PITON Global

EXECUTIVE WHITE PAPER · WP-55

FRAUD DETECTION & MITIGATION · PHILIPPINES

The loss-prevented standard: the executive economics of fraud detection and mitigation outsourcing to the Philippines.

An analysis of why alerts reviewed is an activity vanity metric, how fraud loss prevented and decision precision — never queue throughput — decide the true value of a fraud operation once missed fraud and false-positive friction are counted, and the vendor-selection discipline that catches loss without punishing good customers.

AUTHORS

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

AUGUST 2026
MANILA · ESTABLISHED 2001

Contents

ABOUT THIS SERIES — VOLUME 55

Every fraud-operations proposal is priced per alert reviewed, because queue volume is easy to count. This volume is about the only two outcomes that matter: the fraud caught before the money leaves, and the good customer who was never wrongly blocked. The gap between working an alert and preventing a loss is where missed fraud and false-positive friction 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 fraud-operations capacity for three decades, and learned that the cheap team that clears the alert queue fast is the most expensive one on the market — because it waves through the fraud that drains the account and blocks the good customers who then close theirs. His standing question ignores the queue report: of the alerts you worked, how much fraud did you actually prevent, and how many good customers did you wrongly stop? Applied at PITON-Global, that question is what clients get on every fraud 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 fraud and risk-operations floors, and he judges a fraud team by its decision precision and the loss it prevents, not its alerts-per-hour: whether the adjudication gate holds under a volume spike, and what the false-positive rate does to good customers. The fraud contract in Section 3 codifies what he learned rebuilding teams that cleared queues and missed the fraud that mattered. He now reads fraud providers the way he once ran them — from the missed-loss and false-positive logs 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 fraud team's real decision precision and loss-prevented performance rather than its alerts-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

-60%
Fully-loaded cost per correctly-adjudicated alert vs. US onshore

95%+
Fraud-catch recall on vetted teams, up from 70–82%

-64%
Reduction in false positives blocking good customers

6.4×
First-year ROI, reconstructed case (Section 6)

A fraud operation is bought on the price per alert and paid for on the fraud it misses and the good customers it blocks. The vendor reports a hundred thousand alerts reviewed at a fraction of the onshore cost; the business finds the account takeover that was waved through, the chargebacks that piled up, and the queue of loyal customers whose legitimate transactions were declined until they gave up. The operation hit its queue-throughput number and the business ate the fraud loss and the customer attrition at once. The gap between working an alert and preventing a loss is where the real cost of cheap fraud operations hides, because a fraud missed and a good customer blocked are both invisible on a throughput report.

The Philippines is where fraud operations are run with the most decision discipline, because a deep, analytical, pattern-trained talent pool can hold precision and recall together under the volume spikes fraud work demands. Five findings:

1 **Alerts reviewed is the industry's most flattering number.** Any team can clear a queue. Vetted teams are measured on outcomes: fraud loss prevented, catch recall, false positives held down. Weak providers quote a low per-alert price the business repays in missed fraud and blocked customers. Section 2 shows the gap.

2 **The adjudication gate is the product.** A correct fraud decision comes from working the signals, the history, and the network — not from clearing the alert to hit a queue target. Teams that gate every decision catch the fraud and clear the good customer; teams that gate nothing wave fraud through and block legitimate activity in equal measure.

3 **Precision and recall are one discipline, not a dial to trade.** Catching more fraud by blocking everything is not detection; clearing the queue by waving everything through is not either. Vetted teams hold recall above 95% while cutting false positives, because the business pays for both errors — the fraud missed and the customer lost.

4 **The missed loss and the blocked customer, not the keystroke, are what cost.** A fraud waved through is a chargeback and a write-off; a good customer blocked is friction, complaints, and attrition — many times the per-alert saving. Teams that hold precision drive both down, and that avoided cost — not the offshore rate alone — is where the case study's return is built.

5 **Contract on the loss prevented, not the alert cleared.** The contracting failure of fraud operations is paying per alert and hoping for good decisions. Vetted deals set recall floors, false-positive ceilings, and loss-prevented targets, making sound adjudication the team's problem, which is where it belongs.

The intended reader is a risk, fraud, or operations leader buying fraud operations per alert and paying for it in missed fraud and blocked customers. Sections 2–4 separate the teams that prevent loss from the teams that clear queues; Sections 5–7, the method and the proof.

SECTION 02

The activity mirage: alerts reviewed versus fraud loss prevented

Figure 1 shows one month of a fraud queue two ways: as the alerts-reviewed report a weak provider bills, and as the loss-prevented reality the business inherits. The alerts-reviewed number is identical in spirit to a chatbot's containment rate — a figure engineered to look like protection while fraud slips through and good customers are turned away.

FIG. 1 — ONE MONTH OF ALERTS: LOSS-PREVENTED QUALITY VS. REPORTED THROUGHPUT

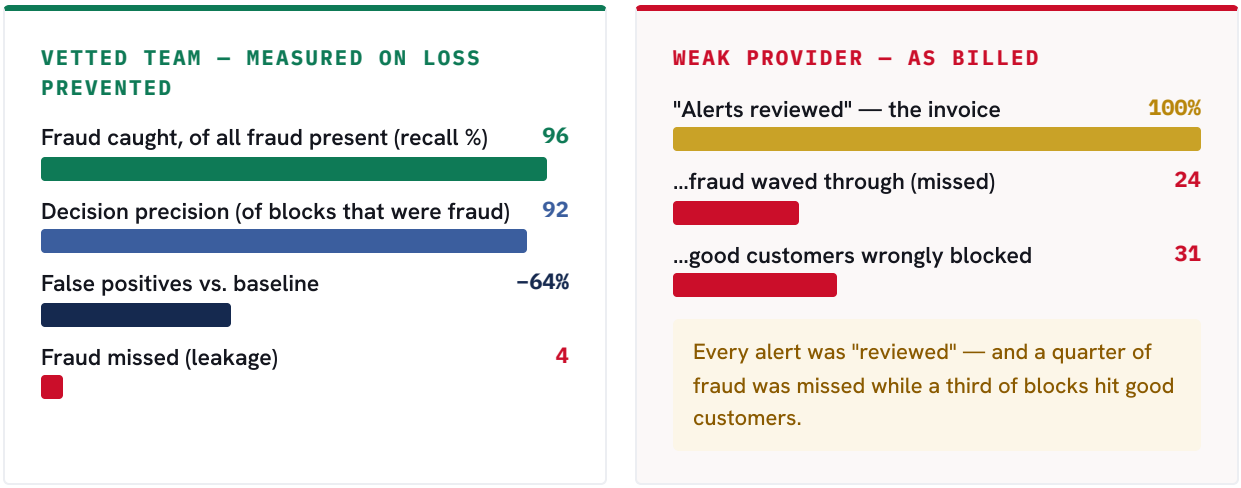

Left: PITON-Global engagement instrumentation, 2025–26, loss-prevented audited. Right: composite of pre-engagement vendor reporting reviewed in diligence.

The economics follow the loss prevented, not the queue cleared. On the left, ninety-six percent of fraud is caught, precision holds at ninety-two, and false positives fall by nearly two-thirds; on the right, a cheap per-alert price misses a quarter of the fraud and blocks a third of good customers, paying in write-offs and attrition at once. A low per-alert price that misses fraud and blocks customers is not a saving; it is a loss on both sides of the ledger. The loss-prevented audit — recall, precision, false positives, missed fraud — is the cheapest due-diligence instrument in this paper.

SECTION 03

The fraud contract: the adjudication gate, precision at recall, and the clean disposition

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

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

GATE THE DECISION

Every alert is worked against the signals, the account history, and the fraud-ring network before a block or clear is issued — with a reviewer on the ambiguous cases. The decision is made on evidence, not on a queue timer, so fraud is caught and the good customer is cleared.

HOLD PRECISION AT RECALL

The team is measured on catching fraud and sparing good customers at the same time, not trading one for the other. Analysts are calibrated to the client's risk appetite and coached on both error types, so recall stays high while false positives fall.

DISPOSITION IT CLEAN

Every case closes with a documented rationale, the action taken, and an outcome loop — confirmed fraud, confirmed good, or escalated — feeding the model and the analyst coaching. A clean disposition is one that improves the next decision, which is how the operation gets sharper over time.

↓ yields ↓

THE LOSS PREVENTED

95%+ RECALL | −64% FALSE POSITIVES | CLEAN DISPOSITIONS

Fraud is caught before the money leaves, good customers pass through unblocked, and every disposition makes the next decision sharper — so loss falls and customer experience improves at once. The operation gets cheaper per correct decision and more protective at the same time, which is the entire promise of outsourcing fraud work to people who measure loss prevented, not alerts cleared.

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

In diligence, the contract is tested against ground truth: replay a labeled set of alerts and measure recall, precision, and false positives; check the disposition rationale on a sample; verify the outcome loop actually feeds the model. 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 loss prevented, not alerts reviewed.

SECTION 04

Cost and performance: cost-per-adjudicated-

alert economics and benchmarks

TABLE 1 — COST PER CORRECTLY-ADJUDICATED ALERT: THREE OPERATING MODELS (US$, 2026)

OPERATING MODEL	COST / ADJUDICATED	VS. US BASELINE
US onshore in-house fraud team	$6.10	baseline
Low-cost offshore alert-clearing mill (per-alert)	$3.80	−38%
Vetted Manila team — adjudication-gated, precision-held	$2.44	−60%

Basis: PITON-Global engagement pricing 2025–26; correctly-adjudicated alert = one decided right (fraud caught or good customer cleared), net of rework; the alert mill looks cheap per alert but missed fraud and false positives cost far more than the gap; US comparator fully-loaded in-house cost.

TABLE 2 — FRAUD-OPERATIONS BENCHMARKS, VETTED VS. BASELINE (2025–26)

METRIC	VETTED TOP TIER	BASELINE	EXECUTIVE READING
Fraud-catch recall	95%+	70–82%	The fraud you actually stop
Decision precision	92%+	60–75%	Blocks that were actually fraud
False-positive rate vs. baseline	−64%	baseline	Good customers spared the friction
Recall held under 3× volume spike	yes	degrades	Protection under attack, or only when calm
Fraud loss prevented vs. baseline	+3–5×	baseline	The number the CFO actually feels

Source: PITON-Global fraud-operations engagement data 2025–26; baseline = typical pre-engagement operations billed on alerts reviewed.

The strategic read: the alert mill's −38% per alert is a false economy — missed fraud and blocked customers cost the business far more than the per-alert gap saves, and a single missed account-takeover ring can dwarf a year of queue savings. The vetted Manila team's −60% is real because it is measured where the value is: fraud caught, good customers spared, decisions that sharpen over time. Buy on alerts reviewed and you buy the right column of Figure 1; buy on cost-per-adjudicated-alert and you buy loss prevented — 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 fraud operations

Fraud diligence adds one hard requirement to every step: a per-alert price is real only when the decision quality behind it survives a labeled-set replay under load. The funnel below is representative of a 25–55-seat fraud-operations search.

FIG. 3 — THE 7-STEP FRAMEWORK AS A FRAUD-OPERATIONS FUNNEL

	STEP	WHAT IT ELIMINATES IN A FRAUD SEARCH	FIELD REMAINING
1	**Requirements definition**	Queue-first scoping: fixes the recall floor, the false-positive ceiling, the loss-prevented target, and the decision-quality standard the operation must meet	1,000+
2	**Market mapping**	Alert mills: operations that bill queue throughput and treat missed fraud as the client's problem, and generalist floors with no fraud-analyst or model depth	~43
3	**Capability screening**	Teams without the machinery: audited recall-and-precision history, an adjudication gate with reviewer, model-and-outcome loop, and a root-caused missed-fraud log	~11
4	**Compliance & security audit**	Weak data governance: customer-PII and transaction-data protection, model-access controls, regulatory competence (AML/KYC as scoped), SOC 2 scope covering the fraud stack	~8
5	**Operational diligence**	The labeled-set replay: run a ground-truth alert set and measure recall, precision, false positives under a volume spike; review disposition rationale; interview on how they work a ring	4–6
6	**Competitive RFP & negotiation**	Per-alert pricing: finalists bid against recall floors, false-positive ceilings, and loss-prevented targets — all carrying credits for fraud missed or good customers blocked	2–3
7	**Transition & launch governance**	Cold cutover: the team calibrates to the client's risk appetite on a shadow set before it owns live decisions, with recall and false-positive floors certified before steady state, PITON-Global embedded buyer-side	1

Field-remaining counts represent a 25–55-seat fraud-operations search, PITON-Global engagements 2024–26.

The standing guarantees hold: **vendor neutrality** — PITON-Global operates no fraud floors, adjudicates no alerts, and takes no placement economics, so the shortlist reflects audited recall-and-precision performance, not inventory; **right-sizing** — teams for whom your volume 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 loss prevented: a 35-seat fraud rebuild, reconstructed

Engagement FR-038 is a US digital-payments company — screening roughly 6 million transactions a month — that had offshored fraud operations to a per-alert vendor. The alert price was excellent; the outcomes were not: recall sat near 78%, false positives were high enough to generate a stream of customer complaints, and a missed account-takeover ring had produced a painful loss quarter. The head of risk wanted the arbitrage and fraud actually prevented. Identifying details are anonymized under NDA.

SELECTION (WEEKS 0–8, LABELED-SET REPLAYS)

Requirements fixed a 95% recall floor, a false-positive ceiling 60% below incumbent, and a loss-prevented target, with a mandatory adjudication gate. Mapping produced 42 claimants; the machinery screen — audited recall-and-precision, adjudication gate, outcome loop — left 11; the data-and-compliance audit left 8. Labeled-set replays under a 3× spike shortlisted 5. A Manila team won at a $2.10 blended per-adjudicated-alert rate, with a loss-prevented audit written into the reporting spec.

TRANSITION (MONTHS 2–6, SHADOW CALIBRATION FIRST)

The team calibrated to the client's risk appetite on a shadow set before owning live decisions; recall and false-positive floors were certified in month 2. The outcome loop fed the model and analyst coaching from day one, and a real volume spike was handled without a recall drop. By month 6 recall reached 96%, false positives fell 64%, precision held at 92%, and quarterly fraud loss fell sharply. PITON-Global chaired the loss-prevented review board to steady state.

TABLE 3 — FR-038 FIRST-YEAR VALUE BUILD-UP (US$)

VALUE COMPONENT	YEAR 1	BASIS
Labor arbitrage on the 35-seat fraud operation	**$1.02M**	65.6k productive hrs × ($26.40 − $10.90)
Fraud loss prevented (recall 78% → 96%)	**$1.34M**	Additional confirmed fraud caught before funds left, net of recoveries
False-positive friction & attrition avoided (−64%)	**$0.41M**	Retained customers and avoided complaint-handling, conservative
Chargeback & dispute-handling reduction	**$0.19M**	Lower chargeback volume and dispute-ops load vs. baseline
Gross value created	**$2.96M**	
Less: engagement & transition costs	**($0.46M)**	Shadow calibration, model-loop integration, analyst ramp; advisory fee $0 (provider-paid model)
Net value ÷ cost = first-year ROI	**6.4×**	$2.96M ÷ $0.46M

Figures rounded; recall, precision, and loss-prevented effects audit-verified. Method in Methodology & notes.

Thirty-four percent arbitrage, sixty-six percent from prevented fraud loss, avoided false-positive attrition, and lower chargebacks — a value-heavy split for the series, earned because the operation was re-based on loss prevented rather than alerts cleared. Every line traces to a diligence step: the recall floor to Step 1, the per-adjudicated-alert pricing to Step 6, the adjudication gate to Step 3's contract. The head of risk's line at the year-one review: "The alert price barely moved. What moved is that we stopped losing money and stopped losing customers."

SECTION 07

The executive playbook: governance for the first 180 days

A fraud transition succeeds when the team is measured and paid on loss prevented and decision precision — client in command throughout.

DAYS 0–30	**Define loss prevented, contract on it**	Set the recall floor, the false-positive ceiling, and the loss-prevented target with risk and finance before launch, and document the risk appetite. Contract on loss prevented and decision precision — never per alert alone — with floors and ceilings carrying credits. Stand up the loss-prevented audit as the program's arbiter of truth.
DAYS 30–90	**Calibrate on shadow, certify precision-at-recall**	Calibrate the team to the risk appetite on a shadow set before it owns live decisions; certify recall and false-positive floors under a spike. Switch on the outcome loop feeding model and coaching from day one. Publish the fraud dashboard: recall, precision, false positives, loss prevented, recall-under-load.
DAYS 90–180	**Own live decisions, certify steady state**	Hand live decisions to the team only after it holds the recall floor and false-positive ceiling through a spike. Review the loss-prevented board monthly with the missed-fraud log on the table. Certify steady state on two months at floor including an attack window; hand over on documented criteria; keep the quarterly labeled-set replay standing.

THE FIVE FAILURE MODES WE ARE HIRED TO PREVENT

Paying per alert while fraud slips through · clearing the queue without an adjudication gate · trading recall for a low false-positive rate or vice versa · a missed-fraud log no one root-causes · owning live decisions before precision-at-recall is proven under load. All five are settled at selection and contract.

The fraud-operations argument, in one sentence: a reviewed alert is activity and a prevented loss is protection, and a vetted Philippine team delivers –60% per adjudicated alert with 95%+ recall and false positives cut by nearly two-thirds — for buyers who audit the loss prevented instead of counting alerts reviewed, and compete the finalists through a fully transparent, fair, comprehensive, and highly competitive RFP. Anyone can clear an alert. We make sure you buy the team that stops the loss.

METHODOLOGY & NOTES

How the numbers were built

Recall, precision, and loss-prevented figures (Fig. 1, Table 2) derive from engagement instrumentation in PITON-Global fraud-operations engagements, 2025–26. Recall and precision are measured against labeled ground-truth sets and confirmed-outcome reconciliation; false positives 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-adjudicated-alert figures (Table 1) are fully loaded — wages, benefits and statutory charges, facilities, technology and fraud-platform tooling, management and QA overhead including adjudication reviewers, attrition-driven recruiting — net of rework; US comparators reflect fully-loaded in-house cost on the same basis. A correctly-adjudicated alert is one decided right — fraud caught or good customer cleared.

The case study (FR-038) is anonymized under NDA; volumes and dates are generalized, figures rounded. Recall, precision, and loss-prevented effects are audit-verified; the false-positive-attrition line is a conservative estimate, not a booked saving. 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; fraud recall-and-precision 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 alert price survive a labeled-set replay?

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

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