London Business School

Tanla: Leveraging the AI-Powered Anti-Scam Telecom Shield

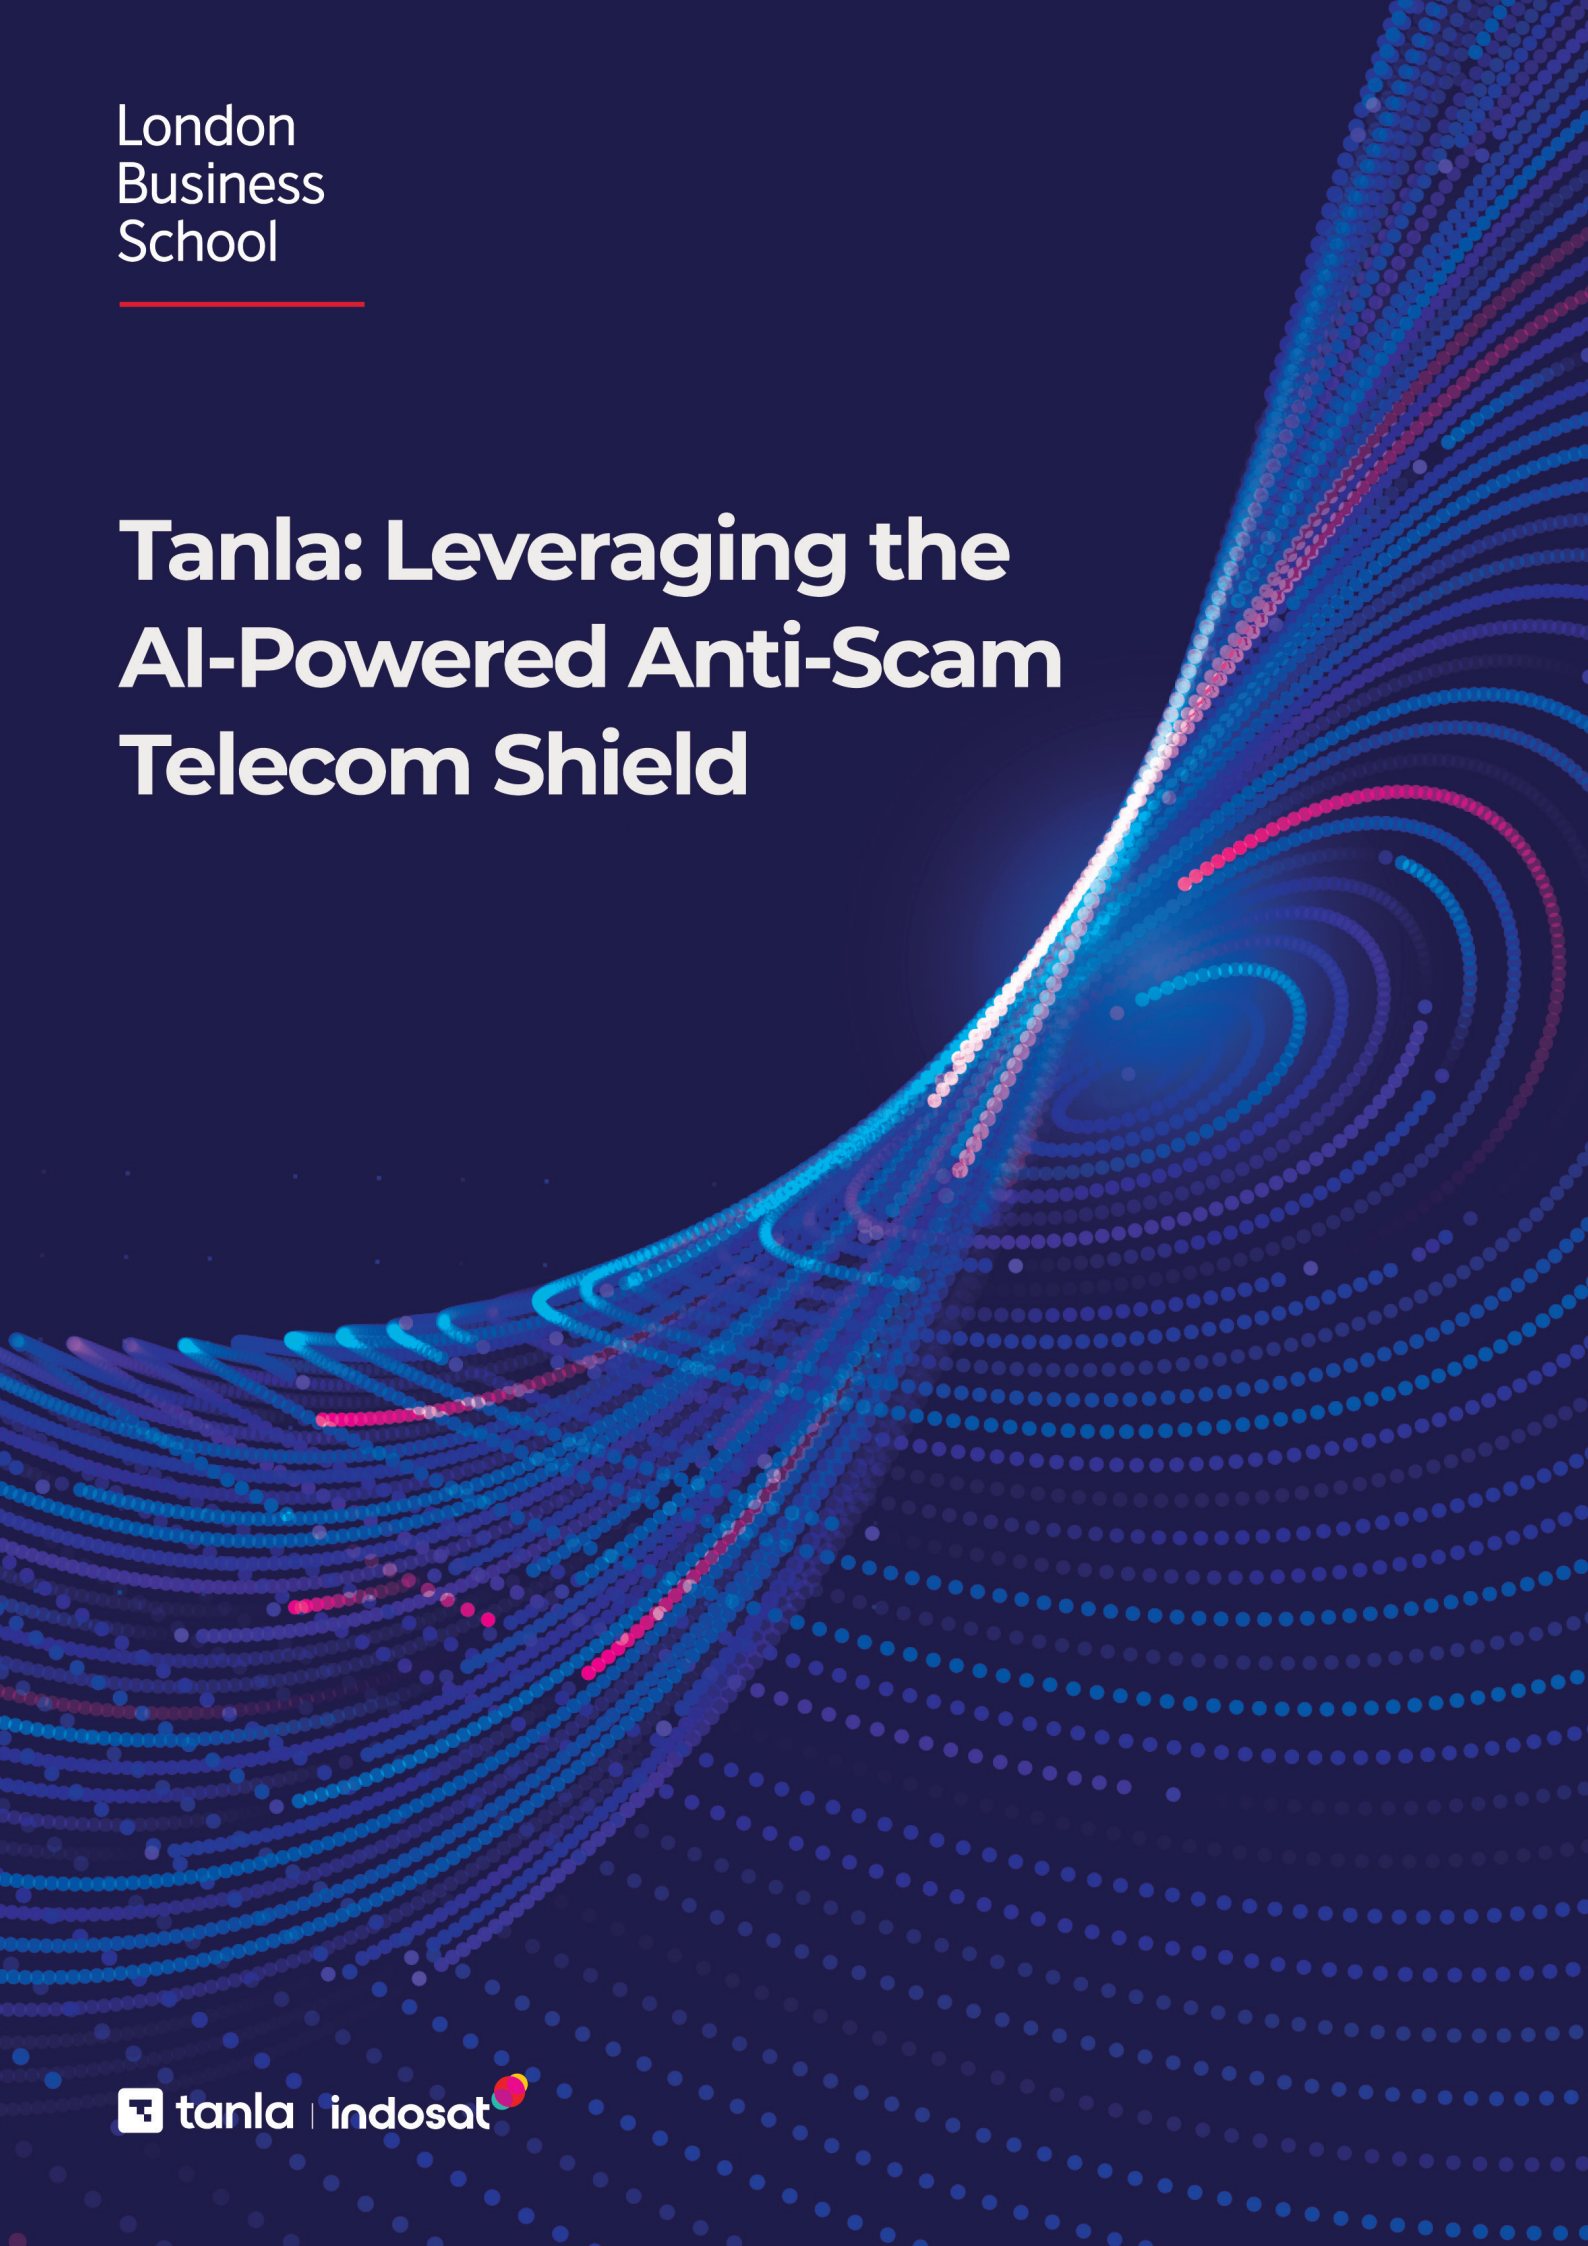

tanla | indosat

London Business School

London Business School
Case Collection

Michael G. Jacobides
M. Dalbert Ma
Shanni Elcock

CS-26-008
May 2026

Tanla: Leveraging the AI-Powered Anti-Scam Telecom Shield (B)

In February 2026, Uday Reddy, Founder Chairman and CEO of Tanla Platforms, reviewed Indosat Ooredoo Hutchison's FY 2025 earnings presentation. In its discussion of AI-powered scam and spam protection, Indosat's leadership had publicly credited Tanla's Wisely Ai platform as central to its strategy of becoming "Indonesia's most trusted telco" (Exhibit A). The presentation reported 100% protection across covered customers, two billion scam and spam attempts prevented, 95 percent positive customer feedback, and "Higher ARPU, Lower Churn."

The results had been building over the preceding weeks. On 6 February 2026, Indosat hosted an Impact Celebration event in Jakarta attended by Nezar Patria, Indonesia's Vice Minister of Communication and Digital Affairs, and Tanla board member Dr. R.S. Sharma, the former head of India's Telecom Regulatory Authority. The accompanying press release reported that the system had detected more than two billion risky calls, messages, and links since its August 2025 launch, that 95 percent of subscribers reported feeling more protected, and that Indosat estimated the system had prevented US$500 million in potential financial losses.

Indosat's Q4 2025 earnings reinforced the picture. ARPU had increased 10.5 percent quarter-on-quarter, from Rp40,000 to Rp44,000—the sharpest sequential increase in the company's post-merger history. The earnings deck attributed this to three factors: "AI-powered, hyper-personalisation initiatives," "easing industry environment," and "customer love through Marvelous Experiences." SATSPAM was not the sole driver, but it featured prominently. Revenue rose 9 percent sequentially, EBITDA grew 12 percent, and normalised net profit increased 51 percent. The following quarter Q1 2026 earnings continued the momentum with ARPU increasing to Rp45000, 15.3 percent growth YoY driving double digit revenue growth (Exhibit B).

Six months earlier, when the partnership was generating tentative indicators rather than hard results, the questions facing both companies had been pointed: Could scam protection generate measurable ROI? Would customers value something they believed should already exist? The early data from November 2025—cautiously encouraging NPS scores, tentative ARPU correlations, growing but unproven Plus+ adoption—had matured into tangible metrics. However, with that success – further questions loomed.

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.

From Tentative Indicators to Proven Metrics

The transformation in the data between November 2025 and February 2026 was substantial. The jointly defined impact framework developed by Tanla's Customer Success function with IOH's leadership provided a consistent lens to measure progress (Exhibit C). Where the picture six months earlier had described a service reaching 20 million users with hypotheses rather than proven results, the updated picture showed a platform approaching national scale with demonstrable commercial impact. By February 2026, as Anshuman Kar, the Chief Customer Officer at Tanla Platforms was reviewing the results, protection had extended across IOH's entire covered subscriber base—a significant acceleration from the 37 percent penetration reported in October 2025. The Plus+ tier, which had reached 1.2 million users by November, continued to grow as the app-based experience was integrated into Indosat's flagship myIM3 and bima+ applications. Meanwhile, the community reporting mechanism had generated over 2.5 million active app users who had collectively logged more than 124,000 reports of fraudulent numbers—creating a feedback loop that strengthened the AI's detection capabilities.

The system's technical capabilities had also matured considerably, and in several cases extended well beyond the original contract scope, which had specified AI prediction on SMS and voice calls only. New features included automatic blocking of scam SMS, detection of VoIP-based high-risk calls, pop-up notifications calibrated to risk level, and in-app summaries of suspicious activity. After discovering that VoLTE penetration was below 30 percent—meaning most users would not receive network-based caller ID alerts—Tanla independently developed the app-based experience, and subsequently extended protection to WhatsApp after operational data revealed scam activity migrating to messaging platforms where call volumes were estimated at three to four times those of traditional telco voice.

The 99 percent detection efficacy rate, maintained across both the IM3 brand (served through SATSPAM) and the Tri brand (through its parallel TRI AI feature), demonstrated that the platform could operate consistently across Indosat's dual-brand architecture inherited from the 2022 merger. Independent market research conducted by Nielsen in November 2025 provided external validation, with the SATSPAM concept scoring significantly above industry benchmarks on liking, relevance, and attractiveness (Exhibit D). The research also surfaced unmet demand: customers' most-requested features included protection against dangerous links and URLs, spam protection on WhatsApp, and defence against harmful apps—capabilities extending beyond the channels the system had originally covered.

Government Validation and Regulatory Momentum

Perhaps the most consequential development was the explicit government endorsement. The Vice Minister's attendance at the Impact Celebration and his public statement encouraging "industry players to implement similar measures" represented a shift from the regulatory ambiguity that had characterised earlier discussions. Where regulators had previously been described as a potential forcing mechanism whose timeline and approach remained uncertain, the February 2026 picture showed active government validation of the operator-led, technology-enabled approach. This validation carried strategic implications for both companies. For Indosat, it strengthened the case that SATSPAM was aligned with national digital policy—providing a degree of insulation against any future regulatory moves that might commoditise protection by mandating narrow technical standards. For Tanla, the Vice Minister's statement was, implicitly, an endorsement of the technology partnership model itself. When a government official encourages other industry players to adopt similar measures, the technology provider behind those measures gains a powerful reference.

Yet the endorsement also sharpened a tension that had been identified earlier. If the government was actively encouraging all operators to adopt similar protections, the differentiation window for Indosat was potentially narrowing. A technology that was initially a competitive advantage risked becoming a table-stakes requirement—

precisely the dynamic that the ecosystem discussions had anticipated.

Tanla's Strategic Crossroads

For Reddy, reviewing the options from Tanla's Hyderabad headquarters, the Indonesia results created a different kind of strategic moment. The proof point was now in hand—and it was stronger than most had anticipated. A major telco had publicly credited Tanla's platform in its earnings presentation. A sovereign government had endorsed the approach. The numbers—two billion threats detected, US$500 million in prevented losses, 10.5 percent ARPU growth—were the kind of metrics that commanded attention across the telecommunications industry. While direct attribution of outcomes is challenging in a complex industry such as telecommunications, a conservative, assumption-driven model pointed to a directional ROI of ~15–20x, with an estimated 20–40% contribution to ARPU uplift and churn reduction.

The question was what to do with this proof point. Four distinct strategic paths had emerged in Tanla's leadership discussions, each with different implications for how the company would grow, compete, and create value.

Option 1: Replicate the Playbook Across Operators

The most immediate path was to leverage the Indonesia success to sign similar deployments with telecommunications operators in other markets. Southeast Asia alone presented multiple candidates: operators across the Philippines, Vietnam, Thailand, and Malaysia faced comparable fraud challenges with similar structural characteristics—large prepaid subscriber bases, fragmented regulatory environments, and growing digital payment ecosystems. India, where Tanla already had deep relationships, offered additional opportunities, as did markets in Africa and the Middle East. Under this approach, each operator would represent a separate deployment, and Tanla would capture value through per-operator licensing arrangements. The Indonesia proof point would serve as the anchor reference case.

This path had clear advantages. It leveraged Tanla's existing go-to-market capabilities—the company already knew how to sell to telco procurement teams. It was commercially straightforward, with revenues scaling with each new operator signed. And it preserved competitive dynamics between operators, with each deployment offering Tanla fresh data and learning that strengthened the platform.But it also carried limitations. Per-operator deployments meant that scam intelligence remained siloed within each client. A fraud pattern detected on Indosat's network would not automatically benefit a deployment in the Philippines—unless Tanla built mechanisms to share anonymised insights across clients. Scaling operator-by-operator was also slow relative to the speed at which fraud evolved and migrated across borders. And the sales cycle for each new operator could be lengthy, particularly in markets where fraud had not yet reached the severity that had motivated Indosat's urgency.

Option 2: Build Toward an Ecosystem Infrastructure Position

Rather than treating each operator as an isolated client, Tanla could position Wisely Ai as the shared intelligence infrastructure layer that the Threat Intelligence Exchange concept had envisioned — orchestrating a multi-party network through which operators, financial institutions, regulators, and digital platforms exchanged threat intelligence in real time. This vision had animated earlier ecosystem discussions, but had previously lacked proof that the underlying technology could deliver at scale. The Indonesia results now provided that proof, though the question remained whether a single-operator deployment could be extended to a multi-party context. If Tanla became the backbone of regional anti-fraud coordination, its strategic position would be considerably more defensible than a series of bilateral operator contracts. Network effects would compound as each new

participant improved detection for all others, creating switching costs that grew with the network — shifting Tanla's role from technology vendor to platform orchestrator.

The risks were equally substantial. Multi-party coordination required governance mechanisms that did not yet exist; competing operators would need to share data through a neutral intermediary; and regulators across jurisdictions held different standards on cross-border data sharing. The commercial model for shared infrastructure remained unresolved, and pursuing the ecosystem play would require Tanla to invest in policy and governance capabilities far removed from its engineering strengths.

Option 3: Deepen the Indonesia Engagement

A third path focused not on geographic expansion but on deepening the value captured within the existing partnership. The Indonesia deployment had proven the anti-scam use case, but the platform's capabilities could extend into adjacent domains: identity verification, enterprise sender authentication, trusted communication channels for government services, and digital commerce protection. Indosat's own strategic roadmap—its "AI TechCo" ambition and investments in NeoCloud, fibre infrastructure, and 5G expansion across 24 Indonesian cities—suggested appetite for technology partnerships that extended well beyond scam protection.

This option had the advantage of building on a proven relationship and a live operating platform. It avoided the risks of entering unfamiliar markets or coordinating multi-party ecosystems. But it concentrated Tanla's exposure in a single client and a single market, and it required Indosat to continue prioritising the partnership—a dependency that carried its own risks.

Option 4: Build a Trusted Sender Marketplace

A fourth strategic path had emerged from an unexpected consequence of the platform's success. As SATSPAM's detection capabilities tightened, legitimate commercial messages were increasingly caught in the crossfire. Banks sending transaction alerts, e-commerce platforms confirming orders, government agencies distributing public health notifications—all relied on the same SMS and voice channels that scammers exploited. Enterprises whose messages were being misidentified had a direct interest in being recognised as trustworthy by the system. If Tanla could build a verification and authentication layer through which legitimate senders established their identity and earned trusted status, it would address the false-positive problem while creating a new revenue stream.

The concept had precedents in adjacent industries. Email spam in the early 2000s gave rise to authentication protocols (SPF, DKIM, DMARC) that allowed senders to verify their identity. SSL/TLS certificates created visible trust signals for websites, with certificate authorities charging businesses for the credential. Google's verified business profiles demonstrated how authentication could be commercialised at scale. For Tanla, a trusted sender marketplace could range from simple verification that whitelisted enterprise messages through to tiered access offering delivery analytics, priority routing, or branded caller ID.

However, the model carried governance risks. If Tanla simultaneously operated the detection system that blocked messages and the verification service that allowed them through, it could face accusations of creating a problem and selling the solution. Transparent authentication criteria, separation between detection and verification operations, and regulatory involvement in standard-setting could mitigate these risks—but getting it wrong could damage relationships with both operators and regulators.

An Industry Watching

By February 2026, the Tanla–Indosat collaboration represented one of the more fully documented cases of AI-powered scam protection delivering measurable commercial outcomes at telecommunications scale. The technology had demonstrated strong efficacy in a large-scale deployment, establishing a credible blueprint for broader adoption. Early evidence suggested that its core architecture, anchored in deep network integration and AI-driven detection could be adapted across regulatory environments and scaled beyond a single operator relationship. As Tanla evaluated the next phase of growth, it was not constrained by a lack of direction but evaluating multiple viable pathways, each building on a now-proven foundation. The choices made in the coming months would shape not only Tanla's trajectory but potentially the broader industry's approach to consumer protection.

FIGURES

Exhibit A

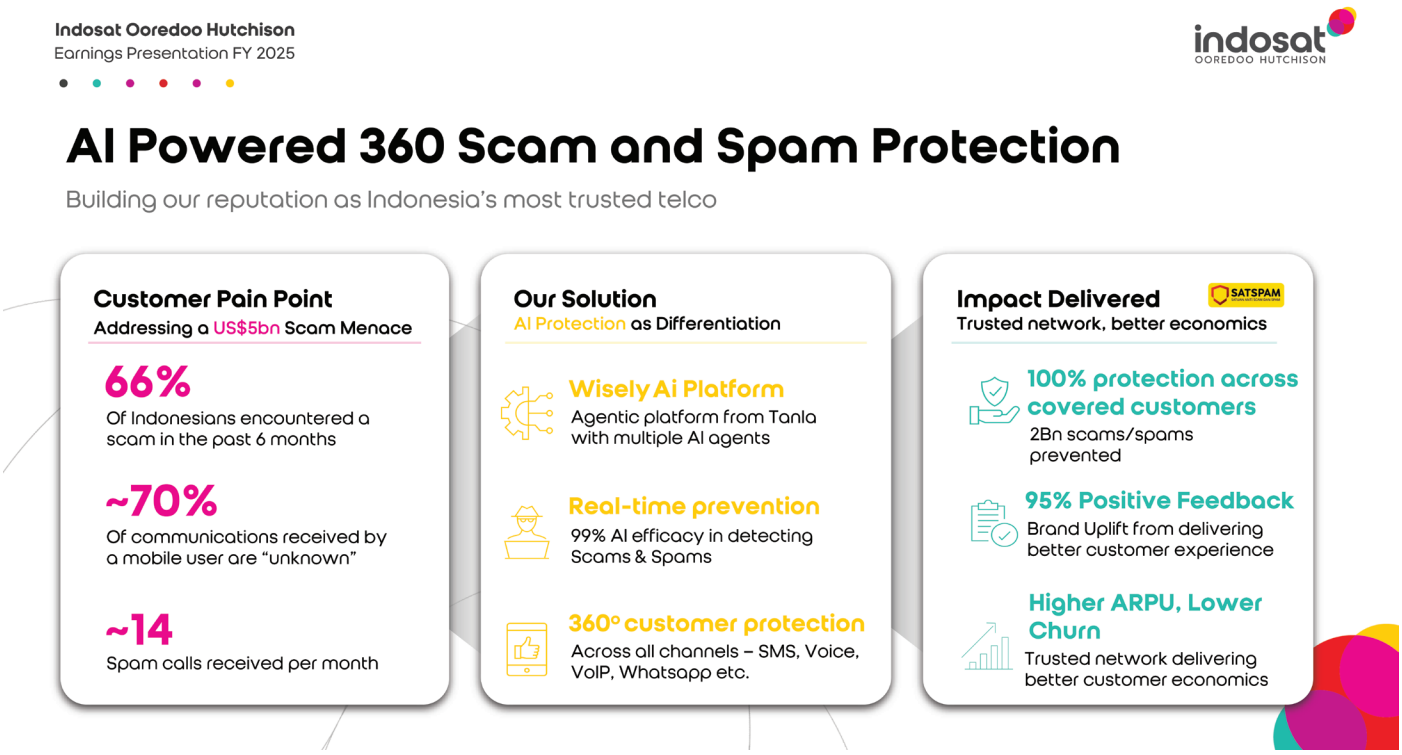

Exhibit B

IOH's ARPU growth journey continues

IOH ARPU (Rp'000)

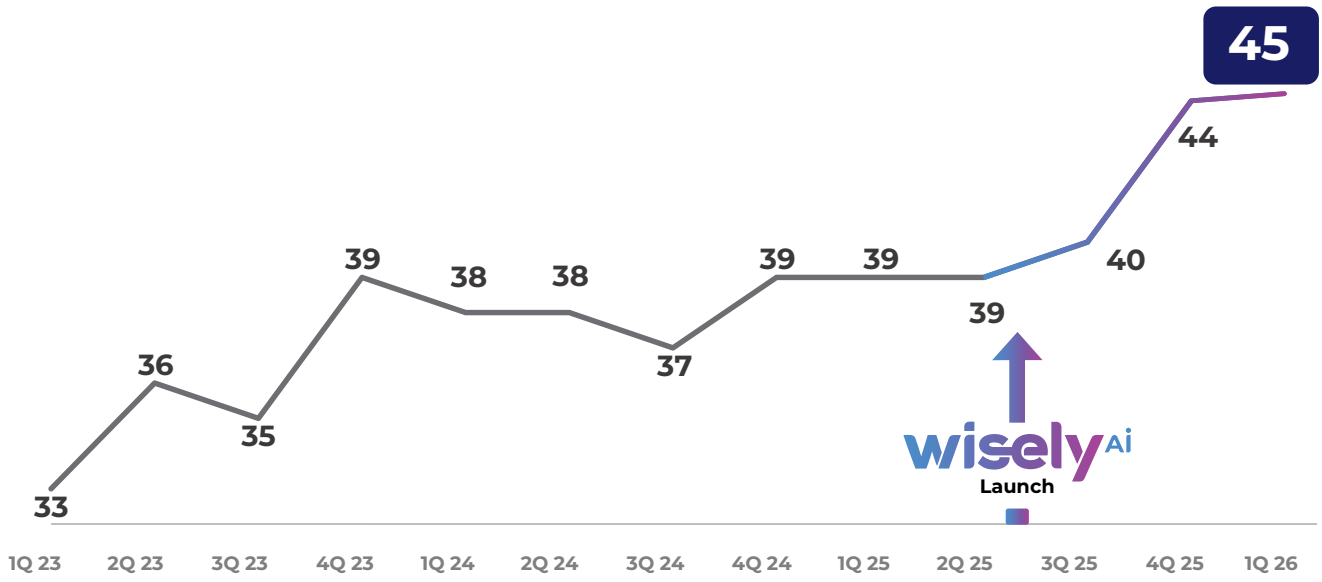

Exhibit C

Customer Success Impact Measurement Framework

Platform KPIs (Input Metrics)

Clean Network

Known Communications

- Scams Blocked (Mn, % Identified scams)
- Spams Alerted (Mn, % Identified spams)
- Traffic with verified sender (Mn, % total traffic)

AI Efficacy

- Identification accuracy (%)
- Time to identify scam (Mins)

User Engagement

- Mobile App MAU (% of subscriber base)
- User feedback reports (#)

Impact (Output Metrics)

User impact

- Users protected (% eligible subscribers)
- App Positive feedback (%)
- Financial loss prevented ($ Mn)

Telco impact

- Telco CSAT (%)
- End user CSAT (%)

Enterprise impact

- Call answer rate for Verified calls (%)
- Enterprise CSAT (%)

Outcome

X%
ARPU increase

Y%
Churn Reduction

Exhibit D

DATA IN %

Is the SATSPAM concept gaining tracking among our target market?

SATSPAM shows strong traction with scores **above norms**, driven by safety and easier to recognize suspicious calls / messages.. Expanding protection to chat platforms (e.g WA, Telegram) and protect harmful apps is the key opportunity for future growth.

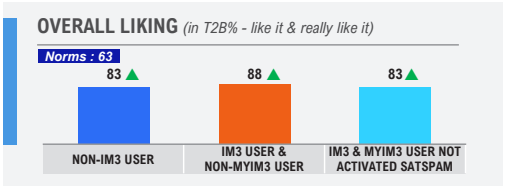

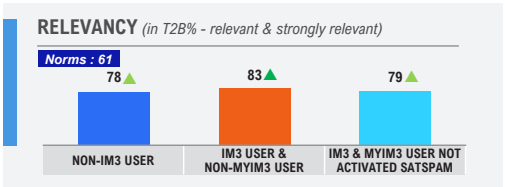

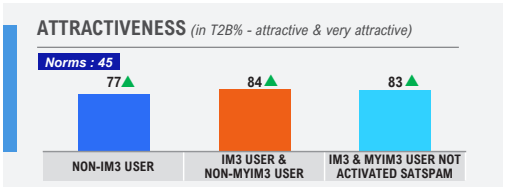

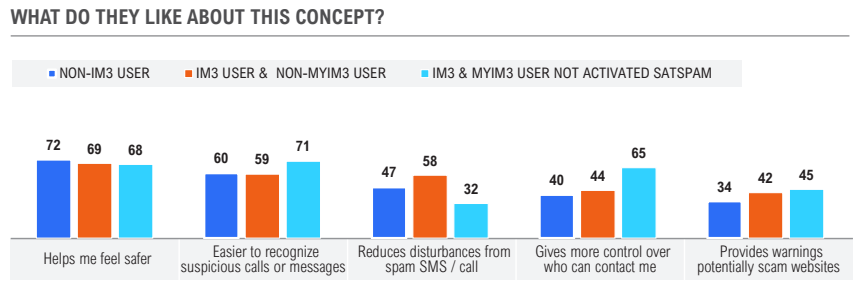

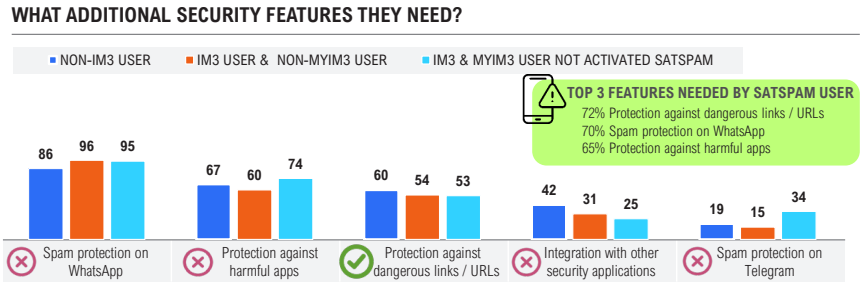

NIQ

Norms on likability based on 17441 sample size from NIQ Telco concept test from 2015 - 2025 database
Norms on relevance based on 15726 sample size from NIQ Telco concept test from 2015 - 2025 database
Norms on intention based on 21966 sample size from NIQ Telco concept test from 2015 - 2025 database

▲▲ Value is significantly higher than NORMS (95%,90% confidence level)
▼▼ Value is significantly lower than NORMS (95%,90% confidence level)

Available on SATSPAM

Not available on SATSPAM

10

References

https://www.telecomreviewasia.com/news/industry-news/28316-indosat-detects-over-2-billion-spam-and-scam-attempts-expands-ai-driven-protection/

https://markets.ft.com/data/announce/detail?dockey=600-202602070345BIZWIRE_USPRX____20260206_BW004579-1

https://ioh.co.id/portal/en/iohcorppressreleasedetail/ioh-safer-digital-ecosystem-feb-2026?_id=10015470

https://ioh.co.id/portal/en/ioh-corporate-governance-detail/presentasi-perusahaan-tw4-2025

London
Business
School

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