INSURANCE CLAIMS PROCESSING · PHILIPPINES

The settlement-accuracy standard: the executive economics of insurance claims processing outsourcing to the Philippines.

An analysis of why claims touched is a volume vanity metric, how settlement accuracy and cycle time — never claims-handling throughput — decide the true cost of a claims operation once leakage, reopened claims, cycle-time drag, and complaints are counted, and the vendor-selection discipline that settles the claim right from FNOL to close.

AUTHORS

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

JULY 2026
MANILA · ESTABLISHED 2001

Contents

ABOUT THIS SERIES — VOLUME 30

Every claims proposal is priced per claim touched, because throughput is easy to count. This volume is about the only claims outcome a carrier's combined ratio actually depends on: the claim adjudicated to the right indemnity and settled inside cycle — not over-paid into leakage, not under-paid into a complaint or reopen, not left to age. The gap between touching a claim and settling it accurately is where leakage and cycle-time drag live. The full series lives at piton-global.com.

THE PITON-GLOBAL LEADERSHIP TEAM

About the authors

John Maczynski

Chief Executive Officer

John has bought claims capacity for three decades, and learned that the fast operation that touches every claim and settles them wrong is the most expensive one a carrier can hire — because every inaccurate settlement is leakage on the over-payments, a complaint or reopen on the under-payments, and cycle-time drag on everything in between. His standing question ignores the productivity report: of the claims you handled, how many were settled to the right indemnity, inside cycle? Applied at PITON-Global, that question is what clients get on every claims sourcing decision — personally, with no 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 claims floors, and he judges a claims team by its settlement accuracy and cycle time, not its claims-per-day: whether the indemnity was right and the claim closed on cycle, or leaked on the payout and reopened later. The claims contract in Section 3 codifies what he learned rebuilding teams that touched claims fast and settled them loosely. He now reads claims providers the way he once ran them — from the leakage audit and the reopen log 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, including each claims team's real settlement accuracy and cycle-time performance rather than its claims-per-day headline —

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

-59%	99%+	-58%	6.3×
Fully-loaded cost per accurately-settled claim vs. US onshore	Settlement-accuracy rate on vetted teams, up from 90–95%	Reduction in claims leakage on vetted teams	First-year ROI, reconstructed case (Section 6)

A claims engagement is bought on the price per claim and paid for on the settlements it gets wrong. The vendor reports five hundred thousand claims touched at a fraction of the onshore cost; the carrier finds the leakage on claims paid above indemnity, the complaints and reopens on claims paid below it, the cycle-time drag as files age in the queue, and the combined ratio drifting the wrong way while activity looked healthy. The operation hit its claims-touched number and the carrier inherited a leakage problem, a reopen backlog, and a cycle-time complaint. The gap between touching a claim and settling it accurately is where the real cost of cheap claims processing hides, because a claim settled wrong is not throughput delivered — it is money leaked or a policyholder aggrieved.

The Philippines is where insurance claims processing is delivered with the most settlement-accuracy discipline, because a deep, insurance-trained, detail-oriented talent pool can run the clean intake, accurate adjudication, and on-cycle settlement that accuracy demands. Five findings:

1. **Claims touched is the industry's most flattering number.** Any team can open a file and log activity. Vetted teams are measured on outcome: settlement accuracy, leakage, and cycle time. Weak providers quote a low per-claim price the carrier repays in leakage, reopens, and slow closes. Section 2 shows the gap.

2. **The accurate settlement, not the fast touch, is the product.** A claim that protects the combined ratio comes from clean intake, disciplined adjudication to the right indemnity, and settlement inside cycle — not from clearing a queue. Teams that settle accurately close once; teams that settle loosely leak on some claims and reopen others, at a cost far above the per-claim price.

3 **Leakage cuts both ways — over-payment and reopen.** Over-settle and the indemnity dollars leak straight out; under-settle and the claim reopens with a complaint, an appeal, and sometimes a regulator attached. Vetted teams hold settlement accuracy above 99% precisely because both tails cost — the money paid away and the claim that comes back.

4 **The leakage and the reopen, not the keystroke, are what cost.** A claim settled wrong costs the leaked indemnity, the rework, the reopen, and the cycle-time drag — many times the per-claim saving. Teams that adjudicate accurately cut leakage and close on cycle, and that protected indemnity — not the offshore rate alone — is where the case study's return is built.

5 **Contract on the accurately-settled claim, not the touched one.** The contracting failure of claims is paying per claim and hoping it settled right. Vetted deals price against settlement-accuracy floors, leakage ceilings, and cycle-time targets, making the accurate settlement the team's problem, which is where it belongs.

The intended reader is a claims VP, COO, or operations leader buying claims processing per claim and paying for it in leakage, reopens, and cycle-time drag. Sections 2–4 separate the teams that settle accurately from the teams that touch volume; Sections 5–7, the method and the proof.

SECTION 02

The volume mirage: claims touched versus accurately settled

Figure 1 shows one month of a carrier's claims two ways: as the productivity report a weak vendor bills, and as the settlement-accuracy reality the combined ratio experiences. The claims-touched number is identical in spirit to a chatbot's containment rate — a figure engineered to look like output while leakage, reopens, and cycle-time drag accumulate underneath.

FIG. 1 — ONE MONTH OF CLAIMS: SETTLEMENT-ACCURACY QUALITY VS. REPORTED PRODUCTIVITY

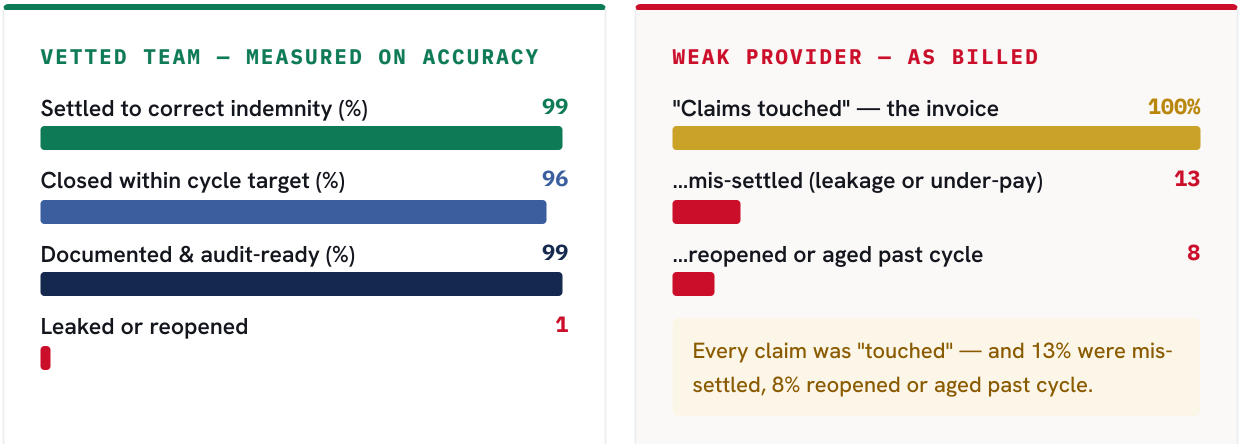

Left: PITON-Global engagement instrumentation, 2025–26, settlement audited. Right: composite of pre-engagement vendor reporting reviewed in diligence.

The economics follow the settlement accuracy, not the claim count. On the left, ninety-nine percent of claims are settled to the right indemnity and ninety-six percent closed inside cycle, so leakage stays near zero and nothing reopens; on the right, a cheap per-claim price mis-settles thirteen percent and reopens or ages eight, paying in leakage, reopens, and cycle-time drag at once. A low per-claim price that mis-settles is not a saving; it is leaked indemnity and combined-ratio damage, delayed and partly permanent. The settlement-accuracy audit — accuracy rate, leakage, cycle time — is the cheapest due-diligence instrument in this paper.

SECTION 03

The claims contract: intake it clean, adjudicate it accurately, settle it on cycle

Figure 2 lays out the three clauses of a working claims contract — the machinery a floor read verifies before any per-claim price is believed.

FIG. 2 — THE CLAIMS CONTRACT, AS OPERATED ON VETTED TEAMS

INTAKE IT CLEAN

Every FNOL is captured complete — coverage confirmed, facts and documentation gathered, the claim triaged and reserved correctly at the start — so adjudication runs on a clean file. The claim is set up right at intake, where a gap can still be closed before it distorts the settlement.

ADJUDICATE IT ACCURATELY

Coverage, liability, and indemnity are evaluated against the policy and the facts, the payout is calculated to the right number, and a questionable claim is investigated rather than paid to clear the queue. The settlement is made correct before the money moves, so it neither leaks nor reopens.

SETTLE IT ON CYCLE

The claim is closed inside its cycle-time target, the settlement documented and audit-ready, and aging kept inside limits — with a settlement report the carrier can read. An on-cycle settlement is a claim closed accurately before it ages, which is the only outcome the combined ratio actually feels.

↓ yields ↓

THE ACCURATE SETTLEMENT 99%+ SETTLEMENT ACCURACY · LEAKAGE HELD LOW · ON-CYCLE

Claims are settled to the right indemnity the first time, leakage and reopens fall, and every claim closes inside cycle and documented — so the carrier's combined ratio stops leaking and its claims experience improves. The operation gets cheaper per accurately-settled claim and tighter at the same time, which is the entire promise of outsourcing claims to people who measure settlement accuracy, not claims touched.

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

In diligence, the contract is tested from the settlement's side: pull a sample of closed claims and re-audit the indemnity against the policy and facts to measure leakage; read the reopen log for claims settled wrong; check cycle time against the aging report. Perhaps one team in ten survives it cleanly. 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 settlement accuracy, not claims touched.

SECTION 04

Cost and performance: cost-per-accurately-

settled-claim economics and benchmarks

TABLE 1 — FULLY-LOADED COST PER ACCURATELY-SETTLED CLAIM: THREE OPERATING MODELS (2026)

OPERATING MODEL	COST / SETTLED	VS. US BASELINE
US onshore in-house claims team	$14.20	baseline
Low-cost offshore claims mill (per-claim)	$9.65	−32%
Vetted Manila team — clean intake, accurate, on-cycle	**$5.82**	**−59%**

Basis: PITON-Global engagement pricing 2025–26; cost per accurately-settled claim = fully-loaded operating cost divided by claims settled accurately and on cycle (mis-settled and reopened excluded); the claims mill looks cheap per claim but leakage, reopens, and cycle drag lift its true cost per accurate settlement; US comparator fully-loaded in-house cost.

TABLE 2 — CLAIMS-PROCESSING BENCHMARKS, VETTED VS. BASELINE (2025–26)

METRIC	VETTED TOP TIER	BASELINE	EXECUTIVE READING
Settlement-accuracy rate	99%+	90–95%	The only claims number the ratio follows
Claims leakage vs. baseline	−58%	baseline	"Adjudicate accurately" — held or not
Cycle time (FNOL to close)	−44%	baseline	Closed on cycle, or aging in queue
Reopen rate	<3%	9–16%	Closed once, or handled twice
Claims-handling complaint rate	−52%	baseline	Fair & fast, or escalated

Source: PITON-Global claims engagement data 2025–26; baseline = typical pre-engagement operations billed on claims touched.

The strategic read: the claims mill's −32% per claim is a false economy — leakage, reopens, and cycle drag lift its true cost per accurate settlement far above the per-claim gap, and indemnity paid away to leakage never returns. The vetted Manila team's −59% is real because it is measured where the value is: the claim settled to the right indemnity, on cycle and documented. Buy on claims touched and you buy the right column of Figure 1; buy on cost-per-accurately-settled-claim and you buy settlement accuracy — the same discipline every volume in this series keeps arriving at from a different door.

SECTION 05

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

Claims diligence adds one hard requirement to every step: a per-claim price is real only when the settlement accuracy behind it survives a leakage re-audit against the policy and facts. The funnel below is representative of a 40–90-seat claims search.

FIG. 3 — THE 7-STEP FRAMEWORK AS A CLAIMS FUNNEL			
	STEP	WHAT IT ELIMINATES IN A CLAIMS SEARCH	FIELD REMAINING
1	**Requirements definition**	Throughput-first scoping: fixes the settlement-accuracy floor, the leakage ceiling, the cycle-time target, and the reopen standard the operation must meet by line of business	1,000+
2	**Market mapping**	Claims mills: operations that bill activity and treat leakage and reopens as the client's problem, and generalist floors with no claims-adjudication depth	~46
3	**Capability screening**	Teams without the machinery: audited settlement-accuracy history, a clean-intake process, an adjudication-QA and leakage-control workflow, and line-of-business coverage across your book	~13
4	**Compliance & security audit**	Weak data governance: PII protection, fair-claims-practice and market-conduct discipline, access controls on the claims systems, SOC 2 scope covering the claims stack	~8
5	**Operational diligence**	The leakage re-audit: re-audit closed claims against policy and facts for true accuracy and leakage; read the reopen and cycle logs; interview on how they adjudicate a complex or questionable claim	4–6
6	**Competitive RFP & negotiation**	Per-claim pricing: finalists bid against settlement-accuracy floors, leakage ceilings, and cycle-time targets — all carrying credits for claims that leak or reopen	2–3
7	**Transition & launch governance**	Cold cutover: the team calibrates against your policies, systems, and claim mix on a sample before it owns the queue, with the settlement-accuracy floor certified before steady state, PITON-Global embedded buyer-side	1

Field-remaining counts represent a 40–90-seat claims search, PITON-Global engagements 2024–26.

The standing guarantees hold: **vendor neutrality** — PITON-Global operates no claims floors, adjudicates no claims, and takes no placement economics, so the shortlist reflects audited settlement-accuracy performance, not inventory; **right-sizing** — teams for whom your operation 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-basing the operation on settlement accuracy: a 70-seat claims rebuild, reconstructed

Engagement CL-045 is a US property-and-casualty carrier — processing roughly 260,000 claims a year from FNOL through settlement across auto, property, and specialty lines — that had offshored claims processing to a per-claim vendor. The claim price was excellent; the settlements were not: settlement accuracy sat near 92%, leakage ran above benchmark, cycle time had drifted long, and reopens were climbing. The claims VP wanted the arbitrage without the leakage and the cycle-time drag. Identifying details are anonymized under NDA.

SELECTION (WEEKS 0–8, LEAKAGE RE-AUDITS)

Requirements fixed a 99% settlement-accuracy floor, a leakage ceiling, and a cycle-time target, with a mandatory clean-intake and adjudication-QA workflow. Mapping produced 46 claimants; the machinery screen — audited accuracy history, clean-intake process, leakage control, line-of-business coverage — left 13; the compliance-and-conduct audit left 8. Leakage re-audits of each finalist's closed claims against policy and facts shortlisted 5. A Manila team won priced against settlement accuracy and leakage, with a re-audit written into the reporting spec.

TRANSITION (MONTHS 2–7, CALIBRATED FIRST)

The team calibrated against the carrier's policies, claims systems, and claim mix on sample files and worked down the aged-claim backlog before owning production adjudication; the accuracy floor was certified in month 3. The leakage-and-QA loop fed adjuster coaching from day one. By month 7 settlement accuracy reached 99.1%, leakage fell 58%, cycle time dropped 44%, and reopens fell below 3%. PITON-Global chaired the settlement review board to steady state.

TABLE 3 — CL-045 FIRST-YEAR VALUE BUILD-UP (US$)		
VALUE COMPONENT	YEAR 1	BASIS
Labor arbitrage on the 70-seat claims operation	$2.14M	131.3k productive hrs × ($26.10 – $9.80)
Leakage recovered (accurate indemnity)	$1.02M	Indemnity dollars retained by lower leakage
Reopen & rework reduction	$0.30M	Reopened-claim and rework volume removed
Cycle-time & complaint value	$0.20M	Faster close and lower complaint/LAE load
Gross value created	**$3.66M**	
Less: engagement & transition costs	($0.58M)	Calibration, aged-claim work-down, claims-system integration; advisory fee $0 (provider-paid model)
Net value ÷ cost = first-year ROI	6.3×	$3.66M ÷ $0.58M

Figures rounded; accuracy, leakage, and cycle effects audit-verified. Method in Methodology & notes.

Fifty-nine percent arbitrage, forty-one percent from recovered leakage, reduced reopens, and faster cycle time — a typical split for the series, earned because the operation was re-based on settlement accuracy rather than claim volume. Every line traces to a diligence step: the accuracy floor to Step 1, the leakage-based pricing to Step 6, the adjudication-QA workflow to Step 3's contract. The claims VP's line at the year-one review: "The claim price barely moved. What moved is that we stopped leaking indemnity — and claims stopped coming back."

SECTION 07

The executive playbook: governance for the first 180 days

A claims transition succeeds when the team is measured and paid on settlement accuracy and leakage — client in command throughout.

DAYS 0–30

Define accuracy, contract on it

Set the settlement-accuracy floor, the leakage ceiling, and the cycle-time and reopen thresholds with claims leadership before launch, and document the intake and adjudication-QA workflow by line of business. Contract against settlement accuracy and leakage — never per claim alone — with floors and ceilings carrying credits. Stand up the leakage re-audit as the program's arbiter of truth.

DAYS 30–90

Calibrate, clear the aged claims

Calibrate the team against your policies, claims systems, and claim mix and work down the aged-claim backlog before it owns production adjudication; certify the accuracy floor and leakage control. Switch on the adjudication-QA loop from day one. Publish the settlement dashboard: settlement accuracy, leakage, cycle time, reopen rate, complaint rate.

DAYS 90–180

Own production, certify steady state

Hand production adjudication to the team only after it holds the accuracy floor and leakage ceiling. Review the settlement board monthly with the reopen log and leakage findings on the table. Certify steady state on two months at floor including a catastrophe or surge cycle; hand over on documented criteria; keep the quarterly leakage re-audit standing.

THE FIVE FAILURE MODES WE ARE HIRED TO PREVENT

Paying per claim while indemnity leaks · touching claims instead of adjudicating them accurately · settling to clear the queue instead of to the policy · letting claims age and reopen · owning production before settlement accuracy is proven. All five are settled at selection and contract.

The claims argument, in one sentence: a touched claim is activity and an accurately-settled claim is indemnity protected and a policyholder treated fairly, and a vetted Philippine team delivers –59% cost per accurately-settled claim with a 99%+ settlement-accuracy rate and leakage cut by well over half — for buyers who re-audit the settlement instead of counting claims touched, and compete the finalists through a fully transparent, fair, comprehensive, and highly competitive RFP. Anyone can touch a claim. We make sure you buy the team that settles it right.

METHODOLOGY & NOTES

How the numbers were built

Settlement-accuracy and leakage figures (Fig. 1, Table 2) derive from engagement instrumentation in PITON-Global claims engagements, 2025–26. Settlement accuracy and leakage are measured by re-auditing closed claims against the policy and facts; cycle time, reopen rate, and complaint rate are measured against the client's pre-engagement baseline. The weak-provider column is a composite of pre-engagement vendor reporting reviewed in diligence, shown for contrast.

Cost-per-accurately-settled-claim figures (Table 1) are fully loaded — wages, benefits and statutory charges, facilities, technology and claims-platform tooling, management and QA overhead including adjudication auditors and leakage specialists, attrition-driven recruiting — divided by claims settled accurately and on cycle, with mis-settled or reopened claims excluded from the denominator; US comparators reflect fully-loaded in-house cost on the same basis. An accurately-settled claim is one adjudicated to the right indemnity and closed on cycle.

The case study (CL-045) is anonymized under NDA; volumes and dates are generalized, figures rounded. Accuracy, leakage, and cycle effects are audit-verified; the leakage-recovered line reflects indemnity retained, and the cycle-time line is a conservative estimate. 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; settlement accuracy and leakage performance 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 provider's claim price survive a leakage re-audit?

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

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