

LLG INSIGHTS · LOGIC LOCK GROUP PTE. LTD.

AI and Cybersecurity in Southeast Asia Q3 2026

Vendor Intelligence

FREE SAMPLE

1 in 13 vendors clear both scoring bars

AI capability and Southeast Asia deployability are weakly — at times inversely — correlated.



18 Vendors · **6** Markets · **5** Categories

Logic Lock Group

AI and Cybersecurity in Southeast Asia Q3 2026 — Vendor Intelligence

LLG Insights · Logic Lock Group Pte. Ltd. — August 2026

About this report

This is the second edition in the *AI and Cybersecurity in Southeast Asia* series. Where the Q2 2026 edition assessed the region's threat landscape, market economics, and strategic imperatives, this edition assesses eighteen vendor entries — thirteen evaluated across five categories on two independent axes, two spotlight profiles, and three regional operators — against the specific conditions of six markets: Singapore, Indonesia, Thailand, Vietnam, Malaysia, and the Philippines.

Lineage. This edition incorporates and extends the regulatory compliance analysis previously published as *the SEA Cybersecurity Vendor Intelligence — Regulatory Compliance Report H1 2026 (VIRC)*. That consolidation reflects how the research itself developed: the regulatory question and the vendor question proved inseparable, and readers of the H1 2026 report will find its jurisdiction-by-jurisdiction analysis continued, updated, and integrated into this edition's market guidance rather than published separately.

How the report is organised. Part 1 summarises. Part 3 argues the four flagship questions in full — references throughout the report in the form “Flagship 3.1” point to those four questions. Parts 5 through 8 carry the vendor and market depth. Part 10 converts the analysis into decisions. Part 2 explains the scoring machinery in ten minutes and is the recommended starting point for readers who begin with the scorecards.

Independence. No vendor assessed in this report sponsored, reviewed, or previewed its contents. Vendor-sourced claims are marked as such throughout, every material claim carries an evidence grade, and the claims register (Appendix A) lists the load-bearing facts with their sources and verification status.

Research cut-off. July 20, 2026. Product availability, hosting locations, preview statuses, analyst positions, and regulatory citations reflect verification as of this date; the claims register (Appendix A) records per-claim status and flags items for re-verification at publication.

Table of Contents

Table of Contents 3

PART 1 — EXECUTIVE SUMMARY 4

PART 2 — HOW TO READ THIS REPORT 6

PART 3 — THE FOUR FLAGSHIP QUESTIONS 8

PART 4 — METHODOLOGY 12

PART 5 — VENDOR EVALUATION BY CATEGORY 16

PART 6 — SPOTLIGHT PROFILES.....20

PART 7 — EMERGING REGIONAL VENDORS.....23

PART 8 — MARKET-BY-MARKET GUIDANCE 24

PART 9 — CROSS-CUTTING THEMES.....28

PART 10 — READER DECISION TOOLKIT.....29

Appendix A — Claims Register 30

Appendix B — Sources 31



PART 1 — EXECUTIVE SUMMARY

Southeast Asia is adopting AI-era security technology faster than the infrastructure, governance, and delivery layers beneath it are maturing — and the eighteen vendor entries examined in this report land on that gap in ways no global analyst publication captures. Of those, this report scores thirteen with full two-axis across five categories on two independent axes, profiles two spotlight vendors and three regional operators, and answers four questions that come before any vendor selection. It is the second edition in [the AI and Cybersecurity in Southeast Asia series](#), and it incorporates and extends the regulatory compliance analysis previously published as the VIRC H1 2026 report — the regulatory question and the vendor questions covered in the VIRC report are reflected in this edition.

The finding that organises everything else is that AI capability and Southeast Asia fitness are weakly — at times inversely — correlated. The vendors with the most genuinely AI-native platforms are frequently the least deployable for the region’s regulated buyers, and the vendors most deployable on sovereign soil are rarely the AI leaders. One vendor in thirteen currently clears both bars. This is why the report scores two axes rather than one: a composite number would average away the exact trade-off a Southeast Asian buyer must navigate.




Figure 1. The SEA Fit × AI Readiness quadrant — all thirteen evaluated vendors. Gold = Leader, blue = Contender.

Four flagship findings, in brief.

Where your security actually lives. “

[..... \(see full report\)](#)

Whether what you already have is good enough. “

[..... \(see full report\)](#)

Why incident-response retainers grow while the agentic SOC works. “

[..... \(see full report\)](#)

The ungoverned AI problem, extended. “

[..... \(see full report\)](#)

Category verdicts, in one line each. In EDR/MDR/XDR, four consumption models map to four buyer situations: CrowdStrike for capability through the managed layer, Trend Micro for sovereignty, Microsoft for the conditional incumbent case, Sophos for the SOC-less mid-market. In SASE, Fortinet’s sovereign deployment model earns the report’s highest SEA Fit score while Zscaler and Palo Alto lead the cloud market it cannot yet serve everywhere. In exposure management, Tenable automates the workflow while XM Cyber’s choke-point analytics answer the talent gap by design. In SIEM, [..... \(see full report\)](#). In ITDR, the category has bifurcated into control-plane and platform models — complementary, not competing — under a residency ceiling no identity vendor yet lifts.

The recommendation pattern the region should name and use: combination. A global AI-native platform, delivered through a regional MSSP with in-country presence, plus the permanent human DFIR layer, assembles from two or three vendors what almost none offers alone — and the pattern already operates at award-winning standard in this region. Vendor selection in Southeast Asia is becoming an exercise in assembling stacks, not choosing champions, and this report’s market-by-market guidance is written accordingly — including a deliberately different playbook for the Philippines, the study’s most underserved market, where regulatory room partially offsets the infrastructure gap.

Who this report is for. The CISO shortlisting vendors or defending an incumbent; the CTO sequencing a regional security architecture; the risk and compliance officer who must answer a regulator’s data-flow question with both locations, not one. Part 10 routes each persona through the report; Part 2, next, explains the scoring machinery in ten minutes.

PART 2 — HOW TO READ THIS REPORT

Two scores, one chart, no composite ranking. Every fully evaluated vendor carries a SEA Fit score (operational deployability in the six covered markets) and an AI Readiness score (the substance behind the AI claims), plotted on the quadrant that appears in Part 1 and recurs per category. The two are deliberately never merged into a headline ranking, because the report’s central finding lives in their divergence: Fortinet (4.55/3.70) and SentinelOne (3.80/4.55) would blend to nearly identical composites while being close to opposite purchases, though the two sit in different categories (SASE and ITDR). Within the same category, the trade-off is just as sharp: Palo Alto Prisma SASE (3.83 Fit / 4.13 AI) and Fortinet Unified SASE (4.55 Fit / 3.70 AI) sit near opposite corners of SASE alone — a blended score would make them look like substitutes when they are close to opposite purchases. No composite score means exactly this: vendor-to-vendor comparisons only hold within a category, never across one. A blended figure exists only to assign category tiers — Leader at 4.00 and above, Contender from 3.50 — and there is no Niche tier because weak vendors were removed at shortlisting, not because the market lacks them.

Applicability bands keep regional scores honest. Regional vendors’ scores apply inside their stated market footprints and are printed with per-market strips. Ensign’s 4.35 is real in Jakarta and not in Manila; the strip says so before the number can mislead.

Evidence grades appear wherever claims do. Grade A is independently observed; B is a vendor-published operational average; C is a best-case “up to” ceiling; D is preview or roadmap. Scores are capped by the grade of the evidence beneath them, which is why a vendor publishing measured averages can outscore one publishing bigger adjectives. Where a vendor does not disclose something material — hosting locations, processing geography, deployment models — the scorecard prints “undisclosed,” and readers should treat the disclosure gap as information.

Statuses. Evaluated vendors receive full two-axis scores. Spotlight vendors (Google, Sangfor) receive argued profiles with named revisit triggers instead of scores, for stated editorial reasons. Regional vendors receive SEA Fit scores plus qualitative AI-posture notes. All eighteen received the same evidence standard.

Scorecard at a glance

Every scored and profiled entity in one table — both axes and tier for the thirteen evaluated vendors, plus the two spotlight profiles and three regional operators. Where each vendor actually applies by market is summarised separately in Part 8.7; full detail sits in each profile.

Vendor	Category	SEA Fit	AI Readiness	Tier
CrowdStrike	EDR/MDR/XDR	■■■	■■■	See full report
Trend Micro (TrendAI)	EDR/MDR/XDR	■■■	■■■	See full report
Microsoft Defender XDR	EDR/MDR/XDR	■■■	■■■	See full report
Sophos MDR	EDR/MDR/XDR	3.65	3.83	Contender
Fortinet Unified SASE	SASE	4.55	3.70	Leader

Palo Alto Prisma SASE	SASE	3.83	4.13	Contender
Zscaler	SASE	■ ■	■ ■	See full report
Tenable One	Exposure Mgmt	■ ■	■ ■	See full report
XM Cyber	Exposure Mgmt	■ ■	■ ■	See full report
Cortex XSIAM	SIEM/SOAR	■ ■	■ ■	See full report
Microsoft Sentinel	SIEM/SOAR	■ ■	■ ■	See full report
SentinelOne	Identity/ITDR	3.80	4.55	Leader
Okta	Identity/ITDR	■ ■	■ ■	See full report
Google (Unified Security)	Spotlight	—	—	Spotlight (argued, not scored)
Sangfor	Spotlight	—	—	Spotlight (argued, not scored)
Ensign InfoSecurity	Regional operator	4.35	Qualitative notes only	Regional
Group-IB	Regional operator	■ ■	■ ■	See full report
LGMS Berhad	Regional operator	■ ■	■ ■	See full report

If you are not (yet) a regulated buyer. The scoring in this report leans on residency and regulatory mechanics because that is where vendor differences are hardest to reverse after purchase — but the audience is not regulated enterprises alone, and the same evidence reads a second way. For non-regulated organisations, large and small, the AI Readiness column and each profile’s AI decomposition describe what a lean security team actually receives from a platform; the delivery-layer findings (P4) identify which vendors can be consumed through a regional MSSP rather than staffed in-house; and the cost contexts flag where consumption metering will shape budgets regardless of sector.

Two reasons the sovereignty half of this report still belongs on a non-regulated buyer’s desk. First, the regulatory perimeter in this region moves outward, not inward: Vietnam’s PDPL and AI statute, Indonesia’s PDP Law, and Malaysia’s expanding NCII designations each pulled organisations into scope that sat outside it a budget cycle earlier — and the cheapest time to inherit a residency posture is before it is mandatory. Second, residency obligations increasingly arrive through the supply chain rather than the statute: regulated customers and government contracts pass their requirements downstream, and a supplier that cannot answer where the data lives can be quietly disqualified from regulated revenue it never knew it was competing for. Non-regulated buyers should read the P2 columns not as compliance homework but as optionality: what would switching cost if scope arrives mid-contract?

The honest limits. Scores are first-edition calibrations against publicly verifiable evidence gathered through mid-2026; several vendors’ scores would move on disclosure rather than on product change, and the profiles say which. Vendor-sourced claims are marked as such throughout. Where this report is wrong, the intent is for that to surface in public, specifically and correctly, — which is more than can be said for adjectives.

PART 3 — THE FOUR FLAGSHIP QUESTIONS

The vendor profiles in this report answer the question every buyer eventually asks: which product should I choose? But the research surfaced four questions that come before vendor selection, and that no vendor's marketing will answer for you. Each is argued here in full; the category chapters and market guidance that follow supply the per-vendor and per-market detail. A reader who stops at the end of this Part will have this report's complete analytical contribution. A reader who acts on it will want the parts that follow.

3.1 Where does your security actually live?

Ask a security vendor where their product runs and you will receive an answer about capability. Ask where your data is stored and you will usually receive an answer about a region list. Ask where your data is *processed* — analysed, correlated, enriched, acted upon — and the conversation frequently changes character. This report treats that third question as the first question, because in Southeast Asia it determines what you are legally and practically able to deploy.

The infrastructure facts are stark and, to our knowledge, have not previously been assembled in one place. Across every category this report evaluated, there is no in-country, cloud-native security service available today in Thailand or Vietnam from any evaluated vendor; the Philippines exited this finding late in the research window, when Fortinet's FortiSASE inventory added an in-country Manila data centre — the market's first evaluated-vendor service edge. Zscaler's Zero Trust Exchange operates full in-country service edges in Singapore, Kuala Lumpur, and Jakarta — a new Kuala Lumpur facility opened in January 2026 (complementing existing local infrastructure) with explicit alignment to Malaysia's National Critical Information Infrastructure agenda — but inspects Thai, Vietnamese, and Philippine traffic from abroad. Palo Alto's Cortex XSIAM offers exactly one Southeast Asian hosting region, Singapore. Okta, the identity control plane for thousands of regional enterprises, operates cells in the United States, Europe, Japan, Australia, Canada, and India — and none in Southeast Asia; a Jakarta bank running Okta has its identity infrastructure in, typically, Tokyo or Sydney. Sophos discloses no Southeast Asian hosting at all.

Three deployment models emerged from the research, and they form the taxonomy this report scores against. The first is cloud-only: the vendor's cloud, wherever it is, is the only place the product runs. The second is the vendor-operated in-country region: the vendor builds local infrastructure and operates its service from inside the customer's jurisdiction — the model Zscaler runs in three SEA markets, Palo Alto in one, and CrowdStrike now operates in Saudi Arabia, India, and the UAE under its Global Data Sovereignty initiative, though not yet in any Southeast Asian market. The third is customer- or partner-hosted sovereign deployment: the full product runs in a data centre the customer or a local partner controls. Only two evaluated vendors offer this today “...

..... (see full report)

The analyst verdict. Sovereignty in Southeast Asia is achievable today, but only through deliberate architecture: sovereign-deployable platforms where they exist, vendor-operated in-country regions where they reach, and contractual clarity about processing location everywhere else. Identity is the unresolved exception. As for local hyperscaler presence, Thailand's map is moving; Vietnam's and the Philippines' are not. And every buyer should now ask every vendor two questions instead of one: where is my data stored, and where is it processed.

3.2 Is what I already have good enough?

Most Southeast Asian enterprises will not select a security vendor this year. They already have one — usually Microsoft, arriving bundled with the productivity estate — and their real question is whether the incumbent is good enough for the AI era, or whether the gap between what they own and what the market now offers justifies the cost and disruption of change. This is the most common question a CISO in this region actually faces, and it deserves a structured answer rather than a vendor pitch.

The economics of the question changed on January 1, 2026, when Microsoft included Security Copilot in Microsoft 365 E5. For the large population of regional enterprises already licensed on E5, agentic security assistance now arrives without a procurement cycle: the Phishing Triage Agent is generally available and directly relevant to a region whose threat profile is dominated by phishing and business email compromise; Copilot Chat runs inside Defender; and a portfolio of triage and analyst agents is in preview. The Sentinel data lake, now generally available, restructures the retention economics that made traditional SIEM painful. On paper, the incumbent has never been more defensible.

The honest caveats are three, and they are structural rather than temporary. First, the portfolio is preview-heavy: several of the most consequential agents are not yet generally available, and this report's evidence-grade discipline scores shipped capability, not roadmap — by that standard Microsoft's agent portfolio carries the highest preview-to-GA ratio of any evaluated vendor, a score that will rise as previews graduate but has not yet. Second, ...

..... [\(see full report\)](#)

The analyst verdict.

..... [\(see full report\)](#)

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

Every vendor in this report now claims an agentic SOC. Sophos markets the world's largest, reporting that 52 percent of its MDR cases are resolved end to end by AI at an average of 89 seconds from alert to automated response. Palo Alto cites up to 98 percent reduction in mean time to respond. SentinelOne shipped zero-click autonomous investigation to its entire customer base in June 2026. Google's triage agent has processed more than five million alerts. If these claims are true — and this report finds the best of them substantially

credible — a puzzle follows: the same vendors maintain elite human incident-response teams, sell incident-response retainers, and are expanding those services, not winding them down. Sophos launched an advisory practice; Microsoft bundled incident response into a new managed suite in January 2026; CrowdStrike's IR business remains central to its identity. If the machines resolve incidents, why is everyone still hiring humans to respond to them?

The resolution is that a case is not an incident, and the industry's marketing language systematically blurs a distinction its own service catalogues make plain. The autonomy statistics describe cases: the high-volume stream of alerts, detections, and routine containments — a phishing email quarantined, a commodity malware infection isolated, an abused credential locked. These tasks are pattern-rich, high-frequency, and low-ambiguity, which is precisely the territory where machine autonomy is real, measurable, and in Sophos's case published to the second. What the statistics do not describe is the major incident: ...

..... [\(see full report\)](#)

The analyst verdict. The agentic SOC is real, and the best implementations are measurably transformative at the case level. The incident level ...

..... [\(see full report\)](#)

3.4 The ungoverned AI problem, extended

In early 2026, LLG's published research argued that Southeast Asia's AI adoption was outrunning its AI governance — that the region was deploying capabilities faster than it was deploying control. This report extends that thesis into the security market with a finding that emerged, unsought, in every category researched: the enterprise AI agent has become simultaneously the newest attack surface and the newest insider, and the security industry has converged — with unusual speed and unanimity — on the same answer. Agents must be governed like employees.

The attacker side of the evidence arrived first. Group-IB's High-Tech Crime Trends analysis for 2026 concludes that identity and trust have become the primary attack surfaces of the industrialised cybercrime economy; Mandiant's M-Trends data shows attacker coordination collapsing from an eight-hour hand-off between initial-access brokers and follow-on actors to twenty-two seconds over three years — machine-speed adversaries operating against human-speed defences. When both the region's leading adversary-intelligence source and the world's largest incident-response practice ...

..... [\(see full report\)](#)

The defender side then shipped, across every tier of this report's vendor universe and mostly within a twelve-month window. Google built the most infrastructure-level implementation: Agent Identity issues unique, scoped, delegation-aware credentials to autonomous agents on an open standard; Agent Gateway enforces policy on agent-to-agent and agent-to-tool traffic; Model Armor protects the models themselves; and Wiz's AI bill of materials inventories every model, framework, and extension in the estate to surface the shadow AI nobody approved. Okta extended the workforce identity fabric to agents as first-class identities. SentinelOne shipped a discovery-and-governance control plane for AI agents plus red-teaming for AI applications. Zscaler placed an AI broker and agent registry inline in the network. Palo Alto bound its agents to the same role-based access control as human analysts, with full reasoning audit trails. Tenable and XM Cyber, from the exposure side, began inventorying MCP servers and scanning agent configurations for embedded credentials. Ensign, from the regional services side, extended the same governance logic to AI-driven physical systems. These are not announcements of intent; with few exceptions they are shipping products, and this report's Part 9 maps them as a coherent product category that did not exist eighteen months ago.

One protocol detail deserves elevation from footnote to finding, because it appeared in every single category of this study: the Model Context Protocol, the emerging standard by which AI agents connect to tools and data, is simultaneously becoming the interoperability layer of the agentic enterprise and its newest attack surface...

..... [\(see full report\)](#)

For Southeast Asia, the governance question carries a regional sharpening. The region's talent constraint makes agentic adoption more attractive here than almost anywhere — the leapfrog logic Part 9 examines — which means the ungoverned-agent risk compounds faster here too: organisations with the least capacity to supervise autonomous systems have the strongest incentive to deploy them. The regulatory frameworks of the six covered markets, examined in Part 8, are at varying and mostly early stages of addressing autonomous-agent accountability, which leaves enterprise governance — agent identity, scoped privilege, audit, and kill-switch authority — as the operative control layer for the planning horizon. The good news this research surfaces is that the tooling for that control layer now exists, at every price tier, from global platforms to regional operators.

The analyst verdict. Treat every AI agent as a new hire with system access: it gets an identity, a scope, a manager, an audit trail, and an offboarding path — and it gets them before it starts work, not after the first incident. The vendors profiled in this report now make that posture implementable. The organisations that implement it will capture the region's leapfrog opportunity; the ones that do not will discover that they have hired, at machine speed, an unsupervised workforce.

PART 4 — METHODOLOGY

4.1 Why two axes

This report scores every fully evaluated vendor on two independent axes rather than one composite number, because the research's central empirical finding is that AI capability and Southeast Asia fitness are weakly correlated — and in several categories, inversely correlated. Fortinet Unified SASE and Palo Alto Prisma SASE illustrate the problem a single score would create within the same category: the former scores 4.55 on SEA Fit and 3.70 on AI Readiness, the latter 3.83 and 4.13 — near-opposite profiles whose blended composites (4.13 and 3.98) sit close enough together to look like a tie. A composite would tell a reader they are near-interchangeable; they are close to opposite purchases for buyers who actually need different things. The two-axis presentation, and the quadrant chart built from it, preserves the trade-off that is this report's actual subject.

SEA Fit (0–5.0) measures operational deployability in the six covered markets, weighted: regulatory and residency fit 30 percent, talent-gap fit 25 percent, delivery layer 25 percent, threat relevance 20 percent. The regulatory pillar carries the top weight because the research found its sharpest and most decision-relevant differentiation there. Threat relevance carries the lowest because, after normalisation (Section 4.5), the global vendors are nearly indistinguishable on it.

AI Readiness (0–5.0) measures the substance of the vendor's AI capability, decomposed per Section 4.2 and weighted: AI-operated 40 percent, AI-analytical 35 percent, AI-securing 25 percent. The operated dimension carries the most weight because agentic execution is what addresses the region's binding constraint — security talent — and is therefore the dimension on which Southeast Asian buyers can least afford to be misled.

A blended figure (60 percent SEA Fit, 40 percent AI Readiness) is computed solely to assign category tiers — Leader at 4.00 and above, Contender from 3.50 — and readers are cautioned that the blend compresses information the axes carry separately. There is no Niche tier in this report; weak vendors were removed at shortlisting under a three-lens test of market presence, technical capability, and operational stability, and the absence of low scores reflects that filter, not the market.

4.2 The three-way AI decomposition

“AI-native” is the most overloaded term in the 2026 security market, and this report does not score it as a single property. Research across the five categories established that vendor AI stories decompose into three distinct capabilities that vendors routinely conflate and buyers routinely need separated.

AI-operated is agentic execution of security work: the platform's AI autonomously performs triage, investigation, remediation, or administration. This is what addresses a missing analyst bench, and it is where the marketing inflation concentrates.

AI-analytical is AI as the detection or prioritisation method itself: the product's core engine is machine learning rather than signatures or rules. This is the oldest and most

verifiable of the three — some analytical engines carry a decade of production history — and it is what the term “AI-native architecture” properly refers to.

AI-securing is governance of the customer’s own AI estate: visibility and control over the enterprise’s models, agents, and agent-protocol traffic. This is a new protection surface, not an operations aid, and vendors’ habit of announcing it under the same “agentic AI” banner as the operated dimension is the single most common conflation this research encountered.

Every evaluated vendor is scored on all three, and the resulting map is more honest than any single scalar: some vendors score strongly on all three; others score on operated and securing but not analytical; some score the inverse — the purest analytical engine in the study attached to a deliberately non-agentic operating model; and for at least one vendor, the securing story is the strongest in the report while its operated score is modest by design. Where a low score reflects philosophy rather than deficiency — a deliberately human-executed remediation model, or an inline-policy-enforcement design — the profile says so in text, because the number alone would mislead.

4.3 Evidence grades and caps

Every autonomy and performance claim in this report carries an evidence grade, and the grade caps the score the claim can support. Grade A — independently observed or audited results — permits scores to 5.0. Grade B — vendor-published operational averages, ideally with named customer corroboration — caps the supported sub-score at 4.5. Grade C — best-case vendor ceilings, recognisable by the words “up to” — caps at 3.5 for whatever rests on the claim. Grade D — roadmap, preview, or capabilities too new for production evidence — is scored on shipped substance only. The cap mechanism makes the scorecard structurally resistant to marketing inflation: a vendor cannot buy a score with an adjective, and a vendor that publishes measured operational averages is rewarded for the risk that honesty entails. The market-level finding bears repeating: nothing in this study earns Grade A for agentic autonomy. Buyers should notice, and procurement teams should begin asking for independently auditable autonomy rates the way they long ago learned to ask for third-party detection testing.

4.4 The regulatory pillar rubric

The regulatory and residency pillar is scored against a published rubric rather than impression. Deployment model is assessed on the three-tier taxonomy from Part 3.1 — cloud-only, vendor-operated in-country region, customer- or partner-hosted sovereign — with cloud-only anchored at a 2.0 baseline. A half-point is added where the vendor operates in-country infrastructure in other geographies, because a shipped sovereignty mechanism that has not yet reached Southeast Asia is materially different from a mechanism that does not exist: the buyer of the former has something concrete to demand. In-country presence within the six covered markets scores progressively above that. Storage and processing locations are checked separately for every vendor, and where a vendor does not publicly disclose, the scorecard prints “undisclosed” — treated as information, not as a gap to be politely skipped. Finally, regulatory stakes are weighted proportionately by category: a SASE platform inspects the customer’s traffic in-line and a

SIEM retains the customer's telemetry, so residency carries full weight; exposure-management products handle metadata about weaknesses rather than the data itself, and their profiles say so rather than applying the pillar mechanically.

4.5 Threat-relevance normalisation

Early scoring drafts awarded differentiated threat-relevance scores to global vendors on the strength of platform-scale telemetry and research reputation. Review exposed the flaw: those are capability arguments, already captured in the analytical dimension of AI Readiness, not evidence of Southeast Asian relevance. The final methodology normalises all global vendors to a common threat-relevance score, with exceptions granted only for verifiable regional evidence — Tenable, uniquely among the globals, publishes exposure research specifically covering all six of this report's markets and is scored accordingly. The benchmark for what full marks require is set outside the evaluated set entirely, by Group-IB, whose monthly regional intelligence, in-country operations, and named-campaign attribution define what "SEA-relevant threat intelligence" means when it is a documented fact rather than a claim.

4.6 The delivery-layer evidence rule

The delivery pillar is scored on countable, named evidence — distributors, MSSP specialisations, and partner operations verifiable per market — rather than on ecosystem adjectives. This rule cuts in both directions and the scorecard shows it: one vendor's score reflects the densest named evidence in the study, while another's decades-deep regional channel scores below a third's because this research could verify named partners for the latter and not the former. Where the rule visibly penalises a vendor whose channel is real but undocumented, the profile flags the score as evidence-limited and revisable — an honest cost of a method that refuses to award points for reputation.

4.7 The frontier-model audit

Because agentic capability increasingly rests on foundation models, every vendor's model strategy was audited across providers as a deliberate neutrality control. The result is a spectrum with sovereignty consequences examined in Part 3.1: vertically integrated (Google, whose Gemini models power its entire security portfolio), multi-model orchestrators (CrowdStrike's ecosystem spans Anthropic, OpenAI, NVIDIA, and Amazon; SentinelOne pairs Anthropic's models with Google threat intelligence under an explicit no-single-model stance; Tenable accepts any model over MCP), self-built specialised models (Trend Micro's open-sourced Cybertron; Sangfor's Security GPT), and single-ecosystem (Microsoft's Copilot stack). Partnerships are treated throughout as evidence of investment, never of capability; the capability test remains shipped behaviour.

4.8 Anticipated objections — how this report handles them

A report that scores vendors should expect their objections and state, in advance, the principles by which each class of objection is answered — with the vendor-specific responses living where the vendor names do, inside the category profiles.

Where a vendor scores below its market reputation, the profile states which evidence gap or structural factor drives the score. The recurring patterns: a validated storage posture without a validated processing posture (residency scores reflect both, and the profile says which is missing); ...

..... (see full report)



PART 5 — VENDOR EVALUATION BY CATEGORY

Five categories, thirteen evaluated vendors. Each opens with a category-scoped quadrant, then the full vendor profiles that carry the evidence behind every score.

5.1 EDR / MDR / XDR

The category question. Endpoint-anchored detection and response is where Southeast Asian security budgets concentrate, where the regional MDR wave is strongest, and where the AI-nativeness spectrum runs widest. The four vendors profiled here span the full range of what a buyer can actually purchase in 2026: (see full report). The chapter’s verdict is deliberately not a ranking. It is a mapping of four consumption models onto four buyer situations — because in this category, more than any other, the right answer depends on what kind of organisation is asking.

Two cross-references frame everything below. Flagship 3.3 (*If the agentic SOC works, why are incident-response retainers growing?*) supplies the severity-tier and evidence-grade discipline applied to every autonomy claim in this chapter. And Flagship 3.1 (*Where does your security actually live?*) supplies the deployment-model taxonomy that separates these four vendors more decisively than any detection benchmark: (see full report).




Figure 2. EDR / MDR / XDR — SEA Fit × AI Readiness (pale dots show the other nine evaluated vendors for context).

CrowdStrike — [SEA Fit & AI Readiness scores: see full report]

..... (see full report)

Trend Micro (TrendAI) Vision One — [SEA Fit & AI Readiness scores: see full report]

..... (see full report)

Microsoft Defender XDR — [SEA Fit & AI Readiness scores: see full report]

..... (see full report)

Sophos MDR — SEA Fit 3.65 · AI Readiness 3.83 · Contender

Overview and stability. Thoma Bravo-owned, with the Secureworks acquisition (completed February 2025, \$859 million) integrating visibly: Counter Threat Unit intelligence merged into Sophos X-Ops, Sophos endpoint bundled into Taegis subscriptions, and an ITDR product launched from the Taegis identity lineage. The watch item — stated as due diligence, not alarm: a post-acquisition workforce reduction and an unresolved long-term consolidation question between the Taegis and Sophos Central platforms. Enterprise buyers should ask for the platform roadmap in writing; mid-market buyers consuming the service layer are less exposed to the answer. A July 2026 development to track: Sophos announced Sophos Fusion, an AI-native defense system unifying SIEM, XDR, and MDR across its platform.

Threat relevance (P1: 4.0). X-Ops incorporating CTU tracks an extensive library of named threat groups (profiles published on Sophos's threat-profiles page), and the platform's MITRE evaluation results are strong on both the Sophos and Taegis lineages. Normalised baseline.

Regulatory and residency (P2: 2.0). The thinnest residency evidence in the category: no disclosed Southeast Asian hosting, no sovereignty mechanism anywhere in the world, and storage and processing locations undisclosed — printed as such per methodology, and treated as information. For buyers in the region's stricter jurisdictions, this is currently the disqualifying pillar; for buyers in permissive ones, it is a contract-clarity task.

Talent-gap fit (P3: 4.5). The strongest "security without a security team" proposition available to the regional mid-market, on the strongest operational evidence any vendor in this study publishes: 52 percent of MDR cases resolved end to end by AI at an average of 89 seconds from alert to automated response, with Sophos's own analysts supervising the AI and owning every outcome — Grade B evidence, measured averages rather than ceilings, across a 40,000 MDR customer base (up 39 percent year on year) with uncapped incident

response included, a breach warranty, and the 4.8/5.0 Gartner Peer Insights rating carried on one of the larger review bases in the MDR category (124 reviews as of July 2026). Flagship 3.3's boundary applies here as everywhere: these are case-tier metrics, the 48 percent remainder still reaches humans, and the bundled incident-response capability is the vendor's own concession that Tier 3 stays human — which, for the buyer this profile addresses, is precisely what they are paying to have handled.

Delivery layer (P4: 4.5). The most channel-native vendor in the category by history and structure, with a two-tier portfolio mapping cleanly onto regional segments: Sophos MDR for the mid-market, typically channel-delivered; Taegis MDR for the enterprise, an open-XDR service that ingests third-party telemetry — including competitors' endpoint agents — which makes it consumable on top of estates the buyer already runs, a pragmatic fit for the region's blended-estate reality. Named Southeast Asian partner verification remains pending; the score is flagged evidence-limited on the same rule applied to Trend Micro.

Cost context (not scored). Publicly reported MDR pricing anchors (median annual contract around \$91,000 in the Secureworks lineage; per-endpoint figures in the \$70–170 range) place the service tier within reach of mid-market budgets that could never staff the equivalent — which is the entire proposition.

AI decomposition. AI-operated 4.5 — the best operational autonomy metrics in the researched market, human-supervised at scale. AI-analytical 4.0 — Taegis ML detection with strong independent showings; a service-first rather than platform-native architecture. AI-securing 2.5 — minimal evidence of capability to govern the customer's AI estate: the portfolio's one clear gap, and an open invitation to publish evidence that would move the score.

Verdict. For a mid-market organisation without a SOC, in a market whose residency rules permit offshore processing, the most directly actionable recommendation in this category: price the alternative of building anything internal against 89 seconds. The completeness caveat belongs beside it: Sophos operates no in-region SOC or disclosed regional hosting, so buyers for whom Pillar 2 matters should pair the service with the residency test below and, where needed, with regional delivery and DFIR partners — Sophos-first is a strong opening position, not automatically the complete set.

Part 5.1 — Who should choose what

The large enterprise with transformation budget and a multi-vendor estate ([see full report](#)); The regulated buyer in Vietnam, Thailand, or Indonesia starts from ([see full report](#)); The E5-licensed, Microsoft-centric estate applies ([see full report](#)). The mid-market organisation with no SOC and permissive residency shortlists Sophos MDR first, prices the alternative against 89 seconds, and completes the set with regional ([see full report](#)); And every buyer in this category should ask each vendor the two-question - ([see full report](#))

5.2 SASE / SSE

The category question. SASE is where the sovereignty question stops being abstract. The product is the network: inspection happens where the point of presence is, and a market without an in-country PoP is a market whose traffic crosses a border to be secured. For two of this report's six markets — Thailand and Vietnam — no evaluated vendor operates an in-country service edge (a Philippine exception was verified at second cut: Fortinet's Manila FortiSASE data centre, noted in the Fortinet profile below), which makes the deployment-model taxonomy from Flagship 3.1 (*Where does your security actually live?*) not one lens among several in this chapter but the organising fact. The three profiled vendors span that taxonomy end to end: a sovereign-deployable hybrid architecture, a cloud platform with the region's best in-country build-out, and a security-depth leader whose regional edge inventory remains undisclosed.

Fortinet Unified SASE — SEA Fit 4.55 · AI Readiness 3.70 · Leader

Palo Alto Networks Prisma SASE — SEA Fit 3.83 · AI Readiness 4.13 · Contender

Zscaler — [SEA Fit & AI Readiness scores: see full report]

..... (see full report)

Part 5.2 — Who should choose what

..... (see full report)

5.3 CTEM / Attack Surface Management

..... (see full report)

5.4 SIEM / SOAR

..... (see full report)

5.5 ITDR

..... (see full report)

PART 6 — SPOTLIGHT PROFILES

Spotlight vendors are assessed against a stated editorial rationale rather than the category scorecard: each earns dedicated attention for a reason the standard grid cannot capture. Neither receives numeric scores; both receive honest verdicts, and both close with the named revisit triggers a tracking posture requires.

6.1 Google (Security Operations + Mandiant + Wiz)

Why spotlighted. The most architecturally compelling converging platform in this study, attached to the least mature regional delivery layer — a combination that a scored grid would flatten into a middling number and that buyers should instead track deliberately, with triggers.

The convergence, no longer assembling. The Wiz acquisition completed in March 2026, and the stack now ships as a single strategy rather than three products in adjacent booths: Wiz detections flow into Google Security Operations and Mandiant Threat Defense with automatic threat-information forwarding; Wiz integrates with Security Command Center for Google Cloud estates; and the whole is packaged commercially as Google Unified Security above tiered SecOps editions. The scoping assumption this research began with — three separately acquired products with visible seams — required updating within the research window itself: seams remain, but the closure rate since the acquisition completed has outpaced expectation.

The agentic evidence. Among the strongest operational numbers in this report: Google's Triage and Investigation agent has processed more than five million alerts, compressing a typical thirty-minute manual analysis to roughly sixty seconds — Grade B evidence, vendor-published volume with named enterprise corroboration ...

..... (see full report)

The governance stack — Flagship 3.4's deepest implementation.

..... (see full report)

Threat intelligence. Mandiant's investigation practice — half a million incident-response hours in the most recent year — anchors this report's adversary-tempo evidence: threat-actor coordination compressing the hand-off between initial-access brokers and follow-on actors from eight hours to twenty-two seconds across three years. VirusTotal's corpus and Google-scale telemetry complete an intelligence depth nobody disputes; ...

..... (see full report)

— Google’s Southeast Asian security presence would strengthen beyond what its current published intelligence shows. This research found no verification that such regional compilation occurs, and the profile carries it as exactly what it is: an unverified structural advantage worth asking Google about directly.

The Bangkok fact. Google Cloud’s Bangkok region — live since January 21, 2026, with three availability zones, in-country storage and in-country processing, and Thailand’s Kasikorn banking group citing Bank of Thailand compliance at launch — is the single most consequential infrastructure development in this report’s coverage window, and the first in-country hyperscaler region in Thailand from any of the big three. The bifurcation caveat from Flagship 3.1 applies in full: ...

..... (see full report)

The constraint, stated plainly. ...

..... (see full report)

Verdict and triggers. Track now; buy when the delivery layer matures — and for Google Cloud-committed enterprises the calculus is already favourable, with ...

..... (see full report)

6.2 Sangfor Technologies

Why spotlighted. The mid-market and MSSP-channel reality of Thailand, Vietnam, and Indonesia that global platform pricing does not serve — and a live test of what AI-augmentation honestly delivers at the affordable tier, which is a datum this market needs published more than it needs another enterprise-platform comparison.

Presence. Shenzhen-listed, with more than 100,000 customers globally and offices in all six covered markets — Singapore, Malaysia, Thailand, Indonesia, the Philippines, and Vietnam — the only vendor below the global tier in this study with that footprint, and with mid-market depth in it the globals lack. Thailand is described by regional leadership as among the company’s fastest-growing markets, with Thai-language presales, a local country organisation, and a working partner ecosystem on public display; published regional customers include Yamaha Motor’s Thai operations and Indonesian banking and education institutions.

The AI story, honestly tiered. Security GPT — a self-developed model — is embedded in Athena XDR as an AI SOC agent framework for triage assistance and investigation acceleration, and the combination ...

..... (see full report)

Delivery. ...

..... (see full report)

Procurement context, stated neutrally. As a China-headquartered vendor, Sangfor is subject to procurement-policy considerations that vary by buyer: some government-linked, critical-infrastructure, and allied-supply-chain organisations restrict or review PRC-origin security software, while other segments across the region weigh it neutrally or favourably on price-performance and local support. This report treats jurisdictional exposure even-handedly — the American vendors in this study carry US CLOUD Act exposure that sovereignty-minded buyers weigh in the same column — and leaves the policy judgment where it belongs: with the buyer, made consciously rather than by default in either direction.

Verdict. For the Indonesian, Thai, and Vietnamese mid-market, ...

..... (see full report)

PART 7 — EMERGING REGIONAL VENDORS

The regional section is presented as an architecture, not a ranking: an intelligence layer, a delivery-and-governance layer, and an incident-response layer that together complement — rather than compete with — the global platforms. Regional vendors carry SEA Fit scores on the same rubric as the evaluated set, a qualitative AI-posture note in place of a numeric AI Readiness score, and per-market applicability strips — because a regional score is real only inside its footprint, and the strip is the honesty mechanism that keeps it true.

7.1 Group-IB — [SEA Fit score: see full report] · *Applicability: SG · TH · VN strong; MY/ID/PH remote-served*

..... (see full report)

7.2 Ensign InfoSecurity — SEA Fit 4.35 · *Applicability: SG · MY · ID strong; TH served-not-staffed; VN/PH outside footprint*

..... (see full report)

7.3 LGMS Berhad — [SEA Fit score: see full report] · *Applicability: MY · SG; ID/PH/TH project-based*

..... (see full report)

7.4 The combination archetype

..... (see full report)

A global AI-native platform, delivered and operated through a regional MSSP with in-country presence, plus the permanent regional DFIR layer, assembles from two or three vendors what almost none offers alone — and the pattern already operates at recognised standard in this region:... .. (see full report)

The recommendation is explicit. Sovereignty-sensitive buyers (see full report); capability-first buyers (see full report); and every buyer (see full report)

Vendor selection in Southeast Asia is increasingly an exercise in assembling stacks, not choosing champions — and this report's market-by-market guidance, next, is written accordingly.

PART 8 — MARKET-BY-MARKET GUIDANCE

Each market section states the regulatory frame, the infrastructure reality, the vendor applicability picture, and differentiated recommendations. Markets diverge from the template where structure demands — and two of the six diverge substantially.

8.1 Singapore — the full-option market, and the market that writes the region's homework

Regulatory frame. The PDPA under an active and communicative regulator; MAS Technology Risk Management guidelines setting the financial-sector bar; IMDA and CSA shaping technology and cyber policy; and the region's most mature AI-governance reference point in the Model AI Governance Framework. Sophisticated but navigable: the binding constraint in Singapore is diligence quality, not option scarcity.

Infrastructure reality. Everything reaches Singapore first, and this report's research quantifies it: Singapore is 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 (operating since 2022), Azure's established region, Ensign's home operations, Group-IB's headquarters. Vendors treat Singapore as the region's landing zone, and the buyer's abundance is real.

The hub, in both meanings. Singapore's regional role carries two distinct implications, and this report treats them separately because they demand different responses. The first is architectural: the regional-headquarters pattern, in which vendor decisions made in Singapore quietly govern subsidiary estates across the region, is Southeast Asia's most common route to accidental non-compliance — ...

..... (see full report)

The second is regulatory, and it is field experience rather than statute: what Singapore's regulators do today, neighbouring regulators cite tomorrow. MAS technology-risk expectations, IMDA's AI governance work, and CSA's frameworks function as the region's reference implementations — several covered markets' supervisory practices track Singapore's with a lag measured in quarters, not decades. The practical consequence for every reader, Singaporean or not: (see full report) — which is the cheapest compliance strategy this report can recommend.

Guidance. Singapore buyers face the region's only genuine abundance problem, and the discipline is architectural: choose on capability and integration; apply the MAS TRM lens to processing location even where storage is local; and — for regional headquarters — run every subsidiary market through this Part before defaulting the regional estate to Singapore-hosted services. The two-question residency test from Flagship 3.1 (*Where does your security actually live?*) should be asked per market, not per vendor.

8.2 Malaysia — the ascending infrastructure story, with an unwritten regional chapter

Regulatory frame. The amended PDPA with breach-notification obligations now in force; the Cyber Security Act framework concentrating obligations on NCII-designated sectors; and an increasingly assertive national cyber agenda that vendors now court openly — Zscaler’s Kuala Lumpur launch was publicly framed inside the NCII protection narrative, ministerial endorsement included.

Infrastructure reality. The fastest-improving map in the region after Thailand: Azure Malaysia West live since May 2025 with a second Malaysian region announced for Johor; Zscaler’s full Zero Trust Exchange hosted in Kuala Lumpur, expanded with a new data centre opened in January 2026; in-country Security Copilot processing arriving within Microsoft’s 2026 wave; and the Johor AI-infrastructure build-out signalling where regional compute gravity is heading.

The regional role its law has not yet caught up with.

..... (see full report)

Guidance.

..... (see full report)

8.3 Indonesia — scale, sovereignty, and the OJK shadow

..... (see full report)

8.4 Thailand — the market where the map is moving, and the risk no vendor fully owns

Regulatory frame.

..... (see full report)

Infrastructure reality — the watershed.

..... (see full report)

The convergence risk no vendor fully owns. Thailand’s manufacturing concentration creates a security condition this report found unaddressed as a whole by any evaluated vendor: the convergence of operational technology and corporate IT...

..... (see full report)

Guidance — with a standing reassessment clause.

..... (see full report)

8.5 Vietnam — the hardest constraint, deliberately navigated

Regulatory frame. The strictest binding environment in the six: the Personal Data Protection Law in force with its implementing decree attaching revenue-linked penalties — transforming residency from a compliance preference into board-level financial risk — alongside the region’s first dedicated AI statute, a combination that makes both where data lives and how AI systems are governed questions of law rather than posture.

..... (see full report)

8.6 Philippines — the differentiated playbook

Regulatory frame — the offset. The Data Privacy Act under NPC administration, with a penalty regime whose sliding scale carries real weight — but a localisation posture materially more permissive than Vietnam’s or Indonesia’s:...

..... (see full report)

8.7 Vendor presence at a glance

A market-level summary of where every evaluated, spotlight, and regional entry actually operates in Southeast Asia, gathered from each profile's own P2 residency finding. This table is presence, not performance — it says nothing about which vendor should win a given market; the SEA Fit and AI Readiness scores in Part 2 and the full profiles in Part 5 do that work.

Vendor	Category	Presence	Notes
CrowdStrike	EDR/MDR/XDR	See full report	See full report
Trend Micro (TrendAI)	EDR/MDR/XDR	See full report	See full report
Microsoft Defender XDR	EDR/MDR/XDR	See full report	See full report
Sophos MDR	EDR/MDR/XDR	Undisclosed	No Southeast Asian hosting disclosed; no sovereignty mechanism
Fortinet Unified SASE	SASE	Customer/partner-hosted sovereign	FortiSASE Sovereign places the full stack in customer/partner facilities
Palo Alto Prisma SASE	SASE	Partial (ingestion only)	On-prem SD-WAN controller and Private Locations; no dedicated SEA hosting region
Zscaler	SASE	See full report	See full report
Tenable One	Exposure Mgmt	See full report	See full report
XM Cyber	Exposure Mgmt	See full report	See full report
Cortex XSIAM	SIEM/SOAR	See full report	See full report
Microsoft Sentinel	SIEM/SOAR	See full report	See full report
SentinelOne	Identity/ITDR	Singapore (since 2022); AU, Mumbai (virtual)	Vendor-operated in-country for one of six covered markets

Okta	Identity/ITDR	See full report	See full report
Google (Unified Security)	Spotlight	Global (Google Cloud)	Least mature regional delivery layer in this study — not scored
Sangfor	Spotlight	All six markets (offices in each)	Mid-market and MSSP-channel focus — not scored
Ensign InfoSecurity	Regional operator	SG · MY · ID strong	Thailand served, not staffed
Group-IB	Regional operator	See full report	See full report
LGMS Berhad	Regional operator	See full report	See full report

PART 9 — CROSS-CUTTING THEMES

Five patterns crossed every category boundary in this research. Each is treated briefly here because the flagships and chapters carry their evidence; each earns its place because a buyer who misses it will feel the omission within a budget cycle.

9.1 Securing the agentic enterprise

..... (see full report)

9.2 Identity as the control point

..... (see full report)

9.3 The frontier-model spectrum

..... (see full report)

9.4 The consumption-cost shift

..... (see full report)

9.5 Where the talent gap goes

..... (see full report)

PART 10 — READER DECISION TOOLKIT

This Part converts the report into decisions. It answers, in order, the questions readers bring to a vendor report before they bring vendor names — each in brief, with pointers into the parts that carry the depth.

Define success before you shop. Detection, response, and resilience are different purchases. A buyer whose real gap is detection coverage weighs the analytical dimension and threat-relevance evidence; one whose gap is response capacity weighs the operated dimension, the managed layers, and the autonomy boundary; one tasked with resilience weighs deployment models, identity architecture, and the DFIR layer as a system. The fastest way to misuse this report is to shortlist Leaders without deciding which of the three you are buying; the quadrant tells you where vendors are, not what you need.

Check your prerequisites honestly. AI-native cloud platforms assume estates they can see: cloud workloads, modern endpoints, healthy telemetry. Blended estates — the regional norm outside Singapore — should weigh the vendors architected for them (the hybrid-architecture platforms, the open-XDR services that ingest third-party telemetry) and treat connector and telemetry hygiene as a funded precondition of any agentic purchase, because agents reason over what the estate feeds them. Cloud-native transformation is not a prerequisite for AI-era security in this region; choosing vendors that match your actual estate is.

Calibrate scope to your mandate.

..... (see full report)

Split the ownership deliberately.

..... (see full report)

Treat cost as posture, not a table.

..... (see full report)

Adjust for industry criticality — in both directions.

..... (see full report)

Know what you control, and plan the rest.

..... (see full report)

Appendix A — Claims Register

Every load-bearing factual claim in this report is registered below with its verification status as of the research cut-off (July 20, 2026). Vendor-published figures are marked as such both here and where they appear in the text; absences and non-disclosures are printed as findings per the methodology. Ninety-two claims are registered; one further item raised during verification (a vendor-relayed analyst projection) was withdrawn when it could not be traced to a primary source, and does not appear in this report.

Infrastructure & residency

..... (see full report)

Autonomy & AI capability

..... (see full report)

Market, channel & recognition

..... (see full report)

Adversary evidence

..... (see full report)

Regulatory (VIRC lineage)

VIRC refers to the [SEA Cybersecurity Vendor Intelligence — Regulatory Compliance Report H1 2026](#), this report's predecessor edition (see [About this report](#)).

..... (see full report)

Structural findings (report-original)

..... (see full report)

Appendix B — Sources

Sources consulted in the verification of registered claims. Per-claim source mapping is maintained in the claims register; vendor-published materials are used only for claims explicitly marked vendor-published in the text.

Vendor primary sources and announcements

..... (see full report)

Regulatory and legal sources

..... (see full report)

Analyst, media, and third-party sources

..... (see full report)