

FREE SAMPLE

And detail supporting this deck is described in the full PDF report

STRATEGIC BRIEFING

AI Integration & Cybersecurity Vendor Fit in Southeast Asia

Q3 2026

An analytical assessment of thirteen evaluated vendors across five categories, scored on operational deployability and AI readiness across six Southeast Asian markets.



PREPARED BY LOGIC LOCK GROUP

Source: AI and Cybersecurity in Southeast Asia — Vendor Intelligence Edition (Q3 2026)


The Core Strategic Trade-off

1 in 13

evaluated vendors currently clear both the SEA Fit and AI Readiness bars.

(see full report)

Palo Alto Prisma SASE (3.83 Fit / 4.13 AI) and Fortinet Unified SASE (4.55 Fit / 3.70 AI) sit near opposite corners — a blended score would make them look like substitutes. They are close to opposite purchases.



Illustrative — full 13-vendor matrix on Slide 4

- AI capability and Southeast Asia fitness are weakly — at times inversely — correlated across every category researched.
- “AI-native” is the most overloaded, inflated phrase in the 2026 enterprise security market. This methodology decomposes vendor architectures into three distinct vectors (Slide 8) rather than scoring the label.
- A single composite ranking is a dangerous metric for a Southeast Asian buyer: AI readiness and operational deployability are frequently inversely correlated, and a blended number averages away the exact trade-off that must be navigated.
- The platforms with the most advanced agentic engines are routinely the least deployable on sovereign soil; vendors with deep localized footprints rarely lead the frontier-model race.

If the agentic SOC works, why are incident-response retainers growing?

52% · 89 sec


of MDR cases resolved end-to-end by AI, at this average alert-to-response time — a case, not an incident; DFIR is still needed for the rest

Because a case is not an incident. Machine autonomy is real at the routine-case tier — and absent at the major-incident tier.

- Best vendor evidence: 52% of MDR cases resolved end-to-end by AI at an 89-second average (Sophos, Grade B, measured — not a ceiling).
- Every autonomy claim is graded on a severity tier (1: triage/routine — 2: complex multi-stage, AI-assisted — 3: major incident, human-led) and an evidence class (A–D).
- Nothing in the market earns Grade A. The closest: SentinelOne's MITRE Managed Services evaluation — 100% detection, best signal-to-noise ratio, independently observed.
- The agentic SOC automates the tier-one/two work SEA cannot staff — it does not automate incident response. Human DFIR remains a permanent, deliberately procured layer.


The Master Market Quadrant

SEA Fit × AI Readiness — all thirteen evaluated vendors



● Leader (blended ≥ 4.00) ● Contender (blended ≥ 3.50)

Endpoint / Managed Security (EDR/MDR/XDR)



VENDOR	SEA FIT	AI RDY	TIER
CrowdStrike			See full report
Microsoft Defender XDR			See full report
Sophos MDR	3.65	3.83	CONTENDER
Trend Micro (TrendAI)			See full report

- (see full report)
- CrowdStrike and Trend Micro sit at opposite ends of this category’s sovereignty spectrum — (see full report).
- (see full report)
- Every buyer in this category should ask the two-question residency test: where is my data stored, and where is it processed.

The Singapore Playbook

The full-option market — and the market that writes the region's homework

1

INFRASTRUCTURE STATUS

- The only covered market with in-country options in every evaluated category: Zscaler's dual service edges, Cortex XSIAM's dedicated hosting region, SentinelOne's Singapore data centre (since 2022), Azure's established region.
- Everything reaches Singapore first — vendors treat it as the region's landing zone.

2

REGULATORY PROGRESS / GUARDRAIL

- MAS Technology Risk Management sets the financial-sector bar; IMDA and CSA shape technology and AI policy under the Model AI Governance Framework.
- Binding constraint here is diligence quality, not option scarcity — sophisticated but navigable.

3

RECOMMENDED ARCHITECTURE

- (see full report)



PITFALL / KEY RISKS

(see full report)

The RFP Verification Framework

Five questions to copy directly into your next vendor RFP — engineered to force disclosure on evidence grade and offshore processing location.

1 Where is our data stored, and — separately — where is it processed?

Storage is not processing. Require both answers, in writing, for every covered market.

2 What evidence grade (A-D) supports each autonomy claim, and is it independently auditable?

(see full report)

3 What is your exact deployment model in each of our markets — cloud-only, vendor-operated in-country, or customer-hosted sovereign?

Map the answer against the three-tier taxonomy; “undisclosed” is itself information.

4 How do you discover, credential, constrain, audit, and offboard AI agents — including your own?

(see full report)


5 What does a month of the intended automation cost at our actual alert volume, under your metering model?

(see full report)

Appendix A — Analysis Methodology


SEA FIT (0–5.0)

Operational deployability across the six covered markets.



AI READINESS (0–5.0)

The three-way decomposition every vendor is scored against.



EVIDENCE GRADES CAP EVERY SCORE



No claim in this study earns Grade A for agentic autonomy. Tiers: Leader at blended ≥ 4.00 ; Contender from 3.50. No Niche tier — weak vendors were removed at shortlisting, not omitted by the market.

What's Inside the Full 25-Slide Deck

1	Section 1 — Macro Governance & the Regional AI Reality	1. Cover · 2. The Core Strategic Trade-off · 3. Regulatory Compliance Pressures · 5–8. Four Flagship Findings
2	Section 2 — Deep-Dive Category Snapshots	9. The Master Market Quadrant (all 13 vendors) · 10. EDR/MDR/XDR · 11. SASE/SSE · 12. SIEM/SOAR · 13. Exposure Management · 14. ITDR
3	Section 3 — Market-by-Market Architectural Playbooks	15. Singapore · 16. Malaysia · 17. Indonesia · 18. Thailand · 19. Vietnam · 20. Philippines
4	Section 4 — Cross-Cutting Themes	21. Five Patterns That Crossed Every Category
5	Section 5 — The CISO Decision Playbooks	22. The Reader Decision Toolkit · 23. Stack Over Strategy: The Combination Pattern · 24. The RFP Verification Framework · 25. Closing

Vendor selection in Southeast Asia is becoming an exercise in assembling stacks, not choosing champions.

Prepared by Logic Lock Group

Source: AI and Cybersecurity in Southeast Asia — Vendor Intelligence Edition (Q3 2026)

Thirteen vendors evaluated on two independent axes across five categories and six markets — Singapore, Malaysia, Indonesia, Thailand, Vietnam, and the Philippines.