

INSURTECH OPERATIONS • PHILIPPINES

The exception-handling standard: the executive economics of insurtech operations outsourcing to the Philippines.

An analysis of why items automated is a volume vanity metric, how exception-handling accuracy — never straight-through-processing rate — decides the true cost of a digital-insurance operation once mishandled exceptions, wrong overrides, escalation gaps, and customer fallout are counted, and the vendor-selection discipline that resolves the exception the automation kicked out.

AUTHORS

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

JUNE 2026
MANILA • ESTABLISHED 2001

Contents

–	About the authors	03
01	Executive summary	04
02	The volume mirage: items automated versus exceptions handled right	05
03	The exception contract: catch it, resolve it right, feed it back	06
04	Cost and performance: cost-per-correctly-handled-exception economics and benchmarks	07
05	Sourcing discipline: the 7-Step Vendor Vetting Framework, applied to insurtech ops	08
06	Case study: a 70-seat insurtech operation re-based on exception-handling accuracy	09
07	The executive playbook: governance for the first 180 days	10
–	Methodology & notes	11

ABOUT THIS SERIES – VOLUME 60

Every insurtech-ops proposal is priced per item processed, because straight-through volume is easy to count. This volume is about the only outcome the automation cannot deliver on its own and a cheap contract quietly mishandles: the exception — the claim the model could not price, the application the rules engine kicked out, the edge case the bot escalated — resolved correctly by the human layer behind the platform. The gap between automating an item and handling the exception right is where wrong overrides, escalation gaps, and customer fallout live. The full series lives at piton-global.com.

About the authors



John Maczynski
Chief Executive Officer

John has bought digital-operations capacity for three decades, and learned that the automation demo that dazzles on the straight-through path is the most expensive one an insurtech can buy when nobody prices the exception queue — because the exceptions the bot cannot handle are exactly the hard, high-value, high-risk items, and mishandling them is a wrong claim payment, a mis-issued policy, or a customer who quits the app. His standing question ignores the STP-rate slide: of the exceptions your automation kicked out, how many were resolved correctly by a human who understood the policy? Applied at PITON-Global, that question is what clients get on every insurtech 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 digital-insurance floors, and he judges an insurtech team by its exception-handling accuracy and its feedback loop into the platform, not its items-per-hour: whether the kicked-out case was resolved right and the recurring failure fed back to fix the model, or overridden fast and left to recur. The exception contract in Section 3 codifies what he learned rebuilding teams that cleared exception queues fast and got the hard cases wrong. He now reads insurtech providers the way he once ran them — from the exception-QA score and the override-error log backward.

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



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 insurtech team's real exception-handling accuracy and feedback-loop discipline rather than its items-per-hour 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

-58%

Fully-loaded cost per correctly-handled exception vs. US onshore

98%+

Exception-handling accuracy on vetted teams, up from 82-90%

-47%

Recurring exceptions removed via feedback into the platform

6.4×

First-year ROI, reconstructed case (Section 6)

An insurtech-operations engagement is bought on the straight-through-processing rate and the price per item, and paid for on the exceptions the automation kicks out and the partner gets wrong. The vendor and the platform demo a dazzling STP rate — eighty, ninety percent of applications, claims, and endorsements flowing through untouched; the insurtech finds that the remaining slice, the exception queue, is where all the hard money and risk live, and that the cheap support layer overrides them fast to clear the queue: the claim paid that should have been investigated, the policy issued on a misread disclosure, the edge case closed wrong and left to recur because nobody fed it back to fix the model. The automation hit its STP number and the insurtech inherited leakage, mis-issued risk, and a customer-experience problem in exactly the cases that mattered most. The gap between automating an item and handling the exception right is where the real cost of cheap insurtech ops hides, because a mishandled exception is not throughput delivered; it is the automation's hardest 15% resolved wrong.

The Philippines — digitally fluent, insurance-literate, and comfortable working inside modern insurtech stacks and APIs — is where digital-insurance exception handling is delivered with the most accuracy and feedback discipline, because that pool, properly trained, can run the judgment-grade resolution and model-feedback loop that a straight-through operation's exception layer demands.

Five findings:

- 1 Straight-through rate is the industry's most flattering number.** A high STP rate says the automation handled the easy items; it says nothing about the exception queue where the value is. Vetted teams are measured on outcome: exception-handling accuracy, override correctness, and feedback into the model. Weak providers ride the STP headline while mishandling the hard slice. Section 2 shows the gap.
- 2 The correctly-handled exception, not the automated item, is the product.** An exception resolved right comes from understanding the policy and the edge case, deciding correctly, and closing to the customer — not from overriding fast to clear a queue. Teams that handle exceptions accurately protect the book; teams that clear them loosely leak money and mis-issue risk on exactly the cases the automation flagged as hard.
- 3 The feedback loop is what separates a cost center from a flywheel.** Every exception carries a lesson: a recurring failure the model or rules engine could learn to handle. Teams that feed exceptions back shrink the queue over time; teams that just clear it leave the same exceptions recurring forever. Vetted teams cut recurring exceptions by nearly half by closing that loop.
- 4 Leakage and mis-issue, not the per-item rate, are what cost.** An exception resolved wrong costs the leaked claim dollar, the mis-issued policy, the complaint, and the recurring re-work — many times the per-item saving. Teams that handle exceptions accurately protect loss ratio and issue integrity, and that protected book — not the offshore rate alone — is where the case study's return is built.
- 5 Contract on the handled exception, not the automated item.** The contracting failure of insurtech ops is paying per item on the STP headline and hoping the exception queue is handled right. Vetted deals price against exception-handling-accuracy floors, override-correctness targets, and feedback-loop standards, making the hard slice the team's problem, which is where it belongs.

The intended reader is a COO, head of operations, or chief claims/underwriting officer at an insurtech or digital carrier buying an ops partner on the STP demo and paying for it in mishandled exceptions, leakage, and customer fallout. Sections 2-4 separate the teams that handle the exception right from the teams that ride the automation; Sections 5-7, the method and the proof.

SECTION 02

The volume mirage: items automated versus exceptions

handled right

Figure 1 follows one month of an insurtech's item flow through the automation waterfall: the straight-through slice the platform handles alone, and — below the line — the exception queue where the support partner's real work lives, shown as the weak vendor bills it versus the exception-handling reality the loss ratio and the app-store reviews experience.



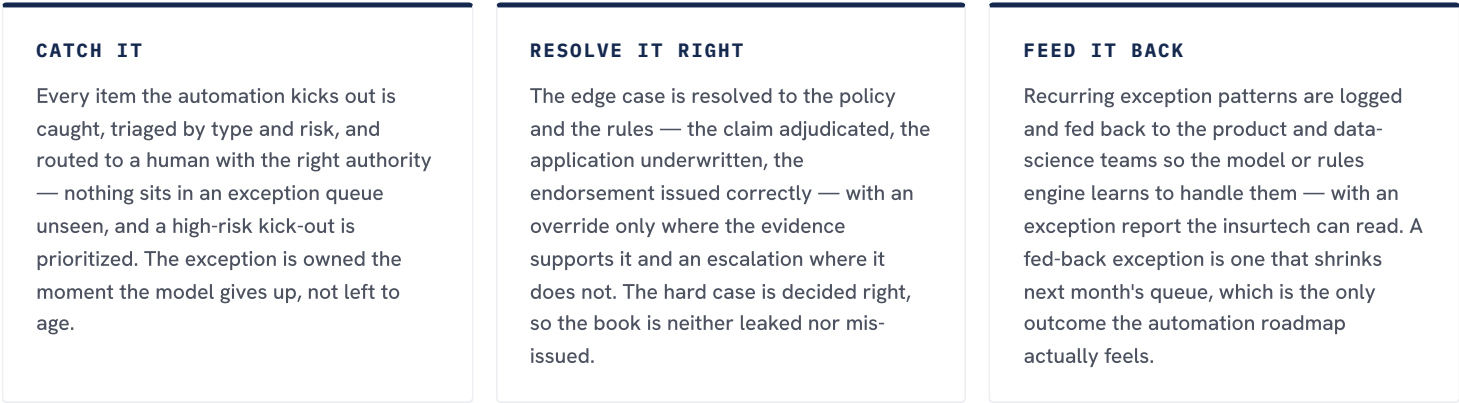
The economics live entirely below the waterfall line. The 84% straight-through slice costs almost nothing and is identical for both providers — the automation does it. All the money and risk are in the 16% exception queue, and there the vetted team resolves ninety-eight percent correctly and feeds most of it back, while the weak provider overrides fifteen percent wrong and lets the same exceptions recur forever. A high STP rate with a mishandled exception queue is not efficiency; it is the automation's hardest cases resolved wrong at scale. The exception-handling audit — handling accuracy, override correctness, feedback rate — is the cheapest due-diligence instrument in this paper.

SECTION 03

The exception contract: catch it, resolve it right, feed it back

Figure 2 lays out the three clauses of a working insurtech-ops contract — the machinery a floor read verifies before any STP-rate demo or per-item price is believed.

FIG. 2 – THE EXCEPTION CONTRACT, AS OPERATED ON VETTED TEAMS



↓ yields ↓

THE CORRECTLY-HANDLED EXCEPTION

98%+ HANDLING ACCURACY

CORRECT OVERRIDES

SHRINKING QUEUE

Exceptions are caught, resolved right, and fed back — so leakage and mis-issue fall, customers on the hard cases are treated well, and the queue shrinks as the model learns. The operation gets cheaper per handled exception and compounds with the platform at the same time, which is the entire promise of outsourcing insurtech ops to people who measure exception handling, not items automated.

Operating parameters from PITON-Global insurtech engagements, 2025-26. The human layer operates inside the client's platform and authority limits.

In diligence, the contract is tested from the exception queue's side: pull handled exceptions and re-decide a sample against policy to measure true handling accuracy and override correctness; read the escalation log for cases pushed through without authority; and audit the feedback loop — because a team that clears exceptions without feeding them back leaves the insurtech paying for the same failures forever. 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 handled exceptions, not items automated.

SECTION 04

Cost and performance: cost-per-correctly-handled-exception economics and benchmarks

TABLE 1 – FULLY-LOADED COST PER CORRECTLY-HANDLED EXCEPTION: THREE OPERATING MODELS (2026)		
OPERATING MODEL	COST / EXCEPTION	VS. US BASELINE
US onshore in-house exception team	\$18.60	baseline
Low-cost offshore queue-clear mill (per-item)	\$12.70	-32%
Vetted Manila team — accurate, fed-back	\$7.81	-58%

Basis: PITON-Global engagement pricing 2025-26; cost per correctly-handled exception = fully-loaded operating cost divided by exceptions resolved correctly and fed back (wrong overrides and un-fed-back excluded); the queue-clear mill looks cheap per item but leakage, mis-issue, and a queue that never shrinks lift its true cost per handled exception; US comparator fully-loaded in-house cost. Straight-through items are excluded from both numerator and denominator – the automation handles them.

TABLE 2 – INSURTECH-OPS BENCHMARKS, VETTED VS. BASELINE (2025–26)

METRIC	VETTED TOP TIER	BASELINE	EXECUTIVE READING
Exception-handling accuracy	98%+	82–90%	The only insurtech-ops number the book follows
Override correctness	98%+	80–88%	Justified by evidence, or leakage
Recurring exceptions fed back vs. baseline	–47%	baseline	"Feed it back" — held or not
Exception cycle time vs. baseline	–43%	baseline	Resolved fast, or the customer churns
Digital-journey CSAT on exceptions	89%+	70–80%	The app's reputation on the hard cases

Source: PITON-Global insurtech engagement data 2025–26; baseline = typical pre-engagement operations billed on items processed / STP headline.

The strategic read: the queue-clear mill's –32% per item is a false economy — leakage, mis-issue, and a queue that never shrinks lift its true cost per handled exception far above the per-item gap, and the hard slice mishandled at scale is a loss-ratio problem no rate recovers. The vetted Manila team's –58% is real because it is measured where the value is: the exception resolved right and fed back so the queue shrinks. Buy on the STP demo and you buy the right column of Figure 1; buy on cost-per-correctly-handled-exception and you buy exception-handling 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 insurtech ops

Insurtech-ops diligence adds one hard requirement to every step: a per-item price and an STP demo are real only when the exception-handling accuracy behind them survives a re-decision of the exception queue against policy. The funnel below is representative of a 40–100-seat insurtech exception-operations search.

FIG. 3 – THE 7-STEP FRAMEWORK AS AN INSURTECH-OPS FUNNEL		
STEP	WHAT IT ELIMINATES IN AN INSURTECH-OPS SEARCH	FIELD REMAINING
1 Requirements definition	STP-headline scoping: fixes the exception-handling-accuracy floor, the override-correctness target, the feedback-loop standard, and the exception-CSAT bar by item type	1,000+
2 Market mapping	Queue-clear mills: operations that bill exception volume and treat wrong overrides as the client's problem, and generalist floors with no insurance judgment or insurtech-stack fluency	~44
3 Capability screening	Teams without the machinery: audited exception-QA history, insurance domain depth, insurtech-platform/API fluency, and a working feedback loop into product and data science	~13
4 Compliance & security audit	Weak governance: customer-data and PII protection, insurance-regulatory alignment, access controls on the platform and admin tooling, SOC 2 scope covering the exception stack	~8
5 Operational diligence	The exception re-decision: re-decide a sample of handled exceptions against policy for true accuracy and override correctness; audit the feedback loop; interview on a genuinely ambiguous edge case	4–6
6 Competitive RFP & negotiation	Per-item pricing: finalists bid against exception-handling-accuracy floors, override-correctness targets, and feedback-loop standards — all carrying credits for wrong overrides	2–3
7 Transition & launch governance	Cold cutover: the team calibrates against your policies, platform, and exception taxonomy on a sample before it owns the live queue, with the accuracy floor and feedback loop certified before steady state, PITON-Global embedded buyer-side	1

Field-remaining counts represent a 40–100-seat insurtech exception-operations search, PITON-Global engagements 2024–26.

The standing guarantees hold: **vendor neutrality** — PITON-Global operates no ops floors, handles no exceptions, and takes no placement economics, so the shortlist reflects audited exception-handling 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 exception-handling accuracy: a 70-seat insurtech rebuild, reconstructed

Engagement IT-085 is a US digital-first insurtech — running roughly 6,400,000 policy, claims, and servicing items a year through an automated platform at an 84% straight-through rate — that had offshored its exception queue to a per-item queue-clear vendor. The STP demo was dazzling; the exception handling was not: exception-handling accuracy on that hard 16% sat near 86%, wrong overrides were leaking claims and mis-issuing policies, nothing was fed back so the same exceptions recurred every month, and digital-journey CSAT cratered precisely on the hard cases. The COO wanted the automation economics without the mishandled exceptions. Identifying details are anonymized under NDA.

SELECTION (WEEKS 0-8, EXCEPTION RE-DECISIONS)

Requirements fixed a 98% exception-handling-accuracy floor, a 98% override-correctness target, a feedback-loop standard, and an exception-CSAT bar. Mapping produced 44 claimants; the machinery screen — exception-QA history, insurance judgment, insurtech-stack fluency, working feedback loop — left 13; the compliance-and-security audit left 8. Blind re-decisions of each finalist's handled-exception samples against policy shortlisted 5. A Manila team won priced against exception-handling accuracy and feedback, with both written into the reporting spec.

TRANSITION (MONTHS 2-7, CALIBRATED FIRST)

The team calibrated against the insurtech's policies, platform, and exception taxonomy on a sample before owning the live queue; the accuracy floor was certified in month 3, and the feedback loop into product and data science ran from day one. The QA-and-feedback loop fed handler coaching and model fixes continuously. By month 7 exception-handling accuracy reached 98.3%, override correctness hit 98%, recurring exceptions fell 47% as fixes shipped, cycle time dropped 43%, and exception CSAT reached 89%. PITON-Global chaired the exception-accuracy review board to steady state.

TABLE 3 – IT-085 FIRST-YEAR VALUE BUILD-UP (US\$)

VALUE COMPONENT	YEAR 1	BASIS
Labor arbitrage on the 70-seat exception operation	\$1.82M	131.3k productive hrs × (\$23.40 – \$9.50)
Leakage & mis-issue recovered (accurate handling)	\$0.98M	Claim leakage and mis-issued risk removed
Queue shrink from feedback loop	\$0.44M	Recurring exceptions eliminated at the model
Retention & CSAT value on hard cases	\$0.22M	Churn avoided on exception journeys
Gross value created	\$3.46M	
Less: engagement & transition costs	(\$0.54M)	Calibration, feedback-loop build, platform integration; advisory fee \$0 (provider-paid model)
Net value ÷ cost = first-year ROI	6.4×	\$3.46M ÷ \$0.54M

Figures rounded; handling, override, and feedback effects audit-verified. Method in Methodology & notes.

Fifty-three percent arbitrage, forty-seven percent from recovered leakage, a shrinking queue, and retained customers on the hard cases — a typical split for the series, earned because the operation was re-based on exception-handling accuracy rather than items processed. Every line traces to a diligence step: the accuracy floor to Step 1, the accuracy-and-feedback pricing to Step 6, the feedback loop to Step 3's contract. The COO's line at the year-one review: "The STP rate didn't change. What changed is that the 16% the robots can't do stopped costing us money — and the queue is actually getting smaller."

SECTION 07

The executive playbook: governance for the first 180 days

An insurtech-ops transition succeeds when the team is measured and paid on exception-handling accuracy and feedback — client in command throughout.

DAYS 0–30

Define exception handling, contract on it

Set the exception-handling-accuracy floor, the override-correctness target, the feedback-loop standard, and the exception-CSAT bar with operations and product leadership before launch, and document the exception taxonomy and authority limits by item type. Contract against exception handling and feedback — never per item on the STP headline — with floors carrying credits. Stand up the exception re-decision audit and the feedback loop as the program's arbiters of truth.

DAYS 30–90

Calibrate, wire the feedback loop

Calibrate the team against your policies, platform, and exception taxonomy on a sample before it owns the live queue; certify the accuracy floor and wire the feedback loop into product and data science. Switch on the QA-and-feedback loop from day one. Publish the exception dashboard: handling accuracy, override correctness, recurring-exception rate, cycle time, exception CSAT.

DAYS 90–180

Own the queue, certify steady state

Hand the live exception queue to the team only after it holds the accuracy floor and the feedback loop is shipping model fixes. Review the board monthly with the recurring-exception and override-correctness trends on the table. Certify steady state on two months at floor including a product-release cycle; hand over on documented criteria; keep the exception re-decision audit standing.

THE FIVE FAILURE MODES WE ARE HIRED TO PREVENT

Buying on the STP demo while the exception queue is mishandled · clearing exceptions instead of resolving them right · overriding without evidence · never feeding exceptions back, so they recur forever · owning the queue before handling accuracy and the feedback loop are proven. All five are settled at selection and contract.

The insurtech-ops argument, in one sentence: an automated item is the easy path and a correctly-handled exception is the automation's hardest 15% resolved right and fed back so it shrinks, and a vetted Philippine team delivers –58% cost per handled exception with 98%+ exception-handling accuracy, recurring exceptions cut by nearly half, and CSAT held on the hard cases — for buyers who audit the exception queue instead of admiring the STP rate, and compete the finalists through a fully transparent, fair, comprehensive, and highly competitive RFP. Anyone can ride the automation. We make sure you buy the team that handles what the automation can't.

METHODOLOGY & NOTES

How the numbers were built

Exception-handling and feedback figures (Fig. 1, Table 2) derive from engagement instrumentation in PITON-Global insurtech engagements, 2025–26. Exception-handling accuracy and override correctness are measured by re-deciding a sample of handled exceptions against policy; recurring-exception, cycle-time, and CSAT effects are measured against the client's pre-engagement baseline. Straight-through items are excluded from the exception measures. The weak-provider column is a composite of pre-engagement vendor reporting reviewed in diligence, shown for contrast.

Cost-per-correctly-handled-exception figures (Table 1) are fully loaded — wages, benefits and statutory charges, facilities, technology and platform tooling, management and QA overhead including exception reviewers and feedback-loop analysts, and attrition-driven recruiting — divided by exceptions resolved correctly and fed back, with wrong overrides or un-fed-back exceptions excluded from the denominator; US comparators reflect fully-loaded in-house cost on the same basis. Straight-through items are excluded from both numerator and denominator. A correctly-handled exception is resolved to policy, overridden only on evidence, and fed back where it recurs.

The case study (IT-085) is anonymized under NDA; volumes and dates are generalized, figures rounded. Handling, override, and feedback effects are audit-verified; the leakage-recovered line reflects claim leakage and mis-issue removed, and the retention 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; exception-handling and feedback performance re-verified May 2026). The human layer operates inside the client's platform and authority limits and augments, never replaces, the client's automation. 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 STP demo survive an exception-queue audit?

j.maczynski@piton-global.com

+1 402-598-8740

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