifying what a thousand-plus providers can really do, gating the field against each client's requirements, running the competitive process, and staying at the client's elbow until the program performs. It holds no delivery assets and accepts no provider economics that could color a recommendation. Tripadvisor, TSYS, Garmin, HealthPartners, Yetipay, Workrise, and Norman Disney & Young are among the companies that have built their Philippine teams on that model.

Executive summary

\$6.20

Cost per resolved T1 ticket (vs. \$19.40 onshore, -68%)

78%

First-contact resolution at vetted providers (vs. ~62%)

-44%

Mean time to resolution vs. client baseline

5.6×

First-year ROI, reconstructed SaaS case (Section 6)

Technical support is the only customer function whose demand is generated by the company's own product decisions — every release, every integration, every device shipped mints future tickets. Yet most organizations price it by the seat and manage it by the queue, which is how a SaaS company ends up with \$80-an-hour engineers resetting passwords and a device maker's NPS hostage to hold music. The honest unit of account is the resolved ticket, at the lowest tier competent to resolve it — and by that unit, the Philippine market is not incrementally better but categorically so.

The country built its technical-support depth serving the world's telcos, ISPs, and device brands through the 2000s, and its SaaS bench through the decade since: agents who troubleshoot across stacks, speak API without flinching at Tier 2, and escalate with reproduction steps attached. AI has intensified rather than diminished the case — bots now absorb the password resets, and what remains is precisely the diagnostic work where the gap between a good agent and a warm body is widest. Five findings frame the volume:

- 1 Price the resolution, not the seat.** Cost per resolved ticket bundles wage, occupancy, FCR, and rework into one honest number — \$6.20 versus \$19.40 at Tier 1 — and it is the number finalists should bid in the RFP.
- 2 Every escalation is a tax.** The spread between an 18% and a 31% T1→T2 escalation rate is engineering payroll doing support work. Top providers attack it with knowledge-base discipline and diagnostic training, not heroics.
- 3 Support quality is priced in NRR.** For subscription businesses, MTTR and FCR move logo retention and expansion revenue in measurable increments — Section 3 does the arithmetic most board decks skip.
- 4 The 24/7 economics mirror travel's.** Software breaks on weekends too; Manila staffs the hours onshore teams dread, at flat rates, with the same time-zone geometry documented in WP-79.

- 5 Genuine technical depth is scarcer than headsets.** Of 1,000+ Philippine operations, PITON-Global's mapping confirms true multi-tier technical practices — T2 benches, KB disciplines, severity-incident machinery — at roughly forty-five. The other providers sell T1 with a thesaurus.

The intended reader runs support, customer experience, product, or finance at a SaaS company, device maker, connectivity provider, or IT organization. Sections 2–4 supply the evidence; sections 5–7, the method and the proof.

SECTION 02

The tier stack: where support cost lives, and the only metric that prices it honestly

Support is a routing problem wearing a staffing costume. Every ticket has a cheapest competent tier, and every organization leaks money in both directions — easy tickets reaching expensive people, hard tickets bouncing between cheap ones. Figure 1 prices the stack as a vetted offshore operation runs it.

FIG. 1 — THE TIER STACK: COST PER RESOLUTION AND THE ROUTING DISCIPLINE (US\$, 2026)

T0 • SELF-SERVE	KB articles, AI assistant, community — absorbs 40–60% of raw demand when the KB is treated as a product	≈ \$0.10 / resolution
T1 • FRONTLINE	Known-issue resolution on documented paths; the FCR battleground — 78% of what arrives here should end here	\$6.20 vs. \$19.40 US
T2 • TECHNICAL	Novel diagnosis: logs, API traces, environment reproduction — the tier that separates real providers from staffing firms	\$14.10 vs. \$42.70 US
T3 • ENGINEERING	Stays with the client. The stack's job is to protect this tier: every ticket stopped below saves an engineering hour	\$45+ / resolution

Fully-loaded cost per quality-passed resolution, PITON-Global engagement data 2025–26. Routing discipline — not wage rate — determines the blended figure.

Two implications follow. First, **the KB is infrastructure, not documentation**: every T0 deflection is a \$6 saving and every stale article is a repeat ticket, so vetted providers staff knowledge management as a named function with article-level analytics — a diligence item most buyers never think to check. Second, **escalation rate is the honest quality metric**: an offshore team that resolves cheaply but escalates freely has merely moved the cost upstream to the client's engineers. Top-tier floors run T1→T2 escalation near 18% with full reproduction steps attached to what does move — which is why client engineering teams learn to trust, rather than re-verify, what arrives.


The work mix across PITON-Global's technical-support engagements splits roughly 55% SaaS and platform support, 25% consumer devices and connectivity, and 20% internal IT help desk — three markets, one economics: the stack above, staffed by people who can actually climb it.

SECTION 03

Support is retention: the churn arithmetic executives skip

Subscription economics turn support quality into a revenue line, and the conversion is calculable. Figure 2 traces it for a representative mid-market SaaS company — \$150M ARR, 6,000 accounts — using effect sizes measured across PITON-Global's SaaS engagement base.

FIG. 2 – FROM TICKET METRICS TO REVENUE: THE CHAIN, QUANTIFIED (\$150M ARR MODEL)




Effect sizes from PITON-Global SaaS engagements 2024-26, conservatively factored; model assumes support-attributable churn at the low end of measured range.

The chain's weakest link is attribution, and disciplined programs fix that in the tooling: churn reasons tagged at cancellation, support history joined to account records, and a quarterly review where the support vendor sits in front of the retention numbers it influenced. Vetted providers welcome that exposure — it is their renewal argument — and the welcome itself is a diligence signal. A bidder who resists revenue-linked reporting is telling you what tier of provider they are.

For device brands and IT help desks the arithmetic changes currency but not logic: warranty-call deflection, return-rate reduction, and employee-hours-recovered play the role of NRR. In every case the executive conclusion is the same one this series keeps reaching from different directions: the cheapest support operation is rarely the lowest-bidding one — it is the one that resolves at the right tier, the first time.

SECTION 04

Cost and performance: role-band economics and top-tier benchmarks



Basis: PITON-Global technical-support engagement pricing 2025-26; fully loaded and net of shrinkage, as defined in Methodology & notes. 24/7 coverage at flat rates.

TABLE 2 – TECHNICAL SUPPORT BENCHMARKS, VETTED VS. BASELINE (2025–26)

METRIC	VETTED TOP TIER	BASELINE	EXECUTIVE READING
First-contact resolution	78%	~62%	Every point retires repeat volume
T1→T2 escalation rate	18%	~31%	Engineering payroll, protected
Mean time to resolution	–44%	reference	The metric churn reasons cite
CSAT / QA accuracy (audited)	90 / 97.8%	~81 / 94%	Quality that survives sampling
KB deflection share of raw demand	40–60%	~25%	The cheapest tier, actually used
Annual attrition, technical programs	24–32%	45–60%	Product fluency compounds in seat

Source: PITON-Global technical-support engagement data 2025–26; baseline = client pre-engagement and onshore/commodity-offshore averages.

The machinery is the series standard — calibrated QA at double market density, certification gates before live queues, tenure economics — plus two disciplines specific to this vertical: **product immersion as a permanent function** (release notes reach the floor before the release does, with KB updates and macro changes shipped in the same sprint), and **severity-incident drills** — because a status-page day is to a support operation what a hurricane is to a travel program, and the providers worth shortlisting have rehearsed it.

SECTION 05

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

The discriminating test in a technical-support search is the T2 bench: anyone can hire communicators, but a provider either has working diagnostic talent — demonstrable in a live troubleshooting exercise — or it has a slide claiming so. The funnel below is representative of a mid-market SaaS search.

FIG. 3 – THE 7-STEP FRAMEWORK AS A TECHNICAL-SUPPORT FUNNEL

STEP	WHAT IT ELIMINATES IN A TECHNICAL-SUPPORT SEARCH	FIELD REMAINING
1 Requirements definition	Seat-count thinking: fixes the tier map, ticket taxonomy, volumes by channel and severity, systems (Zendesk, Intercom, Jira, remote tools), FCR/MTTR targets	1,000+
2 Market mapping	Providers without technical delivery units in the required scale band and 24/7 model	~120
3 Capability screening	No real T2 bench (tested live, not asserted), no KB-management function, no comparable-stack experience	~45
4 Compliance & security audit	Weak access architecture for production systems and customer data: SOC 2 on the delivery unit, least-privilege tooling, session logging, BCP	~22
5 Operational diligence	Escalation rates nobody measures, severity drills never run, references that hedge on incident days, attrition that churns product fluency	6–10
6 Competitive RFP & negotiation	Seat-price bids; finalists quote cost per resolved ticket by tier, with FCR, MTTR, and escalation-rate credits on identical volume scenarios	2–3
7 Transition & launch governance	Knowledge-transfer decay and untested incident response: PITON-Global stays embedded buyer-side through the first severity event, then hands over	1

Field-remaining counts represent a mid-market technical-support search (25–100 seats), PITON-Global engagements 2024–26.

The guarantees are the practice's constants. **Vendor neutrality:** no capacity sold, no placement economics — the live troubleshooting exercise, not the golf outing, decides who advances. **Right-sizing:** a 55-seat SaaS program inside a telco-scale generalist inherits telco-scale indifference; the framework surfaces technical specialists for whom the account is a flagship and the product becomes a second native language. Six to ten best-qualified providers, within 48–72 hours, free to the buyer, competing through a fully transparent, fair, comprehensive, and highly competitive RFP.

The queue that stopped growing: a 55-seat SaaS support program, reconstructed

Engagement TS-044 is a US B2B SaaS company — roughly \$140 million ARR, 5,400 accounts — whose support organization had become the growth plan's hostage: a 19-person onshore team drowning at 62% FCR, weekend queues answered on Monday, engineers losing a day a week to escalations without reproduction steps, and support-tagged churn reasons climbing in every QBR. The CFO's constraint was explicit: fix it without growing operating expense. Identifying details are anonymized under NDA.

SELECTION (WEEKS 0-8)

Requirements fixed a 55-seat, 24/7 scope — 34 T1, 15 T2, 6 KB and QA — on the client's Zendesk/Jira stack, with a live troubleshooting exercise written into Step 3: candidate T2 agents diagnosed three seeded issues in a sandbox tenant, scored blind. Mapping yielded 112 claimants; the exercise and capability screens left 19; the security audit left 11. Diligence produced a shortlist of 7; a three-round RFP concluded with a Manila technical specialist bidding per resolved ticket — \$6.20 T1 / \$14.10 T2 — with FCR, MTTR, and escalation-rate credits.

TRANSITION (WEEKS 9-24)

T1 certified in behind a three-week shadow queue; T2 followed after a joint sprint with the client's engineers that rebuilt the ticket taxonomy and escalation contract (reproduction steps mandatory, logs attached). The KB function shipped 240 new or rewritten articles in its first quarter, lifting deflection from 24% to 41%. The onshore team redeployed to a new premium-tier success function. PITON-Global chaired buyer-side governance through a severity-one incident in month 5 — SLAs held — and to certified steady state in month 6.

TABLE 3 – TS-044 FIRST-YEAR VALUE BUILD-UP (US\$)

VALUE COMPONENT	YEAR 1	BASIS
Labor & overhead arbitrage	\$2.30M	103k productive hrs × (\$35.40 – \$13.10)
Retention effect (support-tagged churn –31%)	\$0.45M	ARR preserved at measured attribution, low-end factor per Methodology
Deflection lift (KB 24%→41% of raw demand)	\$0.28M	Deflected tickets × T1 resolution cost avoided
Engineering time returned (escalations –38%, repro attached)	\$0.22M	Measured engineer-hours recovered × fully-loaded cost
Gross value created	\$3.25M	
Less: engagement & transition costs	(\$0.58M)	Transition, sandbox build, shadow queues, KB sprint, governance overlay; advisory fee \$0 (provider-paid model)
Net value ÷ cost = first-year ROI	5.6×	\$3.25M ÷ \$0.58M

Figures rounded; retention and recovery effects conservatively factored. Method in Methodology & notes.

Seventy-one percent arbitrage, twenty-nine percent operational value — the series ratio again, and again every operational line traces to a selection decision: the retention effect to the live troubleshooting exercise, the deflection lift to a KB function most bidders didn't staff, the engineering recovery to an escalation contract written in Step 1. The client's own summary, quoted with permission from the handover review: "We didn't buy cheaper support. We bought back our roadmap."

SECTION 07

The executive playbook: governance for the first 180 days

Technical-support programs live or die on the interfaces — between tiers, between vendor and engineering, between the KB and the release train. The governance below hardens each one, with the client in command throughout.

DAYS 0–30**Write the escalation contract first**

Define what T1 owns, what T2 owns, and what may reach engineering — with reproduction steps and logs as the price of admission. Contract per resolved ticket by tier with FCR/MTTR/escalation credits. Wire least-privilege access and session logging before the first ticket moves.

DAYS 30–90**Shadow, certify, and ship the KB sprint**

Run shadow queues until audited quality clears the gate; certify T2 on seeded issues, not tenure claims. Fund the KB rebuild as a project with article-level analytics — it is the highest-yield spend in the program. Join the vendor to the release train: notes before launch, macros in the same sprint.

DAYS 90–180**Drill the incident, then chase the root causes**

Run a severity-one drill before reality runs one for you. With SLAs stable for 60 days, point governance at demand itself: top-ticket-driver reviews feeding product fixes, deflection analytics, escalation root-cause work. Certify steady state; hand over on dated criteria; keep the quarterly retention review.

THE FIVE FAILURE MODES WE ARE HIRED TO PREVENT

Buying seats when the product needed resolutions · accepting a T2 bench nobody tested live · leaving the KB as an afterthought while paying T1 rates for its absence · letting escalations arrive bare and burning engineering goodwill · discovering the incident plan during the incident. Each is settled at selection or in the first sprint — never gracefully afterward.

Technical support is where product ambition meets operational truth, one ticket at a time. The Philippine market offers the strongest resolution economics available in 2026 — 68% off the cost per ticket, benchmarks onshore teams cannot staff to, and around-the-clock coverage priced flat. As every volume of this series concludes from its own direction: the market pays buyers who select on evidence. Forty-five providers can climb the tier stack; a live troubleshooting exercise, a verified escalation rate, and a fully transparent, fair, comprehensive, and highly competitive RFP will find the six to ten that belong on your shortlist.

METHODOLOGY & NOTES

How the numbers were built

Operating benchmarks (FCR, MTTR, escalation rate, CSAT, QA accuracy, deflection share, attrition) derive from PITON-Global technical-support engagements active 2025–26, measured on provider ticketing, QA, and knowledge platforms and validated in client governance. Baselines combine client pre-engagement data with onshore and commodity-offshore averages.

Cost figures are fully loaded — wages, benefits and statutory charges, facilities, technology and security stack, management/QA/knowledge overhead, attrition-driven recruiting and ramp — net of shrinkage. Cost per resolved ticket divides the loaded cost pool by quality-passed resolutions at each tier; rework and reopened tickets are charged back to the resolving tier. US baselines reflect Tier-2 metro support labor; Manila figures reflect 2025-26 engagement pricing.

The case study (TS-044) is anonymized under NDA; ARR, account counts, and stack details are generalized, figures rounded. Value factors are conservative: the retention effect uses cancellation-tagged attribution at the low end of the measured range; deflection counts only tickets matched to article views with no subsequent contact; engineering recovery counts logged escalation hours only. $ROI = \text{net first-year value} \div \text{all engagement-related client costs}$. PITON-Global's advisory fee is provider-paid and appears at \$0 in the client cost base; neutrality safeguards are documented at piton-global.com.

Market context draws on IBPAP reporting and PITON-Global's proprietary provider mapping (1,000+ operations; technical-support capabilities re-verified July 2026). Estimates and projections are PITON-Global analysis and should be cited as such.



Which technical-support BPOs belong on your shortlist?

Forty-five minutes with John will tell you. Free to buyers, strictly vendor-neutral — a bench-tested shortlist within two to three business days.

j.maczynski@piton-global.com

+1 402-598-8740

© 2026 PITON-GLOBAL, INC. · EXECUTIVE WHITE PAPER SERIES WP-04 · TECHNICAL SUPPORT OUTSOURCING · PHILIPPINES. This publication is provided for general information. Figures identified as estimates or models are PITON-Global analysis. Reproduction with attribution is permitted.

PITON-
GLOBAL

EXECUTIVE WHITE PAPER SERIES · WP-04 · TECHNICAL SUPPORT OUTSOURCING · PHILIPPINES ·
JULY 2026

PITON-
GLOBAL.COM