London Business School

From a $5B National Threat to an AI-Enabled Telco Solution

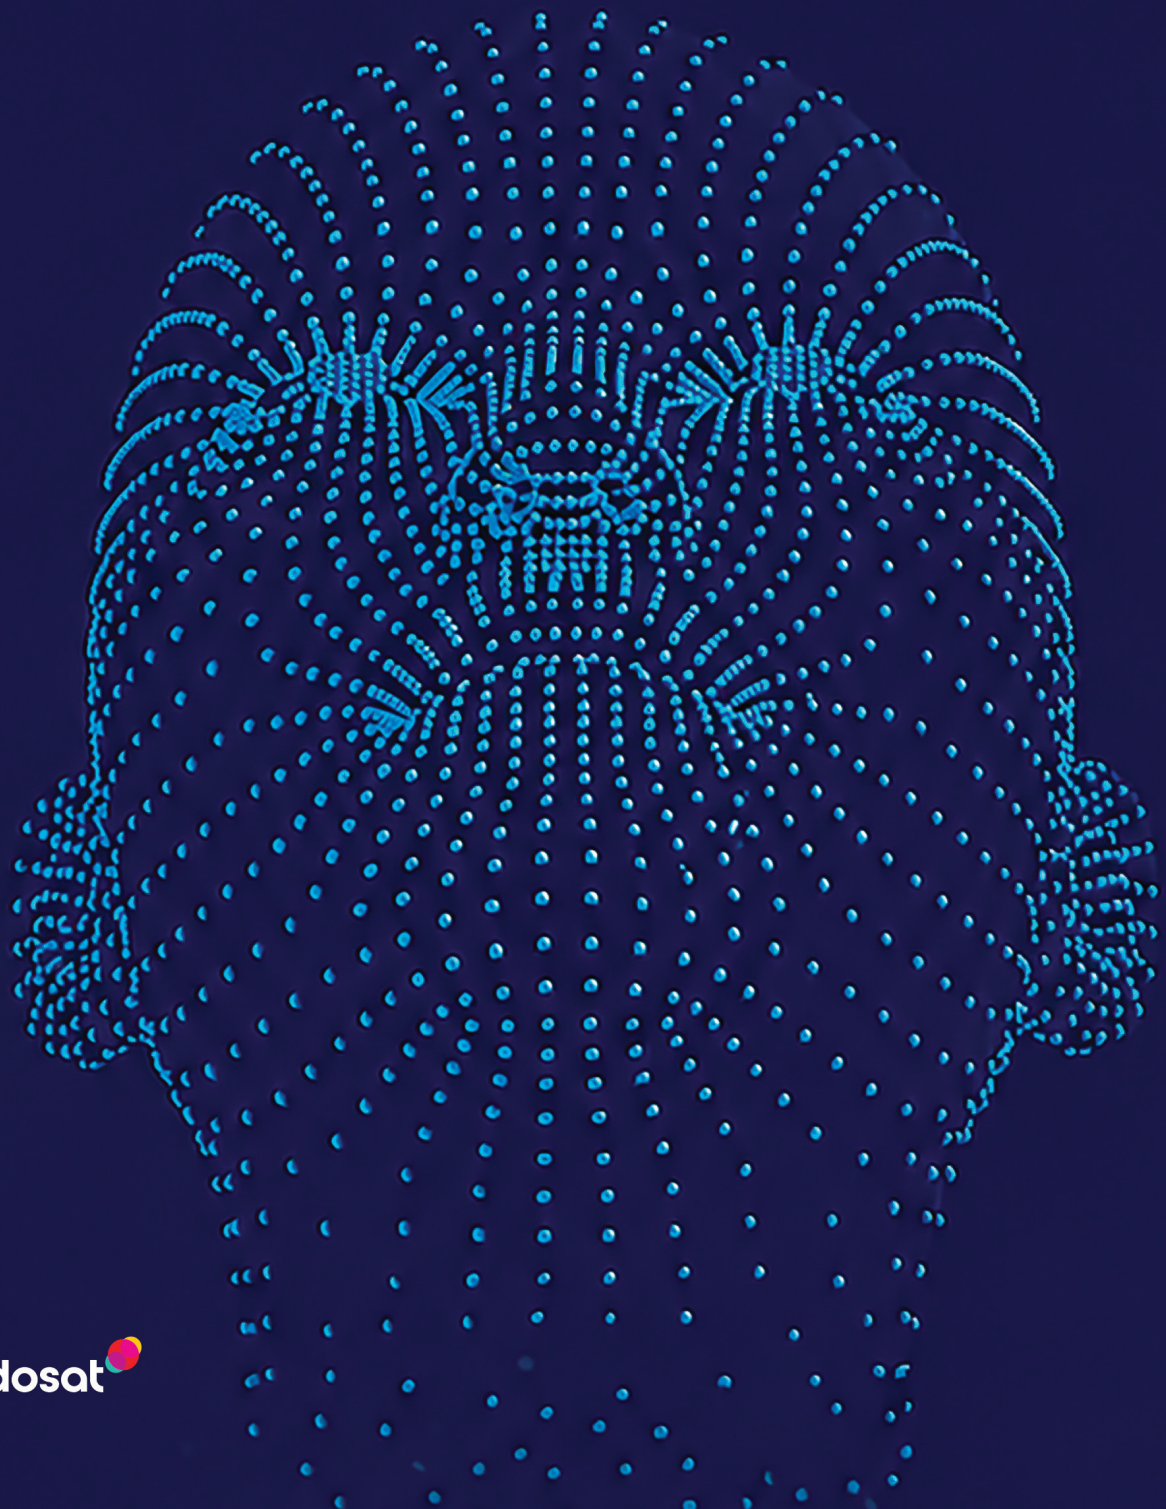

tanla | indosat

London Business School

London Business School
Case Collection

Michael G. Jacobides
M. Dalbert Ma
Shanni Elcock

CS-26-008
May 2026

From a $5B National Threat to an AI-Enabled Telco Solution (A)

In October 2025, Uday Reddy, Founder Chairman & CEO of Tanla Platforms Limited, sat across from Vikram Sinha, President Director & CEO of Indosat Ooredoo Hutchison (IOH), in the latter's Jakarta headquarters. The two had first connected eight months earlier, when Sinha was searching for solutions to create a trusted network to address the epidemic of digital fraud and spam sweeping across Indonesia estimated at ~5B USD lost annually. Now, the early results of their collaboration were in hand. Over the preceding months, IOH had begun rolling out Tanla's AI-powered anti-spam and anti-scam platform to its subscribers — Indonesia's second-largest mobile customer base — with real-time scam and spam alerts already reaching 40 percent of eligible users. The numbers were encouraging: since the launch, the platform had identified 555 million spam communications and 153 million scam attempts, customer satisfaction scores had climbed to 75 out of 100, and regulatory complaints had dropped by 14 percent.

Yet a fundamental tension remained unresolved. Like many services requiring sustained effort to deliver — cybersecurity, insurance, preventive healthcare — the question was not whether the solution worked, but how to commercialise it. The collaboration operated as a B2B2C structure, and each layer faced a different version of this problem. At the consumer end, IOH's customers assumed protection from spam and scams should be a baseline expectation, not a premium service. For IOH, the investment needed an ROI justification — could scam protection generate revenue directly, or would it remain a cost centre justified by churn reduction and brand trust? And for Tanla, the challenge was proving platform value to operators whose end-users struggled to articulate what they were being protected from. The linking question was simple: How do you prove the value of something that, when it works, is invisible?

The technical achievement was real: Tanla's AI platform could identify and intercept scams at scale. But to justify continued investment and rally an ecosystem of complementors around a shared threat intelligence platform, both companies would need to articulate value in terms that resonated at each end of the chain. For IOH, the challenge was providing ROI: could scam protection become a direct revenue-generating service, or would it remain a cost centre justified by second order benefits like churn reduction, or just be a compliance investment? For Tanla, the stakes were strategic: Indonesia represented a proof point for expansion across Southeast Asia, and the metrics emerging from this deployment would shape conversations with other operators in the region. As the meeting concluded, both executives considered the series of choices that had brought them to this moment.

Michael G. Jacobides is the Sir Donald Gordon Professor of Entrepreneurship and Innovation; Professor of Strategy and Entrepreneurship, London Business School. M. Dalbert Ma is a PhD Student in Strategy & Entrepreneurship at London Business School. Shanni Elcock is a Sloan Fellow MSC Student in Leadership & Strategy at London Business School.

A Global Problem, An Indonesian Crisis

By the mid-2020s, digital fraud had become a global epidemic; the Global Anti-Scam Alliance estimated that scammers extracted more than $1 trillion from consumers worldwide in 2024 alone. Indonesia occupied a particularly exposed position. The country ranked first worldwide in the proportion of spam calls and seventh in spam call volume per person, with 89 percent of all unknown calls classified as unwanted and an estimated 35 percent being scam-related. Estimates placed annual losses to Indonesian citizens between $5.0 and $5.5 billion, of which approximately $0.5–0.7 billion stemmed from voice-call scams and $2–2.5 billion from SMS-based phishing. Roughly one in every two SMS messages received was either spam or a scam, and an estimated 15 to 25 million Indonesians fell victim to digital fraud each year.

The pervasiveness of fraud had begun to corrode everyday communication. With roughly 70 percent of inbound calls and messages originating from unknown numbers — spanning scams, unsolicited spam, and legitimate but unfamiliar contacts such as delivery agents — users could not reliably distinguish threat from good communications. Many rationally defaulted to ignoring calls and treating SMS as untrustworthy, including legitimate communications. Promotional SMS click-through rates had fallen to approximately 1 percent, compared to 5–7 percent on OTT messaging apps, while legitimate call answer rates stood at just 26 percent.

Behind these statistics were millions of individuals experiencing tangible hardship. Surveys showed that 51 percent of Indonesian scam victims reported significant stress or trauma, while cultural stigma led many to feel ashamed of having been deceived. Fewer than 7 percent of incidents were ever formally reported — a gap that was not merely a measurement issue but a coordination problem. Any system that relied on victim complaints as the primary trigger for action was, by design, responding after the damage had already occurred.

The Coordination Problem

Structural misalignment between how scams operated and how institutions responded rendered Indonesia's early response largely ineffective. At the heart of the challenge was a severe speed asymmetry. A single scam campaign could push approximately 100,000 SMS messages per hour, cycling rapidly through spoofed numbers and short-lived domains. By contrast, formal redress mechanisms, number takedowns, account freezes, inter-agency coordination, typically took 14 to 16 days. By the time action was taken, an estimated 80 percent of financial damage had already occurred, with perpetrators already migrating away.

Critical information remained fragmented across institutions. Telecommunications providers could observe anomalous traffic patterns but lacked visibility into downstream financial losses. Banks could identify fraudulent transfers but often had little insight into how victims were initially contacted. Digital platforms detected scam advertisements yet operated largely independently of telco and banking systems. Legacy technical controls, static, rule-based firewalls and blacklists designed for traditional spam, proved poorly suited to evolving tactics such as number spoofing, low-volume "grey traffic," and socially engineered messages crafted to evade keyword detection.

Regulatory responsibility mirrored this fragmentation. Komdigi, the Ministry of Communication and Digital Affairs of Indonesia, oversaw telecommunications and online platforms. OJK, the financial regulator, managed fraud reporting related to financial losses. Law enforcement agencies handled criminal investigation. Each operated with its own workflows, evidentiary standards, and response timelines. OJK's public guidance stated that reporting fraud to a bank within twelve hours offered the greatest chance of recovering misappropriated funds, yet formal complaint channels often operated with post-verification service-level agreements of 24 to 48 hours. Most interventions were therefore triggered only after a victim had already suffered losses, placing the burden of detection on consumers who frequently reported incidents late, or not at all.

A series of cross-industry coalitions emerged to address these gaps. SATGAS Waspada Investasi coordinated action against illegal financial activity. Satgas PASTI, launched in November 2023, expanded inter-ministry collaboration. The Indonesia Anti-Scam Centre (IASC), established in November 2024, introduced purpose-built coordination for handling transaction-based scams. While these initiatives improved information sharing, they addressed only part of the problem. Cross-channel coordination still lacked shared service-level agreements and remained unable to operate swiftly. In parallel, many Indonesians turned to device-level solutions such as Truecaller and GetContact, which labelled incoming calls using crowdsourced databases. While these apps could warn or block suspicious calls at the device edge, they were reactive and did not remove scammers from the ecosystem, leaving the underlying coordination failure intact.

What the fragmented system lacked was a mechanism to synthesise dispersed signals into actionable intelligence at speed. The core challenge was integration, and advances in artificial intelligence offered a potential path forward. Machine-learning models trained on network traffic could detect behavioural patterns invisible to rule-based systems, while natural-language processing could identify the linguistic signatures of social engineering in real time. For Indonesia's telecommunications operators, the crisis represented both a threat to customer trust and regulatory standing, and an opportunity for any firm capable of solving a problem that had resisted repeated institutional attempts.

Tanla Platforms: From Connectivity Layer to Intelligent Platform Player

Tanla Platforms, an Indian communications technology company with two decades of experience building messaging infrastructure across emerging markets, believed its AI-native integrated solution could address Indonesia's problem. By 2025, Tanla processed more than 800 billion interactions annually across its platforms, giving the company visibility into messaging behaviour at a scale few others could match. More importantly, Tanla had evolved beyond simply transmitting messages. It had begun actively analysing them, developing AI capabilities that could detect scams not by matching keywords against static blacklists, but by recognising behavioural signatures of fraud across millions of communications in real time.

Tanla also had in its portfolio another business: Communications Platform as a Service (CPaaS) infrastructure that enabled enterprises to send messages through telecommunications networks. When banks issued transaction alerts, e-commerce platforms sent delivery notifications, or ride-hailing apps confirmed bookings, those messages often flowed through platforms like Tanla's. By 2025, Tanla commanded approximately 35 percent of India's CPaaS revenue market and handled roughly 63 percent of the country's A2P SMS traffic, serving more than 2,000 enterprise customers, including Google, Meta, and Truecaller.

Tanla's strategic orientation had been shaped by crisis. In 2012, the collapse of Nokia and BlackBerry wiped out approximately 80 percent of the company's revenues, as the platforms through which most messaging traffic flowed rapidly disappeared. The critical inflection point clarified where Tanla's durable competitive advantage lay — not in application-level features that could be replicated or displaced, but in infrastructure-layer capabilities characterised by deep telco network integrations, stringent regulatory requirements, and deep relationships with telecom operators and large enterprises. This philosophy later became central to Tanla's approach to scam prevention: embedding detection into the network itself complemented by, rather than just another consumer-facing application at the edge.

Indosat and the Tanla Partnership

Indosat Ooredoo Hutchison (IOH) was formed in January 2022 through the $6 billion merger of Indosat Ooredoo and Hutchison 3 Indonesia, creating Indonesia's second-largest telecommunications operator. By 2025, the combined entity served approximately 100 million subscribers. The integration was widely regarded

as successful: total revenue increased 49 percent year-over-year in 2022, while net profit rose 76 percent. But growth had slowed in recent times with the top line flattening for the first 2 quarters of 2025, and IOH was looking for channels to reinvigorate growth.

Beyond financial performance, the merger prompted IOH to articulate a new strategic identity. At its Capital Markets Day in 2024, leadership introduced an "AI North Star" anchored in three ambitions: becoming an AI-Native Telco, an AI TechCo, and an AI Nation Shaper. This agenda positioned IOH as a contributor to "Golden Indonesia 2045," the country's long-term national development vision. Within this framing, the anti-scam and anti-spam initiative was not treated as a narrow product feature but as a flagship application of AI deployed at national scale to address a widespread social problem — IOH's leadership highlighted that more than 65 percent of Indonesians had experienced spam or scams, framing digital fraud as a crisis affecting the majority of the population. The problem statement was deeply relevant for the telecom industry. Users increasingly hesitate to engage with telecom channels like SMS and Voice, undermining the effectiveness of even genuine outreach.

IOH's commitment to addressing the scam problem initiated a period of extensive discovery. The team conducted systematic analysis of leading anti-scam anti-spam solutions globally and identified a consistent pattern: operators in mature markets typically deployed protection in narrow slices — call labelling, SMS filtering, or app-based security — without integrating all three. Many approaches operated either at the network layer or at the application layer, often relying on customer adoption to unlock full functionality, with more limited coverage across voice and cross-channel scam behaviour (see Exhibit A). These gaps reinforced IOH's view that a more integrated, network-native approach might be required. IOH evaluated Tanla's platform alongside incumbent vendors offering AI-enhanced firewalls, point solutions focused on single channels, and bespoke IT-services providers. Because Tanla's approach differed materially from existing offerings, direct comparison proved difficult, and Tanla's team spent significant time educating and deliberating with IOH on appropriate metrics for evaluation.

The partnership began not as a procurement exercise but as a strategic alignment between the two CEOs. Sinha's brief to Reddy was straightforward: he wanted Tanla to help make Indosat the most trusted network in Indonesia, and to develop an AI-native platform that could become a "Global Lighthouse" — a reference deployment for the region. The relationship was structured as a co-development partnership rather than a conventional vendor contract. To validate the approach, the two teams conducted a proof of concept using live network data, requiring deep integration across IOH's systems like CDRs, CRM, Billing giving the AI models access to roughly 200 data signals per interaction. The sensitivity of this data imposed strict constraints: the platform was hosted on-premises in Indonesia to ensure data sovereignty, and both parties agreed to rigorous privacy and security protocols. Beyond validating detection performance, the pilot surfaced the operational dependencies required for AI to function at scale — model accuracy depended on high-quality, real-time data feeds, and formalised service-level agreements governing data availability, latency, and expected outcomes.

Building Wisley Ai: From Rules to Intelligence

In 2019, Tanla launched Trubloq, a compliance-focused spam prevention platform deployed across Indian telecommunications networks. Trubloq addressed a bounded problem: filtering known categories of spam and ensuring that messaging traffic conformed to regulatory requirements. The platform proved effective, reducing spam messages per user by approximately 60 percent and eventually serving more than 60 percent of India's mobile subscribers. At the same time, Trubloq had its own limitations. Built around predefined rules and blacklists, it performed well against familiar threats but struggled with novel scam formats, first-time senders, and low-volume fraud designed to resemble legitimate communication. These limitations were structural rather than incidental: scammers adapted faster than rules could be written, and the system could not learn autonomously

from its own experience. Tightening filters increased false positives; loosening them allowed harmful traffic through.

Tanla's early responses focused on protecting enterprise customers, particularly banks, whose messaging channels were being impersonated. Yet even in these controlled contexts, content-based detection proved insufficient. Scammers learned to neutralise obvious linguistic markers, borrowing the tone and structure of legitimate communications while subtly rotating content to evade keyword-based filters. This realisation prompted a fundamental conceptual shift: from analysing what messages said to understanding how they behaved within the network. A message originating from a new sender, transmitted in sudden bursts to thousands of recipients, and containing links that had appeared across unrelated campaigns exhibited a behavioural signature fundamentally different from a bank's routine transaction alert — even if the textual content appeared similar.

Wisely Ai represented Tanla's effort to build an AI-native platform ground up rather than simply layering machine learning onto existing workflows. In AI-enhanced systems, models might flag suspicious traffic or optimise rule thresholds, but final decisions remained anchored in predefined logic. Wisely Ai was designed differently from the outset. Detection was organised around behavioural pattern recognition, evaluating messages through multiple concurrent signals — textual features, sender behaviour, traffic velocity, link reuse, historical outcomes, and network-level correlations —continuously recalibrated as new data entered the system. The objective was not just perfect accuracy at any single moment, but the capacity to learn faster than adversaries could adapt.

Wisely Ai was built as a network-native system that integrated directly into the telecom operator's infrastructure rather than sitting on top of it as an application layer. This allowed IOH to deploy the platform in roughly three months and, supported multiple solutions — anti-scam, anti-spam and others released later — from a shared infrastructure, reducing the marginal cost of adding new capabilities. The system ran over ten AI agents simultaneously, analysing sender reputation, message semantics, call-to-action patterns, and other signals, each making autonomous classification decisions in real time. The platform was also self-learning: as scammers adapted their tactics, the models updated continuously — a feature considered essential given the pace at which fraud techniques evolved in the Indonesian market. The platform enabled multiple consumer-facing solutions at the network, device, and SIM levels.

Building Wisely Ai required capabilities beyond Tanla's historical strengths in telecom-scale systems engineering. The company made a deliberate organisational choice to combine selective senior hires with intensive internal capability development, embedding AI engineers alongside product, operations, and compliance teams rather than centralising AI as a standalone function. While this approach slowed early development, it reduced friction during deployment and improved coordination between model design, operational constraints, and regulatory requirements.

AI's Tangible ROI Problem

Technical capability alone did not resolve Tanla's most pressing challenge: translating societal value into a sustainable business model. Tanla's traditional revenue model was volume-based, more messages sent through its platforms generated more fees. Scam prevention sat uncomfortably within this logic. Blocking fraudulent messages reduced traffic and protecting consumers did not generate direct revenue from those consumers.

The challenge was structural. Tanla did not sell to consumers; it sold to telecommunications operators, who in turn served consumers. This created a B2B2C value chain with misaligned incentives at each link. At the end of the chain, consumers were the primary beneficiaries of scam protection, yet they did not perceive protection as a service worth paying for. Scam-free communication felt like a baseline expectation rather than a premium offering.

This perception placed telecommunications operators in an awkward position. If consumers would not pay directly for protection, operators had to justify the investment through indirect benefits, reduced churn, improved customer satisfaction, regulatory compliance, or brand differentiation. While these benefits were real, they were diffuse, difficult to quantify, and harder to defend in internal budget discussions than initiatives with explicit revenue lines.

Regulators represented a potential forcing mechanism. A mandate from Komdigi could compel operators to invest in scam protection, eliminating the discretionary nature of the decision. However, regulatory mandates carried their own risks: locking in specific technical approaches, discouraging experimentation, or encouraging compliance-driven box ticking rather than genuine protection.

For Tanla, progress required a partner willing to experiment. The company needed a telecommunications operator prepared to deploy the technology at scale, test monetisation approaches, and generate credible evidence that could shape both commercial strategy and regulatory expectations. Without such a partner, Tanla risked building a technically elegant solution that failed to find durable economic footing.

Commercial Structure

Several monetisation structures were explored through joint workshops involving IOH's digital, commercial, and network teams alongside Tanla's product and strategy leadership. Early discussions reflected the partners' different starting points: IOH evaluated monetisation primarily through customer experience and retention outcomes, while Tanla approached the problem through scalable platform economics. Per-message pricing was quickly set aside because it created a structural contradiction — the more effective the system became at reducing scam traffic, the weaker the revenue signal it generated. Direct consumer subscription models were debated within IOH's digital and marketing teams, but customer research suggested limited willingness to pay for a service perceived as a basic expectation of network safety. Enterprise-funded models were also explored, drawing on Tanla's prior experience with banks and enterprise senders, yet proved incomplete as a significant share of scams originated from peer-to-peer channels rather than enterprise messaging flows. Outcome-based pricing linking revenue to reductions in complaints, fraud incidents, or customer harm gained conceptual support but proved difficult to implement: while Tanla could commit to scam/spam detection performance, factors like churn reduction or brand trust were influenced by multiple variables beyond the platform itself (See Exhibit B).

Ultimately, the parties converged on a blended structure combining licensing elements with usage-based considerations, without mechanically tying revenue to scam volumes. IOH paid on a per-subscriber-protected-per-month basis, anchored by tightly defined service-level agreements and governance mechanisms covering data ownership, model training rights, and escalation protocols. Decision rights over threshold settings, feature prioritisation, and deployment sequencing were shared, reflecting the joint accountability inherent in deploying AI within a regulated national infrastructure. While Tanla retained responsibility for model development, platform evolution, and AI performance, IOH remained accountable for regulatory compliance, customer communication, and network-level enforcement decisions.

Performance discussions increasingly focused on operational metrics alongside commercial outcomes. Teams tracked AI efficacy across both SMS and voice channels, monitoring detection performance while balancing false-positive risks. Platform latency became a central consideration because predictions were executed directly on the transmission path of messages and calls, requiring decisions to be made in near real time without disrupting network performance. These metrics helped clarify accountability boundaries: Tanla committed to technical efficacy and system responsiveness, while IOH evaluated downstream business outcomes such as customer trust, complaint reduction, and brand perception.

Scaling Deployment for National Context

Indonesia proved a defining deployment environment. The country's digital ecosystem characterised by rapid growth, linguistic diversity, and heterogeneous usage patterns tested Wisely Ai beyond initial expectations. Scam messages frequently relied on colloquial phrasing, local idioms, and culturally specific cues that were poorly captured by models trained on other markets. Fraudsters also exploited local payment practices and social norms, producing scam patterns that differed materially from those previously observed.

As Wisely Ai transitioned from pilot environments into live telecommunications networks, Tanla encountered a challenge common to virtually all AI-driven systems: the gap between controlled evaluation and real-world deployment. In pilot settings, conditions are necessarily simplified — data is often partially labelled or retrospectively analysed, latency constraints are relaxed, and false positives carry limited operational consequences. In evaluation, for instance, the model was tested against traffic where scam patterns were relatively distinct. Live deployment introduced far greater ambiguity: legitimate messages from small businesses, sent in irregular bursts, could resemble scam campaigns, while fraud attempts wrapped in culturally familiar language — holiday greetings, informal payment requests — looked much like ordinary communication. Navigating these grey zones required a sensitivity that only real-world exposure could fully develop. At national scale, the operating demands also shifted: Wisely Ai was now required to analyse vast volumes of unlabelled traffic in real time, under strict latency constraints, where decisions carried immediate consequences — incorrectly blocking legitimate messages could disrupt commerce, erode user trust, and expose operators to regulatory scrutiny.

Tanla's teams described the effort as a combination of reuse and reinvention. Approximately 70 percent of Wisely's core: its behavioural models, system architecture, and learning frameworks translated directly. The remaining 30 percent required not just technical adaptation — language-specific features, local scam typologies, revised governance — but sustained human effort. Engineers monitored outputs daily, adjusted thresholds as scam tactics evolved, and exercised judgment on ambiguous cases the models could not resolve on their own. AI, in this context, was not a self-running system but a continuously tended one. This carried direct implications for pricing: the ongoing engineering effort needed to keep the platform effective was a real, recurring cost — yet it was precisely this human work that proved hardest to make visible to the operators and consumers who benefited from it.

SATSPAM Launch and Tiered Offering

In August 2025, IOH launched SATSPAM as a network-level detection and alerting service that identified suspected spam and scam calls and SMS in real time. The name "SATSPAM" was intentionally chosen to anchor the service in a familiar mental model: satpam is the Indonesian term for a security guard. The metaphor positioned SATSPAM as a guard for users' communications channels, rather than "just another telco feature."

The offering was structured into two tiers, reflecting a deliberate trade-off between universal protection and monetisation:

- **Basic Tier:** SMS network protection and Call protection automatically available to all VoLTE-enabled IOH subscribers at no cost. Alerts appeared via voice caller ID and SMS prefixes, warning users before engagement while leaving the decision to answer or ignore with the user.
- **Plus+ Tier:** The SDK innovation created the grounds for a premium tier, designed to make protection inspectable and interactive, the Plus+ tier surfaced within the IOH app. Users could view labelled call and SMS histories, block or report numbers, and interpret colour-coded risk cues (e.g., red for high risk, green

for allowed, orange for caution). The design intentionally prioritised end user experience over technical explanation: users were not asked to understand the underlying algorithms, only to trust what the interface made visible. The experience made the appropriate user action more intuitive.

Consent and permissions quickly emerged as a constraint. To enable protection, customers had to grant access to call logs, and contacts. The team observed immediate drop-off at permission prompts, in part because requests from a “telco app” felt more intrusive than similar requests from specialist apps such as Truecaller. IOH responded by simplifying flows, improving FAQs, reducing the number of required clicks, and remaining explicit about consent boundaries to meet regulatory expectations.

By October 2025, approximately 20 million users (37 percent of IOH’s subscriber base) had been protected in the Basic tier, while 1.2 million users had adopted the Plus+/SDK experience, with the Plus+ tier contributing to ARPU growth through mix shift toward higher-value plans.

Inside IOH, tiering and monetisation were led by the CMO and marketing organisation. Market research validated demand, followed by cohort-based analysis to identify segments likely to accept premium features without suppressing overall adoption. The team favoured bundling Plus+ features into higher-value plans rather than positioning them as a standalone subscription — reflecting a clear philosophy: the Basic tier would ensure universal access, while value creation would come from deeper engagement rather than direct security charges. Plus+ functioned as an implicit upgrade lever rather than an overtly priced service, directly addressing the attribution challenge: users struggled to pay for protection they could not easily perceive, so the approach sequenced perception-building ahead of monetisation.

Continuous Innovation and the Adoption Challenge

Governance was intentionally high-touch during early deployment, with cross-functional teams from both organizations meeting frequently to review performance thresholds, escalation protocols, and rollout sequencing — with CEO-level alignment critical in maintaining momentum, particularly as the initiative was framed as part of IOH’s broader ambition to become a trusted AI-native operator.

Once deployed, the platform’s effectiveness was constrained by infrastructure realities that neither party had fully comprehended. Analysis of IOH’s network revealed that only around 30 percent of subscribers were VoLTE-enabled, meaning network-level caller identification alerts could not reach the majority of the user base for voice calls. Rather than accept this as a coverage limitation, Tanla proposed developing a software development kit (SDK) embedded in IOH’s mobile app, allowing non-VoLTE users to receive scam and spam alerts. A second insight emerged from user feedback: Indonesian consumers made an estimated three to five times more calls over VoIP services such as WhatsApp than over the traditional telco network, meaning the majority of voice interactions would remain unguarded if the platform only protected network calls. This led to the addition of VoIP detection capability within the SDK in November — enabling scam alerts on WhatsApp calls, not just telco network calls. Neither the SDK nor the VoIP capability had been part of the original contract; both emerged from market learning and gaps realized to drive desired impact.

Building new features was only half the problem. The platform could only deliver impact if users actually adopted it, and adoption faced structural constraints: IOH’s mobile app had a penetration of roughly 50 percent of the subscriber base, and activating the SDK required app version updates and explicit permission consent from users. IOH’s initial target for SDK activations was one million within five months. Through a combination of product design changes and targeted campaigns — treating the rollout more like a consumer app launch than a network upgrade — the teams exceeded two million activations by December and were on track to reach five million.

More broadly, Tanla framed its approach to the partnership around a five-stage model of anti-scam maturity: identification of bad actors, alerting users, preventing delivery of malicious communications, elimination of offending numbers from the network, and legal enforcement. Progress through each stage required coordination with Indonesia's telecommunications regulator, with each phase building regulatory confidence before advancing to the next.

Shaping the Regulatory Framework

When IOH initiated the anti-scam programme, Indonesia had no regulatory requirement for telecom operators to protect users from scam or spam communications. Responsibility for fraud prevention fell primarily on financial institutions. There was no established framework for network-level intervention by telcos.

IOH and Tanla engaged proactively with Komdigi to build the case for operator-led protection. The regulator raised concerns around data privacy, user consent, and the implications of AI systems analysing message content. Drawing on experience from deployments in other markets, Tanla worked with IOH's regulatory team to address these concerns directly. Among the mechanisms introduced was a user opt-out facility, tracked and reported regularly to the regulator. Opt-out rates remained in single digits in the first months and declined further — providing evidence of broad user acceptance across millions of subscribers. The Vice-Minister of Komdigi participated in the platform's launch in August 2025.

Initially, the platform was limited to identifying scams and alerting users, but subscribers remained exposed to malicious messages designed to prompt action. IOH and Tanla returned to the regulator to seek approval for expanded enforcement capabilities: blocking SMS identified as scams through AI detection, suppressing malicious domains via DNS-level controls, deactivating scammer phone numbers from the network, establishing a blockchain-based threat intelligence exchange, and defining baseline metrics for joint improvement. Regulatory approval proceeded in stages — alerting only from August 2025, blocking approved in November 2025, and full enforcement completed in January 2026 (see Exhibit C).

The impact extended beyond IOH. In November 2025, Komdigi issued a notice to all Indonesian telecom operators to urgently enhance consumer protection systems, citing the surge in phone and SMS-based fraud and the substantial financial losses it was causing nationwide. IOH's deployment had, in effect, set the regulatory benchmark for the industry.

Transitioning from alerting to automatic blocking carried significant brand and regulatory risk: incorrectly blocked messages could disrupt legitimate communications and expose IOH to complaints. Proceeding required confidence in the platform's detection accuracy, scalability, and willingness to prioritise customer protection ahead of immediate monetisation.

The Value and Attribution Challenge

IOH's go-to-market strategy focused on making prevention tangible and comprehensible without over-explaining the underlying AI or product variants. Public communications relied on press briefings featuring real victim stories and headline statistics on scams flagged and blocked. At the point of sale, SIM card packaging was updated to signal that anti-scam protection was included by default. Within the Plus+ dashboard, colour coding, alerts, and summaries conveyed protection outcomes in the simplest possible terms. The underlying belief was that, for prevention products, comprehension and trust — not feature depth — drive adoption.

Yet customers broadly perceived SATSPAM's benefits less as a premium offering than as a basic obligation

of the mobile network: protection should be implicit, always on, and largely invisible. When safety worked, nothing happened — and nothing happening was precisely what users expected. IOH's leadership increasingly converged on the view that invisibility was not a weakness of the value proposition, but its essence. This tension was not unique to telecommunications: credit card networks had resolved similar challenges by embedding fraud protection into interchange fees, while insurance offered a different logic entirely. The strongest signal of success would be quiet, reliable operation in the background, manifested through fewer scams and a growing sense that the network itself could be trusted.

Customer research conducted in September 2025 revealed this ambivalence (see Exhibits D and E). Among activated users, headline indicators were positive: Net Promoter Scores ranged from +20 to +25, customer satisfaction scores from 75 to 87, and approximately 95 percent described the service as feeling "protected" or "very protected" (See Exhibit F). At the same time, roughly one-quarter of users reported no noticeable difference after activation, approximately 30 percent did not clearly understand what the service did, and others remained sceptical of effectiveness. Users could feel safer in principle while still doubting effectiveness in practice.

To organise measurement, IOH developed a three-tier framework: input metrics captured operational activity such as users protected and scams flagged; output metrics tracked customer response including NPS, CSAT, and regulatory complaints; and outcome metrics focused on business impact — churn, ARPU, and competitive position. The service was framed internally as both a retention lever and a potential acquisition differentiator in a price-competitive market, with migration from Basic into Plus+ serving as a concrete signal of perceived customer value. The central difficulty lay in attribution: proving that anti-scam capabilities caused downstream business outcomes rather than merely correlating with them.

By November 2025, early indicators were cautiously encouraging., Trouble tickets declined to single digits, and regulatory complaints fell meaningfully (see Exhibit G). Coverage expanded from 9 million protected users in August to 20 million by November, with a target of 40 million by March 2026, while Plus+ SDK adoption tripled from 392,000 to 1.2 million users, targeting 5 million by March 2026 (see Exhibit H). Yet fundamental questions remained unresolved: a significant share of users still did not understand the service, and ultimate business outcomes — reduced churn, higher ARPU, improved brand equity — remained hypotheses rather than proven results.

From Firm-Level Protection to Ecosystem Coordination

As deployment progressed, a structural limitation became increasingly apparent. Scams rarely remained confined to a single channel or network. A fraudulent SMS might redirect a victim to WhatsApp, lead to a counterfeit website, and culminate in a bank transfer — each step falling under the jurisdiction of a different institution, none of which had visibility into the full end-to-end journey.

Traffic data illustrated the scale of the challenge. In October 2025 alone, of the 60.5 million scam messages detected, 15.0 million were IOH-to-IOH, 16.5 million were outbound from IOH to other networks, and 29.5 million were inbound from other networks to IOH — meaning cross-network scam traffic exceeded on-network activity, and protection at the level of a single operator could displace rather than eliminate fraud. Other industries facing systemic risk had encountered similar coordination problems: cybersecurity firms pooled intelligence through sector-wide sharing mechanisms, financial institutions collaborated on global payment infrastructure, and lenders exchanged borrower data through credit bureaus. These precedents suggested that when threats were collective, intelligence needed to be collective as well — extending beyond identification and alerting to elimination and legal enforcement.

These insights informed discussions around a proposed Threat Intelligence Exchange: a shared, real-time platform through which telecommunications operators, banks, fintechs, digital platforms, and regulators could

exchange scam intelligence. The ambition was to provide cross-ecosystem visibility, establish verifiable evidence trails for enforcement, and enable coordinated intervention across channels, governed by "give-to-get" principles that aligned incentives to contribute data rather than free-ride. IOH's participation in broader anti-scam alliances signalled that SATSPAM was being positioned not merely as a proprietary product but as part of a national ecosystem response — with shared intelligence as the substrate for faster and more effective cross-channel intervention than any single institution could execute independently.

Designing a sustainable commercial model for such an ecosystem raised additional questions. Membership fees could fund core infrastructure, usage-based pricing could align costs with activity, and contribution-weighted access could reward participants who shared more data. Alternatively, regulators could mandate participation or subsidise the system if scam protection were treated as a public good. Yet success carried a paradox: if shared intelligence became an industry standard, could it erode the differentiation that had motivated IOH's investment?

Strategic Tensions and the Path Ahead

These dynamics presented IOH with a strategic fork. One path was to preserve differentiation by keeping Wisely Ai proprietary, leveraging superior protection to attract and retain customers, and resisting ecosystem-wide standardisation — maximising competitive advantage but leaving the broader scam problem only partially addressed. The alternative was to lead ecosystem development, accepting some erosion of differentiation in exchange for regulatory goodwill, reputational benefits, and a central role in shaping governance and standards. This option aligned closely with IOH's stated ambition to contribute to Indonesia's digital transformation, though its commercial payoff was less immediate.

With credible network-level data in hand, IOH faced further design questions: whether coordination should focus on telecommunications competitors alone or expand cross-sector to include global platforms such as Mastercard, Meta, and Google. A parallel commercial opportunity also emerged on the B2B A2P messaging side. As protection tightened, IOH anticipated developing a model in which enterprise senders could be verified or whitelisted to reduce false positives, enabling tiered commercial arrangements based on duration or message volume. In this framing, SATSPAM was not only a consumer protection tool but also a mechanism for restoring trust in promotional SMS and calls by distinguishing verified senders from malicious actors.

For Tanla, complementary tensions emerged. A fragmented market favoured per-operator deployments; a shared ecosystem model could reduce individual contracts while positioning Tanla as the technical backbone of national infrastructure. Wisely's defensibility lay less in algorithms or architecture which competitors could replicate than in cumulative learning acquired through operating at scale under adversarial conditions, embedded in workflows, and judgement that could not be easily transferred or compressed.

As 2025 drew to a close, the technology had proven capable and early results were promising. Yet success raised as many questions as it answered. How should protection be priced when customers believed it should already exist? What role should IOH play in shaping the broader ecosystem? Could collective intelligence coexist with competitive advantage? The answers would shape not only the future of the Tanla–IOH partnership, but the trajectory of Indonesia's response to its growing scam crisis.

FIGURES & TABLES

Exhibit A

One Platform, One Architecture, Many Solutions

APPLICATION LAYER			
B2C	**B2B**	**B2T**	
• Spam & scam protection on SMS, Voice, and OTT • Auto-block spam and scam calls and SMS	• Verified Business Identity • Threat Assessment API • Authentication services	• SSOT platform • Experience Panel • Analytics & Insights	
↕	↕	↕	↕
AI PLATFORM			
Prediction Engine Real-time threat predictions using AI models	**Validation Engine** Validates against known threat patterns	**Self-Learning Engine** Improves accuracy via continuous feedback loops	**Autonomous** Operates without human intervention
↕	↕	↕	↕
DATA & NETWORK INFRASTRUCTURE			
Data Sources • Multiple integrated data sources • Infrastructure-level security	**AI Infrastructure** • Prediction latency <10ms • NVIDIA Blackwell architecture	**Compliance & Security** • 100% local regulatory compliance • Vulnerability assessment per InfoSec standards	
KEY PLATFORM CHARACTERISTICS			
Deployment Local (on-premises) or cloud	**Multi-Level** Network, device, and SIM levels		**Compliance** 100% compliant with local regulations
Modularity Adaptable to specific business needs	**Technology** NVIDIA Blackwell architecture		**Verification** Blockchain-based verification

Source: Tanla Product Presentation

Exhibit B

The B2B2C Pricing Dilemma: Why Standard Pricing Models Proved Inadequate in This Context

Pricing Model	How It Works	Who Bears the Cost	Challenges & Constraints
Per SMS/Voice Pricing	Tanla charges IOH per message processed or blocked through the platform. Revenue scales with traffic volume.	IOH — passed through as operating cost	Structural contradiction: more effective blocking reduces scam traffic, weakening the revenue signal — the platform is penalised for succeeding.
Consumer Subscription	IOH charges subscribers a recurring monthly fee for scam/spam protection as a distinct product. Debated by IOH's digital and marketing teams before being rejected.	Consumers — direct payment	Customers perceive protection as a baseline expectation of network safety, not a premium service; low willingness to pay.
Enterprise-Funded	Enterprises (banks, e-commerce, government) pay to be verified or whitelisted as trusted senders, addressing false-positive risk for high-volume legitimate senders.	Enterprise senders — B2B revenue	Covers only A2P enterprise messaging scams; a significant share of Indonesian fraud originates from P2P channels, leaving a material gap.
Outcome-Based	Revenue linked to measurable reductions in complaints, fraud incidents, or churn. Gained conceptual support but proved difficult to implement in practice.	Shared risk between IOH and Tanla - Contingent on verified outcomes	Attribution problem: churn reduction and brand trust are influenced by multiple variables beyond the platform itself.

Core dilemma
In B2B2C chains with 'invisible products', investment costs sit with the operator while value accrues to consumers — and the technology provider is upstream of both. Each standard pricing model breaks down because the service is invisible when it works, bundled into the network experience, and resistant to attribution.

Exhibit C

Future implementation roadmap – target to launch blocking by December

Note: non-exhaustive roadmap (illustrative).

Category	Owner	Actions	2025			2026			
			Oct	Nov	Dec	Jan	Feb	Mar	Apr & beyond
Product actions	Tanla	1. VoIP support (Including WhatsApp)		▲					
		2. DNS suppression			▲				
		3. Launch Bureau[1]			▲				
Regulatory actions	IOH + Tanla (regulatory leads)	4. Approval to degrade service for suspected scammers							
Operational actions	IOH + Tanla	5. Ecosystem takedown via collaboration: A. Disable CTAs[2] for IOH users					Potential to start		▶
		B. Tech giants (e.g., Google, WA)				Takedown[3]			
		C. Banks				Secure users from fraud			
		D. LEA and regulators				Enforcement			
		E. Other telcos					Widen reach		▶
IOH	IOH	6. Scale SDK adoption	650K	1 Mn		5 Mn			▶

Notes: ▲ Triangles indicate Planned Milestones; Bars indicate planned activity windows; ▶ Arrowheads indicate activity expected to continue beyond April 2026

1. *Bureau denotes a threat intelligence capability referenced in the original materials*
2. *CTAs denotes Call to Action elements (as used in the original materials)*
3. *Takedown refers to coordinated platform-level removal actions*

Source: Tanla Implementation Roadmap

Exhibit D

Survey results - perceived effectiveness and advocacy for an anti-scam/spam feature

A. How helpful users find the feature (%, by segment and region)

Segment	IM3					3					IM3 (Platinum)
Measure	National	Jaya	Java	Kalsumapa	Sumatera	National	Jaya	Java	Kalsumapa	Sumatera	National
Base (n)	644	152	122	221	146	232	127	110	97	135	86
Very helpful	48%	45%	41%	59%	50%	45%	48%	43%	42%	48%	63%
Helpful	33%	30%	35%	26%	37%	31%	31%	33%	37%	28%	22%
Very helpful + Helpful	81%	75%	76%	85%	82%	76%	79%	76%	79%	76%	85%
Other re-sponses	19%	25%	24%	15%	18%	24%	21%	24%	21%	24%	15%

B. Likelihood to recommend the feature (share of respondents; NPS shown above)

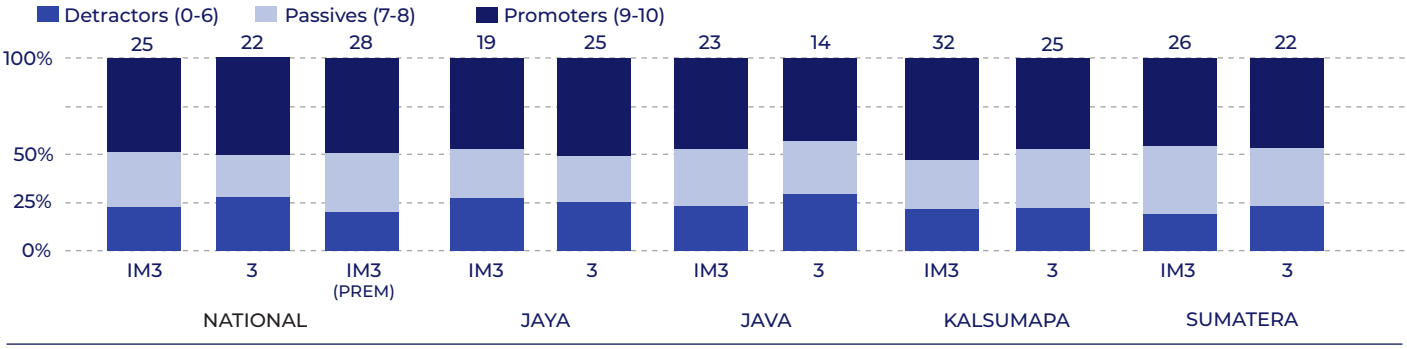

Note: NPS (shown above bars) = % Promoters – % Detractors.
Source: Company survey data; authors' adaptation.

Exhibit E

Customer experience after activating an anti-scam/spam feature

A. CSAT by region and segment (score out of 100)

Region	IM3	3	IM3 (Platinum)
National	79	75	76
Jaya	79	78	—
Java	76	71	—
Kalsumapa	87	81	—
Sumatera	75	71	—

B. Stated reasons for satisfaction (%, multiple responses allowed)

Reason	IM3					3					IM3 (Platinum)
	National	Jaya	Java	Kalsumapa	Sumatera	National	Jaya	Java	Kalsumapa	Sumatera	National
Base (n)	**644**	**152**	**122**	**221**	**140**	**262**	**188**	**63**	**35**	**65**	**48**
Fewer unwanted calls/ messages (unknown numbers)	45%	40%	43%	39%	36%	41%	38%	38%	51%	43%	46%
Protects personal data and digital transactions	36%	34%	42%	31%	36%	33%	34%	29%	37%	30%	36%
Blocks scam/ spam SMS and calls	28%	34%	31%	20%	28%	33%	36%	31%	34%	28%	30%
Feel safer using the service	28%	27%	27%	31%	26%	21%	21%	27%	24%	22%	31%
Easy to use; no complex settings	28%	16%	18%	21%	15%	19%	32%	21%	18%	18%	23%
More comfortable experience	9%	8%	7%	11%	10%	14%	24%	16%	13%	15%	17%

C. Stated reasons for dissatisfaction (%, multiple responses allowed)

Reason	IM3					3					IM3 (Platinum)
	National	Jaya	Java	Kalsumapa	Sumatera	National	Jaya	Java	Kalsumapa	Sumatera	National
Base (n)	644	152	122	221	140	262	188	63	35	65	48
No notice-able change after activa-tion	26%	38%	22%	22%	20%	29%	25%	34%	15%	34%	42%
Still receive spam/scam SMS or calls	23%	29%	22%	7%	27%	27%	30%	25%	26%	24%	33%
Low trust that feature prevents scams	23%	18%	24%	15%	31%	23%	27%	19%	11%	26%	13%
Unclear in-formation on how it works	20%	20%	30%	22%	20%	23%	22%	22%	37%	19%	21%
Does not al-ways detect suspicious nos.	18%	18%	12%	26%	20%	12%	11%	18%	11%	10%	25%

Note: Respondents could select multiple reasons; percentages are shown as reported for each segment/region.
Source: Company survey data; authors' reconstruction and formatting (de-branded).

Exhibit F

SATSPAM user: Is SATSPAM delivering protection to our user and which features stand out?

Most of current SATSPAM users feel secure post-activation especially on features that relevant for them such as SDK caller ID and SMS text prefix. Low visibility of Weekly Digest signals an opportunity to enhance communication and drive feature value.

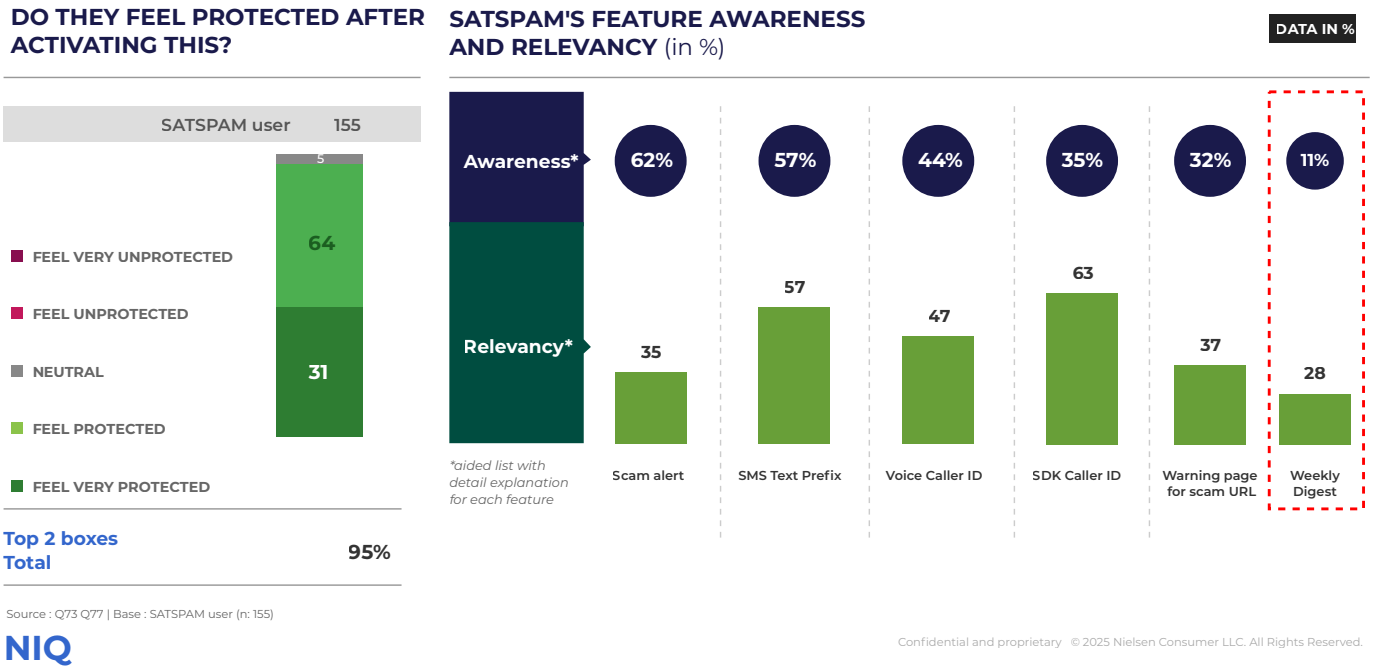

Exhibit G

Monthly interactions and trouble tickets for an anti-spam/scam feature

IM3 — Aug-Oct 2025 summary

Interactions (count)

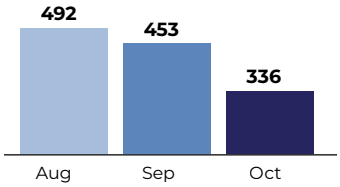

Top interaction types (count)

Category	Aug	Sep	Oct
Info seeking	591	472	307
Service request	31	15	22
Complaint	5	5	7

Top trouble tickets (count)

Category	Aug	Sep	Oct
Incorrect anti-spam/scam plan	1	4	4
Unblocked spam call	0	1	2
Anti-spam/scam app issue	2	0	3
Total	**3**	**5**	**9**

Notes

- Interactions are dominated by information-seeking queries about the feature.
- Service requests relate to enabling/disabling or configuring the feature for a subscriber number.
- Most trouble tickets relate to plan/configuration issues; smaller volumes relate to unblocked spam calls and app issues.

3 - Aug-Oct 2025 summary

Interactions (count)

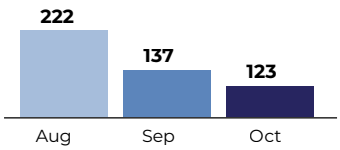

Top interaction types (count)

Category	Aug	Sep	Oct
Info seeking	205	123	110
Service request	11	11	10
Complaint	6	3	3

Top trouble tickets (count)

Category	Aug	Sep	Oct
Incorrect anti-spam/scam plan	2	1	0
Fraud impact	1	0	0
Anti-spam/scam app issue	0	1	0
Total	**3**	**2**	**0**

Notes

- Information-seeking remains the largest interaction type; volumes decline from August to October.
- Service requests are broadly stable month-to-month; complaints show limited change.
- Trouble tickets are concentrated in plan/configuration issues and reported fraud impacts (August).

Source: Company operational reporting; authors' reconstruction and formatting (de-branded).

Exhibit H

SDK activations to date and end-2025 target (selected metrics and plan)

Current Activation Base

SDK-activated customers (count)

1,186,257

As of 14 Nov 2025

Breakdown (count)

IM3: 943, 785

3: 242, 472

•Change +361k •MoM: +44%

Total Activated (in '000)

Change +361k | MoM: +44%

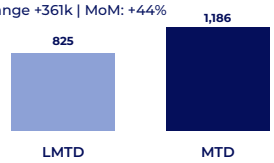

ARPU Comparison (illustrative)

Uplift: 2.3%

	Activated	Not Activated
Sep	103,982	91,248
Oct	117,615	101,128

Values shown in local currency units as reported

Net Adds Avg Run Rate (in '000)

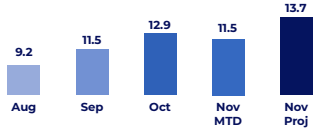

November proj. shown as separate estimates

Avg. Daily Peak Refreshing (in '000)

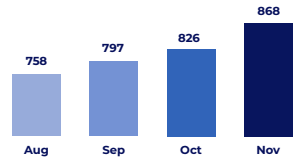

+80k from Aug to Nov (as reported)

Bridge Plan to 3Mn Activations by Dec '25

Start (mid-Nov 2025): 1.2 Mn

Bridge steps:

1. Expanded free trial campaign → **+0.3 Mn**
2. Cohort-based channel maximization → **+1.0 Mn**
3. Reward program → **+0.5 Mn**

Enablers:

- Omni-channel marketing
- Hyper-personalization

Target: 3 Mn by Dec '25

Note: Values reconstructed from provided materials; units shown as labeled (e.g., '000).
Source: Company reporting; authors' reconstruction and formatting (de-branded).

References

Isaac, J. (2025) Indonesia records world's highest number of financial scam reports: OJK, Indonesia Business Post, 3 November.

CybersecurityAsia.net (2025) Indosat Anti-Spam and Anti-Scam Feature Prevents Hundreds of Millions of Potential Digital Fraud Attempts, 12 November.

Ministry of Communication and Informatics (2019) Regulation of the Minister of Communication and Informatics Number 13 of 2019 on Telecommunications Services Licensing (English translation). Available at: https://jdih.komdigi.go.id/produk_hukum/unduhTerjemahan/id/714/t/peraturan%2Bmenteri%2Bkomunikasi%2Bdan%2Binformatika%2Bnomor%2B13%2Btahun%2B2019%2Btanggal%2B25%2Boktober%2B2019

Government of Indonesia (2024) 'Peraturan Presiden Republik Indonesia Nomor 174 Tahun 2024 tentang Kementerian Komunikasi dan Digital', enacted 5 November. (JDIH Komdigi record). https://jdih.komdigi.go.id/produk_hukum/view/id/947/t/peraturan%2Bpresiden%2Bnomor%2B174%2Btahun%2B2024

Pusat Pelaporan dan Analisis Transaksi Keuangan (PPATK) (2023) Annual Report 2023 (PDF). Available at: https://www.ppatk.go.id/backend/assets/images/publikasi/1716195339_.pdf

Otoritas Jasa Keuangan (OJK) (2025) Siaran Pers: Marak Penipuan Keuangan, OJK Bersama Pemerintah Luncurkan Kampanye Nasional Berantas Scam dan Aktivitas Keuangan Ilegal (press release), 19 August. Available at: https://ojk.go.id/id/berita-dan-kegiatan/siaran-pers/Pages/OJK-Bersama-Pemerintah-Luncurkan-Kampanye-Nasional-Berantas-Scam-dan-Aktivitas-Keuangan-Ilegal.aspx

PERBANAS – Perhimpunan Bank Nasional (2026) PERBANAS – Perhimpunan Bank Nasional (organisation page). Available at: https://perbanas.org/

Times of India. (2021) Tanla teams up with Microsoft to launch 'Wisely' secure communication solution, 20 January.

Tanla Platforms Limited. (2019) 'Tanla launches world's first blockchain enabled commercial communication stack, TRUBLOQ, at MWC19 Barcelona' (press release), 26 February; Corporate India (2023) 'Darling of marquee investors: "Wisely" is its business trump card', Corporate India, 15 June. Available at: https://www.corporateind.com/2023/june/first-issue/darling-of-marquee-investors-wisely-is-its-business-trump-card

Tanla Platforms Limited. (n.d.) 'Trubloq: Setting New Benchmarks in Spam Prevention' (case study webpage).

Tanla Platforms Limited. (2021) 'Tanla partners with Microsoft to launch blockchain-enabled CPaaS platform Wisely' (press release), 20 January; CRN India. (2021) 'Tanla partners with Microsoft to launch blockchain-enabled CPaaS platform Wisely', 20 January.

Tanla Platforms Limited. (2021) 'Tanla recognized in the 2021 Gartner® CPaaS competitive landscape' (newsroom post).

NDTV Profit. (2023) 'Tanla Platforms unveils anti-phishing product at MWC Barcelona 2023', 28 February.

London
Business
School

London Business School
Regent's Park
London NW1 4SA
United Kingdom
Tel: +44 (0)20 7000 7000
london.edu