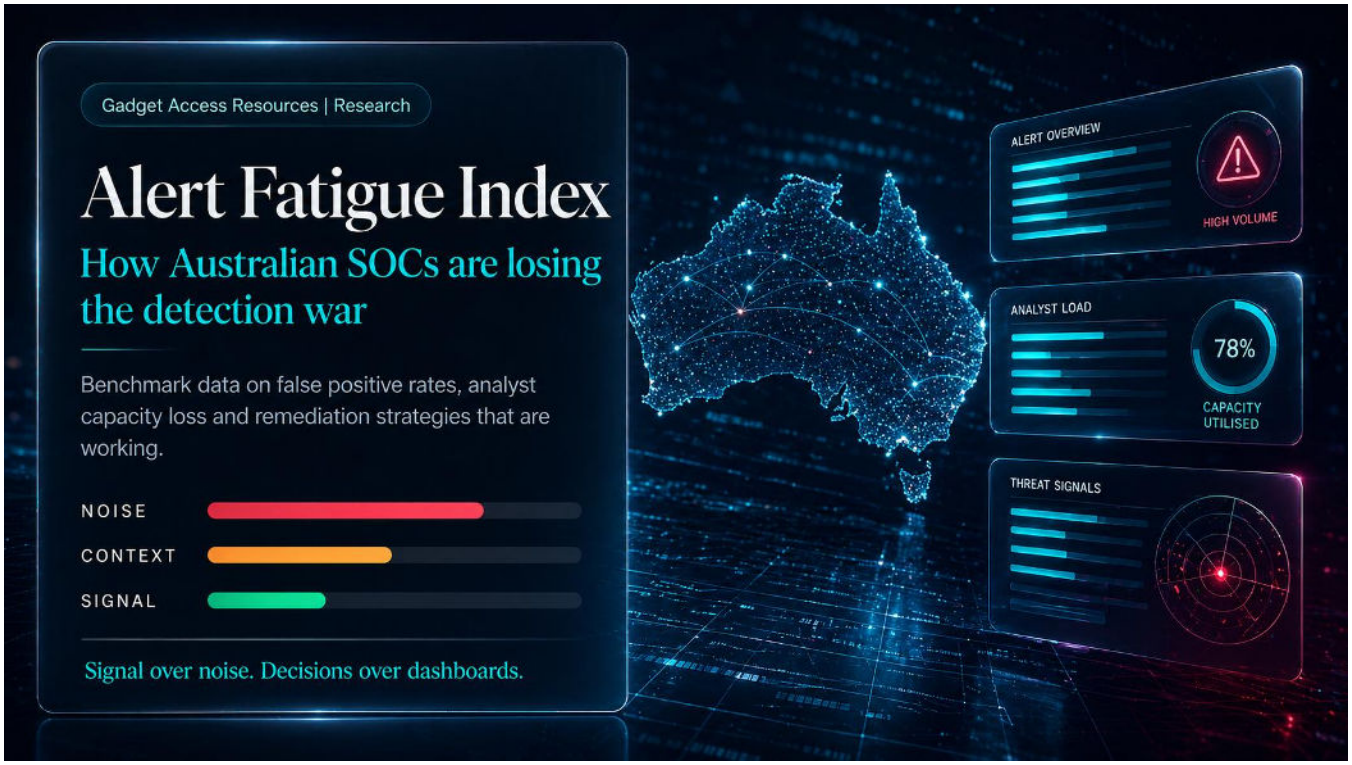

Alert Fatigue Index: How Australian SOCs Are Losing the Detection War

Benchmark data on false positive rates, analyst capacity loss and the remediation strategies that are working.

Gadget Access Resources | Research | Estimated read time: 5 minutes

Core thesis: alert fatigue is not a human performance problem. It is an operating model defect. The fix is not simply fewer alerts. It is better decisions, better context and faster proof of remediation.

Why alert fatigue is now a governance problem

Australian security operations centres are not failing because analysts cannot read alerts. They are failing when the operating model asks people to do mathematically impossible work and then calls the backlog a resourcing issue. Alert fatigue is the polite term. A better description is detection capacity leakage.

The threat side of the equation is growing. ASD's Australian Cyber Security Centre reported more than 84,700 cybercrime reports in FY2024-25, roughly one every six minutes, alongside more than 1,200 cyber security incidents and more than 1,700 notifications to entities of potentially malicious cyber activity [1]. That volume does not fall neatly into one SIEM queue, but it does explain why boards are asking for more visibility, more reporting and more assurance. The trap is that more visibility can create more decisions, not better decisions.

Detection tooling has improved. Australia is full of well-meaning EDR, identity, cloud, NDR, SIEM, SOAR, CTI, vulnerability, ticketing and GRC stacks. The problem is not the existence of tools. It is the cost of stitching them into a workflow where each alert carries enough context to be trusted, assigned and fixed. Without that context, the SOC becomes a very expensive inbox with nicer logos.

The benchmark: what the data says

The numbers are uncomfortable. The SANS 2025 Detection and Response Survey found that false positives were the leading challenge in detection, cited by 73% of respondents, while more than 60% encountered false positives frequently or very frequently [3]. That is not background noise. It is the sound of analyst capacity being spent on decisions that rarely move risk.

Australian managed SOC data tells the same story in local language. Triskele Labs' State of Cyber SOC 2025 reported 247,922 processed alerts and 10,442 ticketed cases, meaning only 4.2% of processed alerts became tickets. Of those ticketed cases, 40.08% were true positives, 59.52% were false positives and 0.39% were indeterminate. The report also notes that nearly 70% of processed alerts were not actionable [2]. This is a vendor benchmark, not a census of every Australian SOC, but it is still a useful proxy for the operating pressure many teams recognise immediately.

The workforce picture completes the equation. The SANS 2025 SOC Survey found that 79% of surveyed SOCs operate 24/7, while two to ten people is the most common fully staffed SOC size, 62% of SOC professionals say their organisation is not doing enough to retain top talent, and 69% report metrics manually or mostly manually [4]. Put bluntly, many SOCs are trying to run an airline schedule with a go-kart pit crew.

False positives are not always false

A false positive label can hide several different failure modes: a rule that is too broad, a true but benign event, insufficient identity context, duplicate telemetry, an alert from an asset nobody owns, or a valid issue that cannot be fixed inside the business window. That distinction matters. An alert does not need to be technically wrong to waste time. It merely needs to be non-actionable.

USENIX research into SOC analyst perspectives found that many perceived false positives are benign triggers explained by legitimate behaviour. The same work proposed five alert quality properties that make validation easier: reliable, explainable, analytical, contextual and transferable [5]. Those words should be printed above every detection engineering backlog. They are a better north star than simply counting how many rules fired overnight.

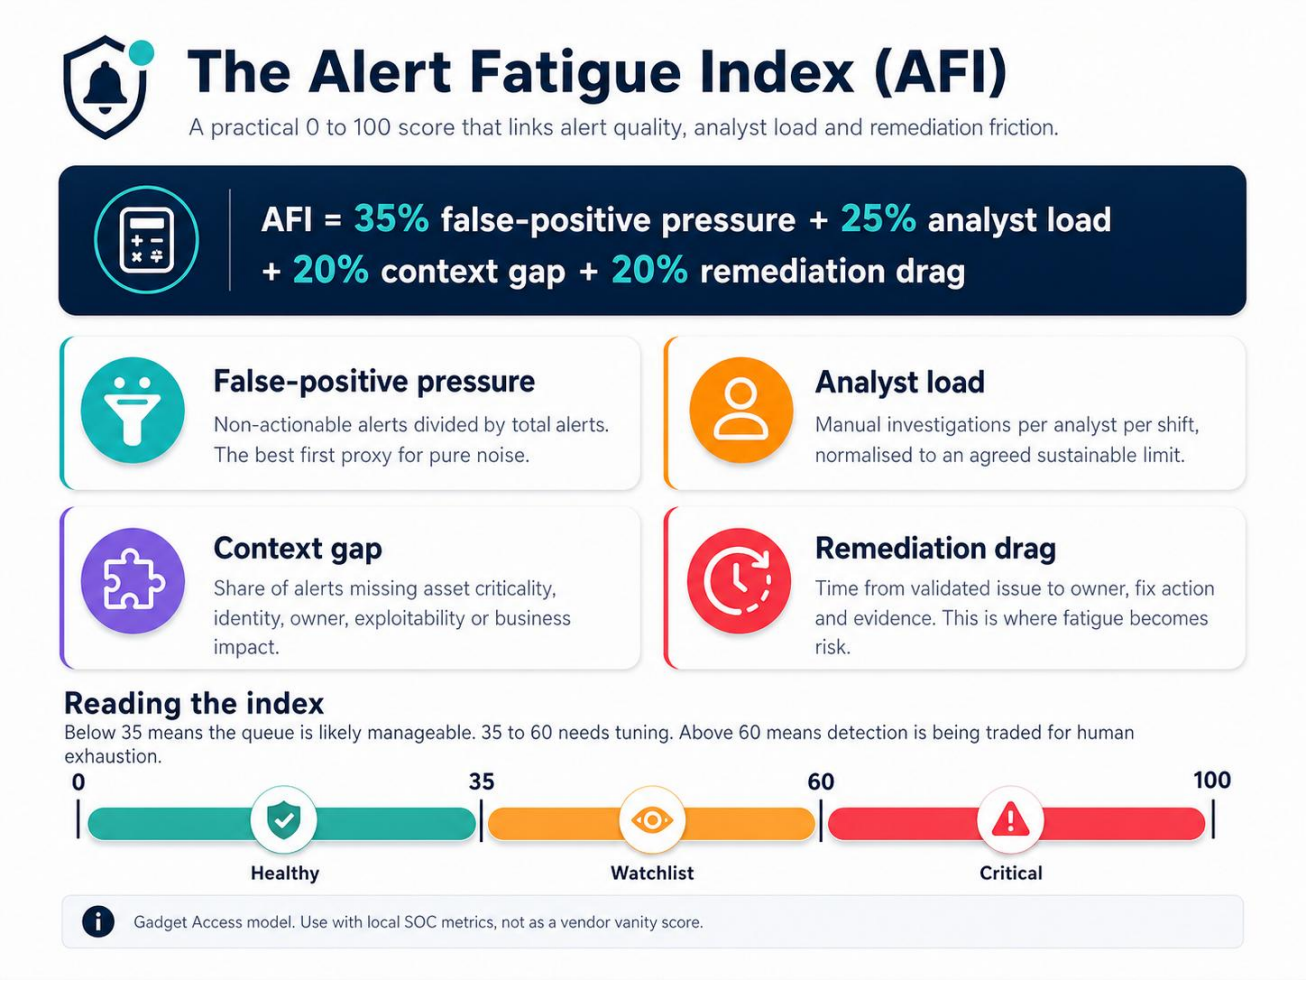

Figure 1. The Gadget Access Alert Fatigue Index model turns alert noise into a management metric.

The Alert Fatigue Index

Gadget Access recommends starting with an Alert Fatigue Index, or AFI, before buying another platform or asking analysts to work faster. The AFI is not a vendor score and it is not intended to shame a SOC. It is a management metric that links alert quality, analyst load, context gaps and remediation drag into a practical 0 to 100 score.

The model uses four inputs: false-positive pressure, analyst load, context gap and remediation drag. The weighting is intentionally pragmatic. False-positive pressure receives the highest weighting because it is the most visible source of waste. Analyst load matters because the same alert volume means something different in a five person team than it does in a mature 24/7 operation. Context gap measures whether an analyst can see asset criticality, identity, owner, exploitability and business impact without opening five more systems. Remediation drag captures the time between a validated issue and proof that the risk has been reduced.

The last input is often the most revealing. Many SOC dashboards stop at triage, but the attacker does not care when a ticket was opened. Risk reduces when an owner fixes the issue, evidence is captured, and detection is tuned so the lesson is retained. Below 35, the queue is probably manageable. From 35 to 60, the operating model needs tuning. Above 60, the organisation is effectively trading detection for human exhaustion.

Where capacity disappears

The AFI becomes powerful when translated into capacity. Using the Australian managed SOC benchmark, 247,922 processed alerts and a 59.52% false-positive rate implies roughly 147,500 false-positive determinations. At five

minutes each, that is about 12,300 hours. At ten minutes each, it is about 24,600 hours. At twenty minutes each, it becomes about 49,200 hours. Based on 1,920 productive hours per analyst-year, that is roughly 6, 13 and 26 analyst-years respectively.

This stress test is illustrative. It does not rebuild Triskele's staffing model, and it does not assume every alert is investigated in the same way. The point is simpler: false positives are not free. Every low-value decision displaces threat hunting, purple teaming, detection tuning, vulnerability validation, control evidence and the quiet thinking that good investigations require.

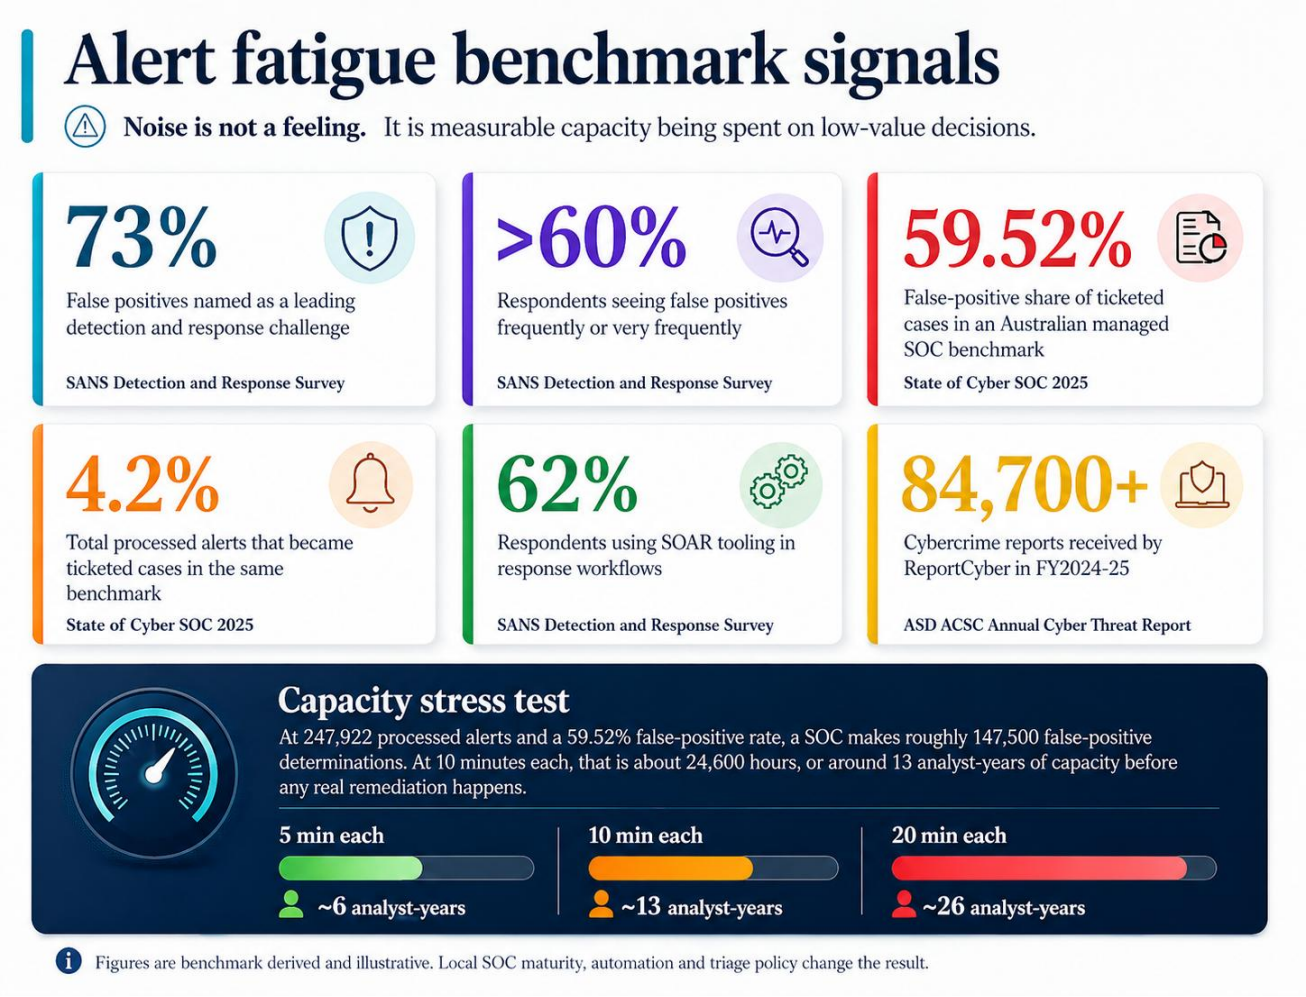

Figure 2. Benchmark signals and a simple capacity stress test show how quickly alert noise becomes analyst-year loss.

What remediation strategies are working

The first strategy is detection ownership. Every rule should have an owner, a hypothesis, a data source, an expected response, a tuning history and a retirement condition. If nobody owns a rule, the SOC eventually becomes a museum of good intentions. Detection engineering is not a one-off configuration exercise. It is a product management discipline for signals.

The second strategy is enrichment before triage. Alerts should arrive with asset criticality, identity, business service, exposure, known vulnerabilities, privilege level, recent change and relevant threat intelligence. Context turns a queue into a decision system. Without it, analysts become professional tab-switchers, and no board should be funding tab-switching as a cyber strategy.

The third strategy is automating the repeatable, not the judgement. The SANS 2025 Detection and Response Survey found that 66% of respondents use partial automation and 13% use full automation in response workflows, with

predefined playbooks and SOAR among the common mechanisms [3]. The sensible pattern is to automate collection, enrichment, deduplication, containment guardrails and evidence capture. Human judgement should be reserved for ambiguity, business impact and adversary behaviour.

The fourth strategy is measuring validated remediation instead of queue movement. A closed ticket is not the same thing as reduced risk. Good SOC programs connect alerts to vulnerability management, identity controls, endpoint hygiene, cloud posture, change management and compliance evidence. This is where Gadget Access' Continuous Compliance framework becomes useful: it keeps the control obligation, technical finding, owner action and evidence trail close enough that uplift can be measured continuously rather than reconstructed under audit pressure.

Why AI changes the stakes

AI will not make alert fatigue disappear by magic. It will change the shape of the bottleneck. Anthropic describes Project Glasswing as a defensive coalition in which partners receive access to Claude Mythos Preview to find and fix vulnerabilities in foundational systems. Anthropic says Mythos Preview is available as a gated research preview and has identified thousands of zero-day vulnerabilities across critical infrastructure [7]. Microsoft has also described plans to incorporate advanced AI models such as Claude Mythos Preview into secure development lifecycle processes to identify vulnerabilities and develop mitigations earlier [8].

That is good news for defenders, but there is a sting. If AI lowers the cost of vulnerability discovery, the scarce resource becomes prioritisation, validation, ownership, safe remediation and proof. A future SOC will not only receive more endpoint and cloud alerts. It may also receive more AI-generated findings, more exploitability claims and more demands for immediate action. Without operating discipline, the new AI era becomes a faster conveyor belt into the same tired queue.

The economics still favour intelligent automation. IBM's 2025 Cost of a Data Breach research reports that extensive use of security AI and automation reduced breach lifecycle by an average of 80 days and saved an average USD 1.9 million compared with organisations that did not use it [6]. The lesson is not to avoid AI. The lesson is to deploy AI where it reduces decisions, adds context and shortens the path to verified remediation.

Where Gadget Access and CiBRAI fit

For Gadget Access, this is the design point for CiBRAI: not another blinking console, not another alert source, and definitely not another reason for analysts to keep twelve browser tabs open. CiBRAI should act as a simplified cyber operating layer that brings alerts, assets, obligations, vulnerability intelligence, investigation notes, owners and evidence into one working model.

By integrating the latest on-premise and cloud large language models, CiBRAI can help summarise evidence, correlate telemetry, translate findings into remediation paths, map controls to obligations, and keep the human analyst focused on judgement rather than clerical archaeology. Paired with Gadget Access consulting, virtual CISO support and technical uplift capability, that operating model can be used across Essential Eight, PSPF, ISM, ISO 27001:2022, NIST and local risk frameworks without turning compliance into a quarterly spreadsheet festival.

This is the difference between having tools and having an operating model. Tool sprawl asks analysts to assemble context after the alert fires. A continuous compliance operating model ensures the context is ready before the decision is made.

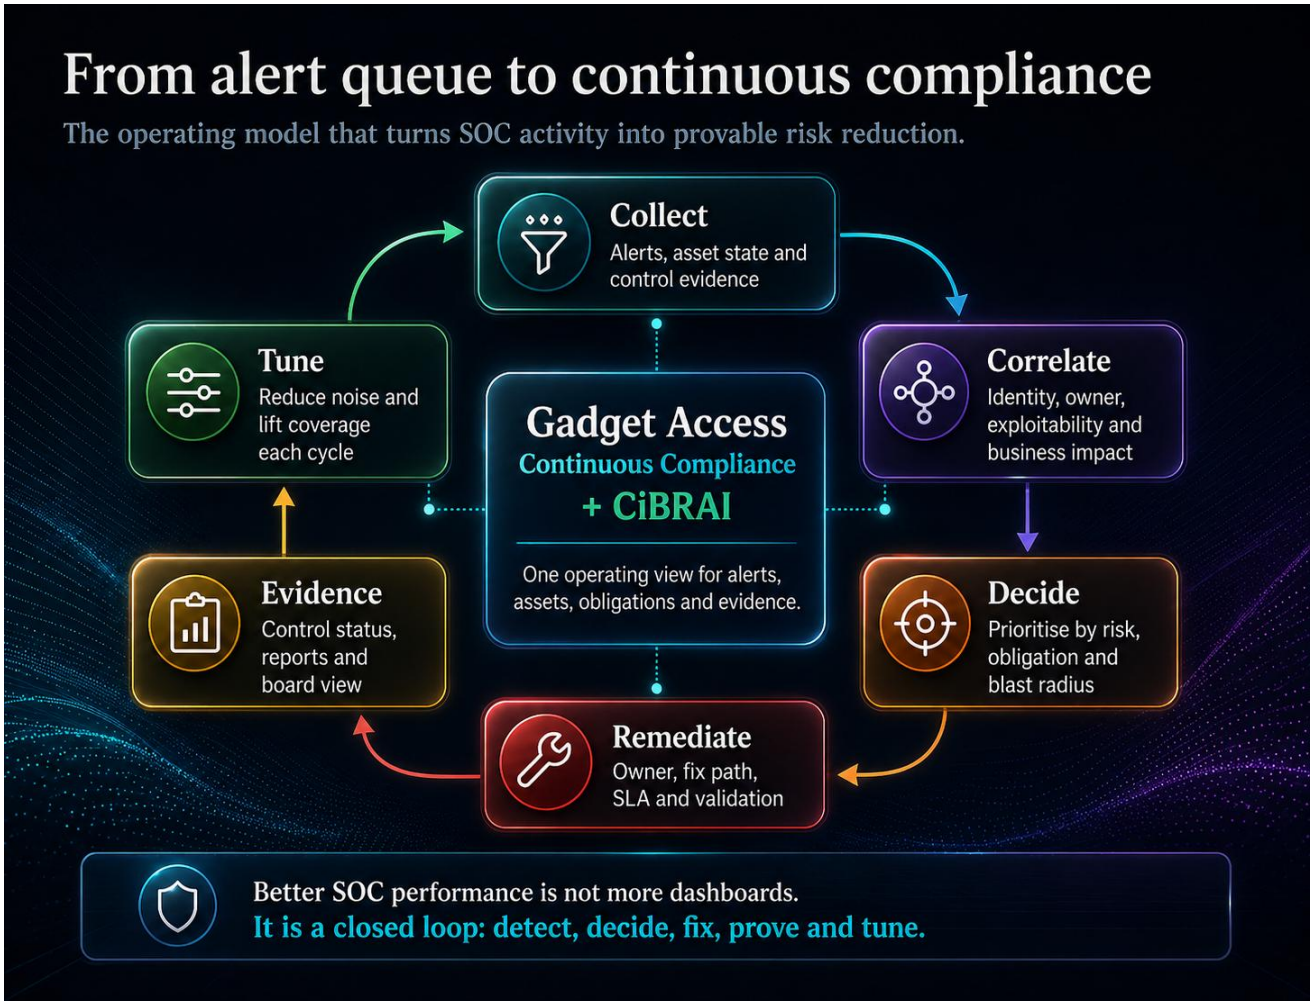

Figure 3. Continuous Compliance connects detection, remediation, evidence and tuning so the SOC becomes a closed loop.

Conclusion: signal over noise

Australian SOCs are not losing the detection war because defenders lack effort. They are losing ground when leaders accept a queue that cannot be worked, a SIEM that stores more than it explains, and compliance reporting that arrives after the risk has already moved on.

The Alert Fatigue Index reframes the conversation. It turns alert volume into capacity leakage, false positives into measurable cost, and remediation drag into governance risk. The winning SOC of the next two years will measure noise as a cost, automate safely, tie every detection to business context, and treat remediation evidence as part of the same workflow. It will not ask analysts to be heroic every shift. It will reserve human judgement for the alerts that deserve it.

Signal over noise. Decisions over dashboards. Evidence over excuses. That is the practical path back to advantage.

References and further reading

[1] Australian Signals Directorate's Australian Cyber Security Centre. Annual Cyber Threat Report 2024-25. https://www.cyber.gov.au/about-us/view-all-content/reports-and-statistics/annual-cyber-threat-report-2024-2025

[2] Triskele Labs. State of Cyber SOC 2025: Detection, Response and Alert Trends. https://www.stateofcyber.com.au/report/mdr

[3] SANS Institute. SANS 2025 Detection and Response Survey. https://www.opswat.com/resources/reports/sans-detection-response-survey

[4] SANS Institute. SANS 2025 SOC Survey. https://www.sans.org/white-papers/sans-2025-soc-survey

[5] Alahmadi, B. A., et al. 99% False Positives: A Qualitative Study of SOC Analysts' Perspectives on Security Alarms. USENIX Security 2022. https://www.usenix.org/conference/usenixsecurity22/presentation/alahmadi

[6] IBM. Cost of a Data Breach Report 2025. https://www.ibm.com/reports/data-breach

[7] Anthropic. Project Glasswing. https://www.anthropic.com/project/glasswing

[8] Microsoft Security. AI-powered defense for an AI-accelerated threat landscape. https://www.microsoft.com/en-us/security/blog/2026/04/22/ai-powered-defense-for-an-ai-accelerated-threat-landscape/

Prepared for Gadget Access Resources. Separate reusable image assets are included in the accompanying image asset package.