PITON Global

EXECUTIVE WHITE PAPER · WP-08

EMAIL SUPPORT OUTSOURCING · PHILIPPINES

The patient queue: the executive economics of email support outsourcing to the Philippines.

An analysis of how the asynchronous time buffer reprices service labor, the craft of SLA ladders and backlog control, judgment-led templating, and vendor-selection discipline for the channel every customer uses and no executive watches.

AUTHORS

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

AUGUST 2026
MANILA · ESTABLISHED 2001

Contents

ABOUT THIS SERIES — VOLUME 8

Email is the channel executives stopped watching around the time they started watching chat — yet it still carries a third or more of most support volume, and its asynchronous nature makes it the most schedulable, most improvable operation in CX. This volume prices the channel nobody benchmarks. The full series lives at piton-global.com.

THE PITON-GLOBAL LEADERSHIP TEAM

About the authors

John Maczynski

Chief Executive Officer

Before founding what became PITON-Global's client practice, John spent over three decades signing off Fortune-500 CX budgets — including the unglamorous email queues that never made a QBR slide until they broke. He learned to read a support operation by its backlog curve, and that instinct now shapes how the firm scopes, shortlists, and governs programs. Clients work with him directly; there is no junior bench between his name on this paper and their engagement.

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

Ralf Ellspermann

Chief Strategy Officer

Ralf ran Philippine delivery floors for twenty years, and the email queues were always his favorite instrument: the one channel where staffing, batching, and SLA design show up in the numbers within a week. The ladder-and-backlog disciplines in Section 3 are ones he implemented in production long before writing about them. Today he reads providers' queue dashboards the way he once built them — which is what a floor read is for.

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

PITON Global

ABOUT PITON-GLOBAL

Since 2001, PITON-Global has served one side of the outsourcing table: the buyer's. The firm maintains a working map of the Philippine provider market — more than 1,000 operations, capabilities re-verified continuously — and uses it to place companies with the six to ten partners genuinely fit for their program, then run the competition among them. No seats sold, no supply-side consulting, no placement commissions shaping the shortlist. Companies from Tripadvisor and Garmin to TSYS, HealthPartners, Yetipay, Workrise, and Norman Disney & Young have sourced through the practice.

SECTION 01

Executive summary

-62% Fully-loaded cost per productive hour vs. US onshore	**91%** Achievable occupancy on async queues, vs. ~74% on voice	**1.18** Touches per resolution on vetted floors, vs. 1.7–2.3 baseline	**6.1×** First-year ROI, reconstructed case (Section 6)

Email is the only service channel with a built-in time buffer — the customer expects an answer, not an answer *now*. That buffer is an economic gift most operations throw away. Used well, it lets a floor run fifteen points hotter on occupancy than any voice queue, batch similar work into faster and more accurate handling, schedule to demand instead of to peaks, and put the hardest cases in front of the most qualified people without a customer waiting on hold. Used badly, it becomes a place to hide a backlog until the backlog becomes the brand.

The Philippines handles more asynchronous support volume than any country on earth, and its written-English depth is the workforce fact that makes the channel exportable at quality. Five findings:

1. **The time buffer is worth 25–35% beyond the wage delta.** Higher occupancy, demand-shaped scheduling, and batching mean an async hour produces more resolutions than a synchronous one — value that stacks on top of the -62% labor arbitrage. Section 2 does the math.

2. **Touches per resolution is the metric that matters.** A queue answering every message within four hours but taking 2.3 exchanges to resolve costs more — in labor and in customer patience — than one answering in eight hours and resolving in 1.2. Most SLAs contract the wrong number. Vetted floors run 1.1–1.3.

3. **Templates should carry judgment, not replace it.** The best floors treat their template library as structured starting points — maintained weekly, personalized per contact, retired when root causes die. The worst paste them whole. Customers can tell, and so can a floor read: Step 3 screens for template governance the way other verticals screen for licensure.

4 **Backlog is a governance failure before it is a staffing failure.** Every email disaster follows the same arc: volume rises, first-response slips quietly, agents start double-touching apologetic customers, productivity falls exactly when demand peaks. The fix is a ladder of SLA tiers with automatic surge rules — designed at contract, not improvised at crisis.

5 **AI drafting makes the human layer more selective, not smaller first.** Draft-assist tools lift throughput 20–30% on vetted floors — and amplify template-pasting on weak ones. The differentiator has moved from typing speed to editorial judgment: what to keep, what to verify, when to abandon the draft. Selection now screens for editors.

The intended reader runs support, operations, or finance at a company where email and tickets carry real volume. Sections 2–4 price the channel; sections 5–7 show how the value is captured and kept.

SECTION 02

The asynchronous advantage: how a time buffer reprices service labor

A voice queue must be staffed for the worst fifteen minutes of the day; an email queue can be staffed for the average of the day — the buffer absorbs the peaks. Figure 1 shows what that difference does to paid time across the three service channels, measured on production floors.

FIG. 1 — WHERE A PAID HOUR GOES, BY CHANNEL (VETTED PHILIPPINE FLOORS, 2025–26)

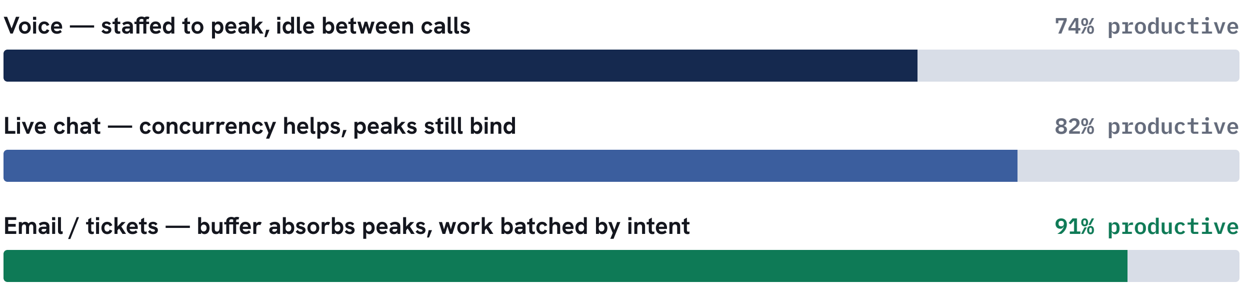

Productive share of paid hour = handling + wrap, net of idle and availability wait. PITON-Global engagement instrumentation, 2025–26; voice figure at healthy 80/20 service level.

Seventeen points of occupancy is only the visible part. Batching adds more: a queue sorted by intent lets an agent handle twelve consecutive refund requests with the policy loaded in working memory — measurably faster and more accurate than the same twelve scattered through a shift. Demand-shaped scheduling adds again: overnight Manila hours land answers in US customers' morning inboxes, converting time-zone offset from a coordination tax into a service feature. Stack the three effects and an asynchronous productive hour delivers 25–35% more resolutions than its synchronous equivalent — before the wage delta is counted.

The buffer has one price: it hides failure as gracefully as it absorbs load. A voice queue collapses loudly, in real time; an email queue collapses silently, one slipped SLA tier at a time. That asymmetry is why this channel — the easiest to run well — is so often run worst, and why Section 3 treats backlog control as the discipline that separates providers.

SECTION 03

Queue craft: SLA ladders, judgment-led templates, and backlog control

Figure 2 lays out the queue disciplines a floor read verifies — the machinery that keeps the patient queue patient instead of neglected.

FIG. 2 — THE THREE QUEUE DISCIPLINES, AS OPERATED ON VETTED FLOORS

THE SLA LADDER

Tiers by intent and stake — account-at-risk in 2 hours, order issues in 4, general in 8, low-stakes in 24. Surge rules trigger automatically at defined depth, not at someone's discretion.

TEMPLATE GOVERNANCE

A curated library with owners, weekly review, and retirement rules. QA scores personalization and accuracy, not template use. Paste-rate is tracked — and coached down.

BACKLOG CONTROL

Queue depth and age published hourly; oldest-first sweeps scheduled daily; a standing rule that no message ages past 2× its tier without escalation to the program lead.

↓ yields ↓

THE ZERO-AGED-BACKLOG OPERATION 1.18 TOUCHES / RESOLUTION 98%+ SLA ATTAINMENT

Every message answered within its tier, resolved in nearly one touch, by an agent whose queue was sorted to make that likely. The operation is boring to watch — which is the point, and what the buyer is paying for.

Operating parameters from PITON-Global email-support engagements, 2024–26.

None of this machinery appears in a proposal, because every proposal claims it. The floor read settles it in an hour: ask for the live queue-age dashboard (weak floors don't have one), pull last month's SLA attainment by tier (weak floors report a single blended number), and read ten outbound replies against their templates (paste-rate is obvious on sight). Providers who survive that hour tend to survive the contract. Steps 3 and 5 of the framework exist to run it before signature.

SECTION 04

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

TABLE 1 — FULLY-LOADED COST PER PRODUCTIVE HOUR BY ROLE BAND (US$, 2026)

Basis: PITON-Global email-support engagement pricing 2025–26 vs. US metro contact-center labor, fully loaded; follow-the-sun coverage at flat rates.

TABLE 2 — EMAIL SUPPORT OPERATING BENCHMARKS, VETTED VS. BASELINE (2025–26)			
METRIC	VETTED TOP TIER	BASELINE	EXECUTIVE READING
SLA attainment, by tier	**97–99%**	blended only	Reported per tier or not trusted
Touches per resolution	**1.1–1.3**	1.7–2.3	The real unit cost, and the real CSAT driver
Resolutions per productive hour	**6.5–8.0**	3.5–5.0	Batching + one-touch quality compounding
Aged backlog (>2× tier)	**~0%**	8–20%	The silent-collapse indicator
QA score on written quality	**92–96**	80–88	Personalization and accuracy, not paste-rate

Source: PITON-Global email-support engagement data 2025–26; baseline = typical pre-engagement operations, onshore or unvetted offshore.

The strategic read: resolutions per productive hour is where the two tables multiply. A vetted floor at $10.70 producing 7.2 resolutions per hour delivers a unit cost around $1.49 per resolution; the onshore baseline at $28.30 and 4.2 resolutions runs $6.74. That 4.5× unit-economics gap — not the 62% rate delta — is the number that belongs in the board deck, and it is entirely a function of which floor wins the RFP.

SECTION 05

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

Email diligence is unusually verifiable: the artifacts — queue dashboards, tier reports, outbound replies — either exist or they don't. The funnel below is representative of a 20–60-seat email and ticket search.

	FIG. 3 — THE 7-STEP FRAMEWORK AS AN EMAIL-SUPPORT FUNNEL		
	STEP	WHAT IT ELIMINATES IN AN EMAIL-SUPPORT SEARCH	FIELD REMAINING
1	**Requirements definition**	Volume-only staffing: rebuilds requirements from intent taxonomy, tier stakes, SLA ladder design, and the written-voice standard the brand actually wants on the page	1,000+
2	**Market mapping**	Voice-first floors that treat email as overflow filler, and giants for whom a 30-seat written program is nobody's priority	~55
3	**Capability screening**	Floors without queue machinery: per-tier SLA reporting, template governance with owners and review cadence, written-English assessment in hiring, draft-assist editorial controls	~15
4	**Compliance & security audit**	Weak data posture for a channel full of account details and attachments: SOC 2 on the delivery unit, DLP on outbound mail, retention rules, access governance on the ticket system	~10
5	**Operational diligence**	The artifact read: live queue-age dashboard, per-tier attainment exports, ten outbound replies read against templates, touches-per-resolution trend, backlog history through the last volume surge	6–9
6	**Competitive RFP & negotiation**	First-response-only SLAs: finalists bid with per-tier attainment floors, touches-per-resolution targets, aged-backlog zero rules, and written-QA gates — all carrying credits	2–3
7	**Transition & launch governance**	Voice drift and backlog surprises at cutover: reply-quality certification against a golden set, tier dashboards live from day one, PITON-Global embedded buyer-side to steady state	1

Field-remaining counts represent a 20–60-seat email and ticket search, PITON-Global engagements 2024–26.

The standing guarantees hold: **vendor neutrality** — PITON-Global operates no floors and takes no placement economics, so the shortlist reflects artifacts, not inventory; **right-sizing** — providers for whom your written channel is a flagship program. 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

From five-day backlog to eight-hour answers: a 30-seat email turnaround, reconstructed

Engagement ES-023 is a US online resale marketplace — about $180 million in revenue, two-sided disputes, seasonal spikes — whose onshore-plus-freelancer email operation had quietly collapsed: five-day first response at peak, 2.4 touches per resolution as apologies begat replies, and a support Reddit thread doing the brand's talking. Identifying details are anonymized under NDA.

SELECTION (WEEKS 0–7, INCL. ARTIFACT READS)

Requirements rebuilt the queue from a 22-intent taxonomy with a four-tier SLA ladder — disputes in 4 hours, payout issues in 6, general in 12, low-stakes in 24 — and a written-voice standard drafted with the brand team. Mapping produced 52 claimants; the machinery screen left 13; the DLP-and-access audit left 9. Artifact reads — live dashboards, tier exports, ten replies read cold — shortlisted 6. A Manila floor with marketplace experience won at $10.70 blended, bidding 1.25 touches and zero aged backlog with credits attached.

TRANSITION (MONTHS 3–6, BACKLOG BURNED DOWN FIRST)

A two-week surge team burned the 11,000-message backlog to zero before steady-state routing began — oldest first, one-touch goal, no apology boilerplate. Agents certified against a 60-reply golden set for voice and accuracy; draft-assist tooling switched on in month 4 with paste-rate coaching. Follow-the-sun scheduling put Manila's overnight on the US inbox morning. PITON-Global chaired weekly governance to steady state, then quarterly.

TABLE 3 — ES-023 FIRST-YEAR VALUE BUILD-UP (US$)		
VALUE COMPONENT	YEAR 1	BASIS
Labor arbitrage on the 30-seat program	$0.99M	56.3k productive hrs × ($28.30 – $10.70)
Touch collapse (2.4 → 1.2) — capacity released	$0.31M	Avoided seats at steady volume, priced at blended Manila rate
Dispute-cycle acceleration — seller payouts unblocked 3.1 days faster	$0.22M	Reduced churn among high-volume sellers vs. prior-year cohort, margin basis
Chargeback and goodwill-credit reduction from one-touch accuracy	$0.13M	Fewer duplicate credits and escalated disputes, audited quarterly
Gross value created	**$1.65M**	
Less: engagement & transition costs	($0.27M)	Backlog surge team, golden-set certification, dual-running, governance overlay; advisory fee $0 (provider-paid model)
Net value ÷ cost = first-year ROI	6.1×	$1.65M ÷ $0.27M

Figures rounded; retention and dispute effects cohort-verified and taken at margin. Method in Methodology & notes.

Sixty percent arbitrage, forty percent operational value — the touch collapse alone funded the transition. Each line traces to a diligence step: the tier ladder to Step 1, the one-touch bid to Step 6's contractible targets, the zero-backlog steady state to Step 5's artifact read. The COO's summary at the year-one review: "We used to measure how fast we said something. Now we measure how rarely we have to say it twice."

SECTION 07

The executive playbook: governance for the first 180 days

An email transition succeeds when the queue's honesty is engineered in from day one — client in command throughout.

DAYS 0–30

Contract the ladder, burn the backlog

Put per-tier attainment floors, touches-per-resolution targets, and an aged-backlog zero rule in the contract with credits attached. If a backlog exists, fund a surge burn-down before steady-state routing begins — never launch a new floor into an old queue. Stand up the tier dashboard as the program's single source of truth.

DAYS 30–90

Certify the written voice, watch the touches

Certify agents against a golden set of replies for voice, accuracy, and completeness — not speed. Track touches per resolution weekly by intent; every intent above 1.5 gets a root-cause review. Introduce draft-assist tooling only after the one-touch habit is established, with paste-rate coached from the first week.

DAYS 90–180

Kill root causes, certify steady state

Feed intent volumes into a monthly review with product — the cheapest roadmap input the company owns; every intent you delete is capacity returned. Verify the queue through one real volume surge before certifying steady state on two clean months. Hand over on documented criteria; keep the quarterly artifact read standing.

THE FIVE FAILURE MODES WE ARE HIRED TO PREVENT

Contracting first-response speed while touches multiply · launching a new floor into an old backlog · letting templates replace judgment instead of carrying it · reporting one blended SLA number that hides the aged tail · switching on AI drafting before the one-touch habit exists. All five are settled at selection and contract.

The email argument, in one sentence: the channel with the time buffer should be your calmest, cheapest, most accurate operation — and whether it is depends almost entirely on which floor runs it. A vetted Philippine queue delivers a 4.5× unit-cost advantage at 97%+ tier attainment — available to buyers who read artifacts instead of proposals and compete the finalists through a fully transparent, fair, comprehensive, and highly competitive RFP. The patient queue rewards a patient buyer. We compress the patience into three days.

METHODOLOGY & NOTES

How the numbers were built

Channel productivity figures (Fig. 1) come from PITON-Global engagement instrumentation, 2025–26: productive share = handling plus wrap over paid time, net of idle and availability wait, with voice measured at a healthy 80/20 service level for fair comparison. The 25–35% async premium stacks occupancy, batching, and demand-shaped scheduling effects measured separately.

Operating benchmarks (SLA attainment, touches per resolution, resolutions per hour, aged backlog, written QA) derive from provider exports validated in client governance. Cost figures are fully loaded — wages, benefits and statutory charges, facilities, technology, management and QA overhead, attrition-driven recruiting — net of shrinkage; US comparators reflect metro contact-center labor on the same basis.

The case study (ES-023) is anonymized under NDA; volumes and dates are generalized, figures rounded. Seller-churn and chargeback effects are cohort-verified against prior-year baselines and taken at margin. 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 are documented at piton-global.com.

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

How many touches does your queue really take?

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

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