

FINANCIAL SERVICES OUTSOURCING • PHILIPPINES

The 2026 RAYVI Standard.

Why Risk-Adjusted Yield per Verified Identity has replaced cost-per-FTE as the governing metric for Philippine KYC/AML outsourcing — and how continuous PCI-DSS 4.0 architecture eliminates the 270-day compliance exposure window.

AUTHORS

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

JUNE 2026
MANILA • ESTABLISHED 2001

Contents

01 Executive summary

Five findings that reframe KYC/AML outsourcing as a yield decision, not a cost decision

02 The metric shift: from cost-per-FTE to RAYVI

Why the cheapest verification contract carries the most expensive loss curve

03 The RAYVI calculation model

The 40,000-application cost cohort behind the –81% per-clean-verification figure

04 The 270-day exposure window

PCI-DSS 4.0 Requirement 10.7 and continuous compliance architecture versus the periodic-audit failure model

05 The Tripartite Composite

The synthetic-identity typology defeating rules engines — and the Agentic + HITL forensic layer that intercepts it

06 Sourcing discipline: the 7-Step Framework, applied

How a field of 1,000+ Philippine providers collapses to roughly twenty-five with genuine regulated-delivery depth

07 Case study: a 60-seat digital-lender program, reconstructed

Before and after, with the first-year value decomposition

08 The executive playbook: governance for the first 180 days

Two clocks, five failure modes, and the handover criteria

– Methodology & notes

How the numbers were built, and how to cite them

ABOUT THIS SERIES

This series exists because outsourcing decisions are usually made with less evidence than any other spend of comparable size. Volume by volume, it publishes the material PITON-Global builds for clients privately — cost mechanics, verified benchmarks, provider-field analysis, and governance patterns — so that executives can pressure-test the business case before a single vendor conversation takes place.

THE PITON-GLOBAL LEADERSHIP TEAM

About the authors



John Maczynski
Chief Executive Officer

John's career divides into two acts with one constant. In the first, he spent thirty-plus years running outsourced operations for household-name US enterprises — banks and card processors among them — accountable for the results vendors delivered. In the second, he crossed the table to represent buyers full-time. The constant is the standard: every shortlist PITON-Global issues under his leadership is one he would sign a contract from himself, and he stays in the engagement until the program is live and proven.

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



Ralf Ellspermann
Chief Strategy Officer

Manila has been Ralf's operating base for more than twenty years — long enough to have built and run the kinds of regulated delivery floors this paper describes, for banks, card processors, insurers, and fintech ventures on three continents. He leads the research side of the practice: the provider dossiers, cost curves, and diligence protocols that every PITON-Global engagement draws on, and that give this series its data.

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

ABOUT PITON-GLOBAL

A quarter century in the Philippine market has given PITON-Global something no directory or marketplace can replicate: firsthand knowledge of how more than a thousand providers actually perform — who keeps commitments after the sales team leaves the room, whose certifications survive a floor walk, whose attrition numbers are real. The firm, decorated with numerous industry awards over the years, puts that knowledge exclusively at the service of buyers: scoping requirements, gating the field, running competitive RFPs, and staying on the client's side of the table through launch. Tripadvisor, TSYS, HealthPartners, Yetipay, Workrise, and Norman Disney & Young are among the companies that have built Philippine teams this way.

SECTION 01

Executive summary

-81%

Cost per clean verification, RAYVI lens (Section 3 cohort)

-74%

Fully-loaded cost vs. US onshore at standard role mix

270 d

Compliance exposure window a quarterly-audit model generates (Section 4)

60 s

Anomaly-to-alert SLA under continuous Req. 10.7 monitoring

For two decades, financial-services outsourcing has been bought on one number: cost per FTE. The number is easy to compare, easy to negotiate, and — for regulated identity work — systematically wrong. A verification desk is not a cost center that occasionally touches risk; it is a risk instrument that happens to carry a payroll. Buying it on payroll alone is how institutions end up with the paradox this paper opens with: the lenders holding the cheapest KYC contracts carry the highest fraud losses, the slowest originations, and the largest remediation spend in their cohort.

The correction is a metric, not a slogan. **Risk-Adjusted Yield per Verified Identity (RAYVI)** prices a verification operation by what it actually produces — clean, activated, durable identities — net of what it silently costs: fraud losses admitted, false positives that destroy good applicants, remediation and re-work, and the compliance exposure between audits. Measured this way, the Philippine top tier does not merely match the onshore baseline at lower cost; it beats it on the loss curve while costing 74% less to run. Five findings frame this volume:

- 1 Cost-per-FTE and RAYVI rank the same vendors in opposite order.** In the Section 3 cohort, the cheapest bidder per seat is the most expensive per clean verification once fraud admittance and false-positive abandonment are counted. The metric you contract on determines the operation you get.
- 2 Certification is a photograph; compliance is a film.** PCI-DSS 4.0 Requirement 10.7 made continuous control monitoring mandatory. The quarterly-audit model still run by most providers leaves an aggregate ~270-day window per year in which control failures surface at the next audit instead of the next minute.
- 3 Synthetic identity is now an assembly problem, not a forgery problem.** The dominant 2026 typology — the Tripartite Composite of Section 5 — passes every single-document check because every document is real. It fails only against cross-source coherence analysis: an agentic screen at machine speed, adjudicated by forensic analysts.
- 4 AI re-prices analysts; it does not replace them.** The Agentic + HITL architecture clears routine volume at machine speed and routes composite-identity adjudication, jurisdictional coherence, and regulatory sign-off to Step 04-vetted humans. The agent proposes; a qualified human disposes, and signs.
- 5 RAYVI-grade providers are a narrow slice of the market.** Of 1,000+ Philippine BPO operations, PITON-Global's mapping confirms genuine regulated-delivery depth — live KYC/AML programs, continuous-monitoring floors, forensic escalation desks — at roughly twenty-five. Finding them is a research problem before it is a procurement problem.

The intended reader is a COO, CRO, head of financial crime, or head of operations at a bank, lender, payments platform, or insurer — at the point where verification volume is scaling faster than the team that clears it. Sections 2-5 present the evidence; Sections 6-8, the method and the proof.

SECTION 02

The metric shift: from cost-per-FTE to RAYVI

Cost-per-FTE survives because it is legible to procurement. It fails because it prices the input while ignoring the output. A KYC seat produces four things simultaneously: approved-and-clean identities (revenue), approved-and-fraudulent identities (losses, discovered later), rejected-good applicants (abandonment, invisible on any operations dashboard), and casework that must be re-done (remediation). Only the first is yield. A contract that compensates seats is indifferent among all four.

FIG. 1 – THE RAYVI DEFINITION

$$\text{RAYVI} = (\text{Net value of clean, activated identities} - \text{fraud losses admitted} - \text{false-positive abandonment} - \text{remediation \& exposure cost}) \div \text{verification operating spend}$$

Each deduction is measurable on client-side systems: charge-off attribution for admitted fraud, funnel telemetry for abandonment, QA re-work logs for remediation, control-failure findings for exposure. Method in Methodology & notes.

The paradox this produces is best seen side by side. Figure 2 compares two anonymized digital lenders from PITON-Global's engagement base — one contracted on cost-per-FTE with a commodity offshore provider, one on RAYVI-linked terms with a vetted Manila specialist. The FTE buyer pays 22% less per seat and 2.4× more per clean verification.

FIG. 2 – TWO LENDERS, TWO CONTRACTS: THE SAME QUARTER, PRICED BOTH WAYS

METRIC	FTE-CONTRACTED	RAYVI-CONTRACTED
Cost per seat-hour	\$8.90	\$11.40
Fraud admittance rate	0.86%	0.21%
False-positive rejection of good applicants	11.2%	3.8%
Median review turnaround	16.8 hrs	4.1 hrs
Cost per clean, activated verification	\$31.70	\$13.10

Anonymized engagement data, normalized to the same quarter and application mix. Fraud admittance measured by 9-month charge-off attribution.

The mechanism is incentives, not competence. A provider paid per seat optimizes occupancy: clear the queue, keep handle time low, escalate nothing. A provider whose credits ride on RAYVI components — admittance, abandonment, turnaround, audited accuracy — optimizes exactly what the buyer's P&L feels. Tie compensation to RAYVI and you spend less *by not absorbing the losses the cheap model silently generates*. That contractual inversion, not any single technology, is the 2026 standard this paper is named for.

The RAYVI calculation model: the 40,000-application cohort

To move RAYVI from definition to decision, PITON-Global reconstructed a full cost cohort: 40,000 consumer-credit applications processed under three delivery models — US onshore in-house, commodity offshore (FTE-contracted), and vetted Manila top tier (RAYVI-contracted) — with every deduction priced from client-side data. Table 1 shows the result; the -81% headline is the last row.

TABLE 1 — 40,000-APPLICATION COHORT: FULL VERIFICATION ECONOMICS BY DELIVERY MODEL (US\$, 2026)			
COMPONENT	US ONSHORE	COMMODITY OFFSHORE	MANILA VETTED
Verification operating spend	\$1.31M	\$0.42M	\$0.49M
Fraud losses admitted (9-mo attribution)	\$0.74M	\$1.62M	\$0.38M
False-positive abandonment (lost contribution)	\$0.51M	\$0.93M	\$0.31M
Remediation, re-work & exposure cost	\$0.22M	\$0.41M	\$0.09M
Total cost of verification	\$2.78M	\$3.38M	\$1.27M
Clean, activated identities produced	30,900	28,400	33,700
Cost per clean verification	\$90.00	\$119.00	\$37.70

Basis: composite of PITON-Global financial-services engagements 2025-26, normalized to a 40,000-application consumer-credit cohort. Manila vetted vs. commodity offshore: -68%; vs. US onshore in the RAYVI lens with loss deductions at cohort rates: -81% per clean verification (peak-quarter basis). Method in Methodology & notes.

Three readings deserve emphasis. First, the commodity model is the most expensive of the three despite the lowest operating spend — its losses are 3.9× its payroll. Second, the vetted model produces *more* clean identities from the same 40,000 applications: lower false-positive rejection is yield, not just courtesy. Third, the onshore column is not incompetent — it is simply structurally unable to staff forensic review density at US labor rates, so it buys accuracy with turnaround time, and pays for it in abandonment.


In the plain FTE lens — for buyers who must still budget that way — the same cohort prices at **70–78% below an onshore build** depending on role mix, 74% at the standard blend of verification analysts, AML casework, and QA. The RAYVI lens simply adds what the FTE lens forgets to count.

SECTION 04

The 270-day exposure window: PCI-DSS 4.0 and continuous compliance architecture

PCI-DSS 4.0 — fully mandatory since March 2025 — changed the grammar of compliance. Requirement 10.7 obliges entities to detect and respond to failures of critical security controls *as they happen*, not at the next assessment. Most offshore providers still operate the model the old standard tolerated: certify annually, self-assess quarterly, and treat the space between as presumptively compliant. The arithmetic of that model is unforgiving. Four quarterly checkpoints leave roughly 270 days per year in which a failed control — a logging gap, a privilege drift, an unpatched jump host — exists, is invisible, and is accruing liability that surfaces only at the next audit. That is the exposure window; the certificate on the wall says nothing about it.

FIG. 3 – CONTROL-FAILURE VISIBILITY: PERIODIC AUDIT VS. CONTINUOUS MONITORING



What continuous architecture looks like on a vetted Manila floor is concrete, not aspirational: **biometric MFA** on every session; **non-persistent VDI** so no data survives a session; **dedicated client environments** rather than shared floors; **100% interaction logging** feeding the telemetry layer; and control-state dashboards the client's own risk team can read. The certification tells you what one good day looked like; the telemetry covers day 91.

THE DILIGENCE QUESTION THAT SORTS THE FIELD

"Show me the alert, the ticket, and the closure record for your last critical-control failure." A continuous-architecture provider produces all three in minutes, because failures are routine, caught, and evidenced. A periodic-model provider offers the certificate instead — which is precisely the problem. PITON-Global asks this question on every floor walk in Step 4 of the vetting framework; roughly four in five providers cannot answer it.

SECTION 05

The Tripartite Composite: synthetic identity and the Agentic + HITL forensic layer

Synthetic identity fraud stopped being a forgery problem around 2023. The dominant 2026 typology — PITON-Global's diligence protocols call it the **Tripartite Composite** — assembles one fictitious person from three authentic parts: a **real national identifier** (typically a dormant or minor's SSN with no credit footprint), a **real address history** (rented mailbox chains that survive utility checks), and a **real device-and-behavior trail** (aged social accounts, warmed phone numbers, human-plausible session behavior). Every component passes its own check because every component is genuine. The fraud exists only in the *combination* — and single-document verification, by construction, never examines the combination.

FIG. 4 – WHY EACH LAYER PASSES, AND WHERE THE COMPOSITE FAILS

COMPONENT	WHY IT PASSES ALONE	WHERE COHERENCE ANALYSIS CATCHES IT
Real national identifier	Valid issuance, no fraud flags, thin file reads as "young credit"	Identifier age vs. digital-footprint age diverge by a decade or more
Real address history	Survives utility and postal verification; matches documents	Address chain correlates with known mailbox clusters across other applications
Real device & behavior trail	Aged accounts, human session patterns, no emulator signals	Behavioral fingerprint recurs across "different" identities in the portfolio

Composite of PITON-Global financial-crime diligence protocols and client-side fraud attribution, 2024–26.

Catching the combination requires two capabilities no rules engine has: cross-source coherence analysis at portfolio scale, and judgment about what an incoherence *means*. That is the division of labor in the **Agentic + HITL architecture** the vetted tier now runs. The agentic layer screens every application against the full coherence surface — identifier age curves, address-cluster graphs, behavioral fingerprints — at machine speed, clearing the routine 90%+ and scoring the remainder. The human layer is not a queue of generalists:

it is Step 04-vetted forensic analysts who adjudicate composite-identity cases, test jurisdictional coherence, and carry regulatory sign-off. The agent proposes; a qualified human disposes, and signs — because a false positive here is a destroyed good customer, and a false negative is a charge-off with a nine-month fuse.

90%+

of volume cleared agentially, at machine speed, with full evidence trail

100%

of composite-suspect cases adjudicated and signed by a forensic analyst

0.21%

fraud admittance at the RAYVI-contracted lender of Section 2 — 4x better than its FTE-contracted peer

SECTION 06

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

Financial services compresses the framework's funnel harder than any vertical except healthcare: the field collapses at capability screening (live KYC/AML programs with forensic depth are rare) and again at the compliance audit (Section 4's continuous-architecture test is pass/fail). The stages below show a representative mid-market lender search.

1	Requirements definition Work types, volumes, data classes, regulatory regimes, RAYVI targets — fixed before any vendor contact	1,000+ →
2	Market mapping Eliminates providers without financial-services delivery units in the required hubs and scale band	~65 →
3	Capability screening Live KYC/AML programs, forensic escalation desks, coherence-analysis tooling — evidenced, not asserted	~25 →
4	Compliance & security audit The continuous-architecture test applied on the floor: Req. 10.7 telemetry, biometric MFA, non-persistent VDI, alert-ticket-closure evidence	~11 →
5	Operational diligence QA calibration, attrition cohorts, analyst-tenure economics, reference calls with comparable institutions	6–10 →
6	Competitive RFP & negotiation	2–3 →

Finalists compete on transparent per-productive-hour terms; credits tied to RAYVI components — admittance, abandonment, TAT, audited accuracy

7

Transition & launch governance

PITON-Global stays embedded buyer-side from signature to certified steady state — then hands over

1

As throughout the practice, the framework is **vendor-neutral by construction** — no inventory, no placement quotas, no way to buy a shortlist position — and **right-sizing-led**: for a 20-to-80-seat financial-services program it deliberately surfaces specialist mid-sized providers for whom the account is a flagship, carrying the certifications of the majors with the governance access of a boutique. Output: 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.

SECTION 07 • ANONYMIZED CASE STUDY

Before and after: a 60-seat digital-lender program, reconstructed

Engagement FS-041 is a US digital lender originating roughly \$900 million annually across personal and small-business credit. Two years of growth had pushed KYC review turnaround past 17 hours, fraud charge-offs to 0.9% of originations, and the in-house verification team of 41 into permanent backlog. The board approved offshore sourcing with two conditions: continuous compliance posture acceptable to the partner bank, and a measured RAYVI improvement — not just cost relief — within two quarters. Identifying details are anonymized under NDA.

SELECTION (WEEKS 0–7)

Requirements fixed a 60-seat scope — 26 KYC/KYB verification, 18 AML casework and screening, 16 fraud review — with continuous Req. 10.7 monitoring, dedicated environment, and the client's decisioning stack. Mapping yielded 63 claimants; capability screening on live coherence-analysis operations cut the field to 21; the floor-walk compliance audit left 10. Diligence produced a shortlist of 7; a three-round RFP concluded with a Manila financial-crime specialist at \$13.60 per productive hour, credits tied to admittance, abandonment, TAT, and audited accuracy.

TRANSITION (WEEKS 8–20)

Three certification-gated waves. Fraud review transitioned first behind a two-week parallel run scored on the client's own attribution data; KYC followed with dual-review on every decision for the first month; AML casework moved last with the partner bank observing the calibration sessions. The provider's coherence-analysis layer replaced the client's rules engine for composite screening in week 14. PITON-Global chaired weekly buyer-side governance through month 6; steady state was certified in month 7.

FIG. 5 – FS-041 AT MONTH 9: BEFORE → AFTER

KYC REVIEW (MEDIAN) 17.2 hrs → 4.1 hrs	FRAUD ADMITTANCE 0.90% → 0.24%	FALSE-POSITIVE REJECTION 10.4% → 3.9%
COST PER CLEAN VERIFICATION \$84.10 → \$35.90	AUDITED CASE ACCURACY 95.1% → 98.8%	COMPLIANCE EXPOSURE WINDOW ~270 days → 0

Client-side attribution and funnel data, month 9 versus trailing pre-engagement quarter. Backlog cleared in week 15.

FIRST-YEAR VALUE COMPONENT	YEAR 1	BASIS
Labor & overhead arbitrage	\$2.41M	113k productive hrs × (\$34.90 – \$13.60)
Fraud losses avoided (admittance 0.90%→0.24%)	\$1.94M	Charge-off attribution at observed average loss severity, conservatively discounted
Recovered originations (false positives 10.4%→3.9%)	\$0.66M	Incremental approved-good applicants × first-year contribution, discounted
Remediation & backlog-triage cost avoided	\$0.28M	Accuracy lift + backlog clearance vs. trailing run-rate
Gross value created	\$5.29M	
Less: engagement & transition costs	(\$0.71M)	Transition, parallel runs, security validation, governance overlay; advisory fee \$0 (provider-paid model)
Net value ÷ cost = first-year ROI	6.5×	\$4.58M ÷ \$0.71M

Figures rounded; loss-avoidance and conversion effects conservatively factored. Method in Methodology & notes.

Note the composition: arbitrage is under half of gross value. The majority came from the RAYVI components — fraud not admitted, good applicants not destroyed, backlog risk retired. A rate-card process would have captured the arbitrage and left roughly \$2.9 million of loss-side value on the table, while contracting the incentives that generate those losses in the first place.

The executive playbook: governance for the first 180 days

Regulated verification programs live under two clocks — the regulatory clock (SAR timelines, partner-bank reviews, exam cycles) and the origination clock. The governance below keeps both, with the client in command throughout.

DAYS 0–30

Wire the telemetry before the volume

Validate the continuous-monitoring floor in person: alert, ticket, closure record. Provision least-privilege access with immutable logging before data flows. Contract per productive hour with credits on RAYVI components — and lock the named roster that won the RFP. Give the partner bank read access to the control dashboard before go-live, not a PDF after.

DAYS 30–90

Certify, parallel-run, then release volume

Run fraud and KYC pods behind parallel runs scored on your own attribution data; release each wave on certification, never on dates. Chair weekly governance with the RAYVI dashboard first — admittance, abandonment, TAT, accuracy — and the deadline dashboard second. Calibrate the agentic layer's escalation thresholds against your loss tolerance, in writing.

DAYS 90–180

Optimize the loss curve, not just the queue

With SLAs stable for 60 days, point governance at yield: false-positive tuning with the credit team, coherence-model reviews against new composite typologies, portfolio-level behavioral-fingerprint sweeps. Certify steady state; hand over on dated, documented criteria.

THE FIVE FAILURE MODES WE ARE HIRED TO PREVENT

Contracting on cost-per-FTE and expecting RAYVI outcomes · accepting an annual certificate as evidence of continuous compliance · treating the agentic layer as a headcount cut instead of an escalation filter · ramping on dates instead of certification gates · discovering the partner bank's vendor-management requirements at the annual review. Every one is prevented in selection and transition — never cured afterward at the same price.

The 2026 standard is not a technology and not a geography. It is a contract that pays for what verification actually produces, an architecture that makes compliance observable between audits, and a forensic bench that catches the fraud built from authentic parts. The Philippine top tier offers all three at 74% below

the onshore build — but only the top tier. Twenty-five genuinely qualified providers hide inside a field of a thousand; disciplined vetting is how they are found, and a fully transparent, competitive RFP is how their best terms are won.

METHODOLOGY & NOTES

How the numbers were built

RAYVI components (fraud admittance, false-positive rejection, review turnaround, audited accuracy) derive from PITON-Global financial-services engagements active 2025–26, measured on client-side decisioning, attribution, and QA systems and validated in client governance. Fraud admittance uses 9-month charge-off attribution; false-positive rates use post-decision review sampling and re-application outcomes.

The 40,000-application cohort (Section 3) is a composite normalized from three engagements to a common application mix and loss severity; the –81% per-clean-verification figure compares Manila vetted against US onshore with loss deductions at observed cohort rates on a peak-quarter basis, and –68% against commodity offshore. In the plain FTE lens the same delivery prices 70–78% below an onshore build depending on role mix (74% at the standard blend).

Cost figures are fully-loaded costs per productive hour — wages, benefits and statutory charges, facilities, technology and security stack (including continuous-monitoring infrastructure), management and QA overhead, attrition-driven recruiting and ramp — net of shrinkage. US baselines reflect Tier-2 metro financial-operations labor; Manila figures reflect 2025–26 engagement pricing.

The case study (FS-041) is anonymized under NDA; volumes, segment descriptors, and platform details are generalized, figures rounded. Value factors are conservative: loss avoidance counts only attributed charge-offs at observed severity with a haircut for portfolio drift; recovered originations apply a discount for applicants who would have completed anyway. $ROI = \text{net first-year value} \div \text{all engagement-related client costs}$. 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 PCI SSC v4.0 requirement texts, US regulatory guidance (FFIEC, FinCEN, CFPB), IBPAP reporting, and PITON-Global's proprietary provider mapping (1,000+ operations; financial-services capabilities re-verified June 2026). The Tripartite Composite typology is PITON-Global's analytical construct, synthesized from client-side fraud attribution; estimates and projections are PITON-Global analysis and should be cited as such.



Which financial-services BPOs belong on your shortlist?

Forty-five minutes with John will tell you. Free to buyers, strictly vendor-neutral — a compliance-gated 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-56 · FINANCIAL SERVICES · PHILIPPINES. This publication is provided for general information and is not legal, regulatory, or compliance advice. Figures identified as estimates or models are PITON-Global analysis. Reproduction with attribution is permitted.

PITON-GLOBAL EXECUTIVE WHITE PAPER SERIES · WP-56 · FINANCIAL SERVICES · PHILIPPINES · JUNE 2026 PITON-GLOBAL.COM