WHITE PAPER

NEXTGEN SOCS: BUILDING HYPERAUTOMATED, AI-FIRST SECURITY OPERATIONS

Abstract

Security Operations Centres (SOCs) are approaching a breaking point. Escalating threat velocity, the emergence of offensive AI, analyst burnout, and increased regulatory pressure are overwhelming traditional operating models. This whitepaper outlines a vision for the next generation of SOCs: hyperautomated, AI-first operations built to scale, adapt, and defend in an environment where attackers increasingly leverage automation and large language models (LLMs).

It explains why this evolution is now essential, the architecture required, and how to blend AI, automation, and human expertise into a resilient modern SOC.

themissinglink®
An Infosys company

Infosys®
Navigate your next

Industry problem: Why traditional SOCs are no longer enough

1. Exponential growth in AI-enabled threats

The past three years have seen an explosion in LLM capabilities for both defenders and adversaries[1,2]. Attackers can now deploy thousands of autonomous LLM agents to attempt exploits, craft phishing lures, automate reconnaissance, and scale campaigns to unprecedented levels. The limiting factor is no longer skill; it's simply cost.

People were shocked to learn of an adversary using Anthropic's services at scale to perform human-in-the-loop attacks against multiple organisations simultaneously[3]. Now consider what a nation state is capable of when hosting their own LLMs, removing constraints that limit their civilian counterparts.

2. Escalating operational and governance pressure

Security leaders face growing obligations across:

ASD Essential Eight maturity uplift: This framework has moved from being a government-focused baseline to a cross-sector expectation in Australia[4]. Progressing through the maturity levels, especially levels 2 and 3, requires demonstrable evidence of consistent security controls. Manual processes rarely deliver the reliability and scalability needed for this uplift.

ISO 27001:2022 operational controls: The 2022 update to ISO 27001 introduced more rigorous requirements around risk assessment, performance evaluation, and continuous improvement[5]. To meet these, organisations need the ability to map security controls to real time threat activity and produce audit ready reports quickly.

NIST CSF 2.0 outcome-driven performance metrics: The NIST Cybersecurity Framework now emphasises measuring security outcomes rather than simply checking off activities[6]. This shift means SOCs must provide evidence of effectiveness, such as detection rates, dwell time reduction, or improved recovery metrics, something legacy, manual processes struggle to support.

Board-level reporting mandates: Executives and boards increasingly expect SOCs to provide contextualised, business aligned insights. This includes understanding risk posture, regulatory exposure, and how incidents may impact operations or reputation. Without automation and contextual enrichment, these insights are difficult to generate at scale[7].

Tightened breach notification requirements (globally): With breach disclosure windows tightening globally - often to 72 hours or less - organisations must rapidly detect, classify, and document security incidents[8,9]. Manual investigation pipelines are too slow and pose compliance and reputational risk.

Manual processes cannot meet these expectations at scale[7].

3. Tool sprawl and loss of security context

Data lives across SIEMs, EDRs, identity systems, SaaS logs, cloud telemetry, and ticketing systems.

This distributed architecture increases capability but creates complexity. Each tool sees a different slice of the environment, requiring human analysts to manually correlate and piece together the picture. As adversaries move faster, this delay becomes a critical vulnerability.

Without automation and AI-driven context building, no analyst can see enough of the environment fast enough to beat an AI-powered adversary.

Modern SOCs must unify context in real time to support rapid triage, response, and reporting. This is no longer a nice-to-have; it is essential to outpace machine-driven attacks.

Recommended solution: The hyperautomated, AI-first SOC

A next generation SOC is not defined by an XDR, a SIEM, or a set of playbooks. It is defined by four core principles:

Principle 1: AI-driven detection, enrichment, and decisioning

Traditional SIEM correlation rules often miss nuance, require extensive tuning, and produce overwhelming volumes of alerts. LLMs and AI-driven engines offer a fundamentally different approach: interpreting, reasoning, and enriching data with speed and depth. This principle shifts the SOC from static rule or scoring based detection to dynamic, adaptive signal interpretation.

LLMs are uniquely suited to:

Interpreting unstructured logs: LLMs excel at making sense of raw, disparate log formats, helping identify anomalies without requiring strict schemas.

Summarising incidents: Instead of sifting through long chains of alerts, analysts can be presented with natural language summaries that explain who, what, where, and how, accelerating decision making.

Enriching alerts with organisational context: AI can cross reference internal data sources (e.g. asset criticality, business ownership) to prioritise threats.

Correlating disparate signals: Rather than relying on rigid correlation rules, AI can identify behavioural patterns across toolsets and timeframes.

Understanding Standard Operating Procedures (SOPs): LLMs can be trained or prompted to understand and execute SOPs, enabling automation with contextual nuance.

Principle 2: Hyperautomation across Tier 1 and Tier 2 workflows

The SOC of today is drowning in alerts, repetitive tasks, and procedural overhead. Hyperautomation is about removing inefficiency, not people. By automating Tier 1 and Tier 2 activities, from triage to evidence collection, teams can reduce response times and focus human effort where it matters most.

Automation replaces human repetition across:

Alert triage: Rule-based and AI-augmented triage ensures only relevant threats reach human analysts.

Containment actions: Automatically isolate affected hosts, reset credentials, or block malicious IPs where risk is low and predefined.

Common threat hunting routines: Automation can execute hunts using behavioural and signature-based logic, flagging anomalies for human follow up.

Evidence collection: Gathering forensic artefacts from endpoints, cloud workloads, or identity providers no longer requires manual effort.

Incident documentation: Automatically compile case files, timelines, and summaries for incident post-mortems and audit readiness.

Principle 3: Human-in-the-loop by design

AI may be fast, but security is ultimately a risk-based, business-aligned function. Human judgment remains critical, particularly when actions have regulatory, reputational, or strategic impact. This principle embeds analysts into the loop as decision-makers and supervisors, ensuring AI acts as an assistant, not a replacement.

While AI takes on mechanical and analytical tasks, humans remain decision makers for:

- **Strategic risk choices:** Not all actions can or should be automated. Decisions involving business risk require human context and discretion.
- **Ambiguous or high impact actions:** When data is inconclusive or the potential impact is significant, escalation to human review is essential.
- **Regulatory interactions:** Disclosures, communications with authorities, and compliance actions require human oversight.
- **Proactive threat hunting:** Humans are best at spotting unknown unknowns, using intuition and hypothesis driven exploration.

No model, regardless of capability, should remove humans entirely.

Principle 4: Single unified platform

Tool sprawl leads to context loss, analyst fatigue, and poor decision making. A unified platform brings telemetry, reasoning, and workflows into a single ecosystem, reducing friction and enabling faster, more consistent response in the face of AI-powered threats.

Fragmented tools and dashboards slow analysts and increase operational overhead. Next-generation SOCs use unified platforms where SIEM, SOAR, EDR/XDR, threat intelligence, case management, and AI reasoning engines operate seamlessly.

This reduces the need to context switch between consoles, accelerates response, and enables true end-to-end visibility.

These principles are reflected in the reference architecture shown in Figure 1.

Reference architecture for an AI-first Security Operations Centre

An AI-first SOC requires a clear architectural separation between intelligence inputs, security telemetry, analytics, automation, and human decision making. The architecture below shows how threat intelligence and raw security data are consolidated, analysed, and enriched before flowing into automated response workflows and analyst led operations. Each layer has a defined function, ensuring signals are contextualised, prioritised, and actioned in a controlled manner. This structure supports faster detection and response while preserving governance, auditability, and clear accountability across SOC operations.

Figure 1: Reference architecture for an AI-first security operations centre

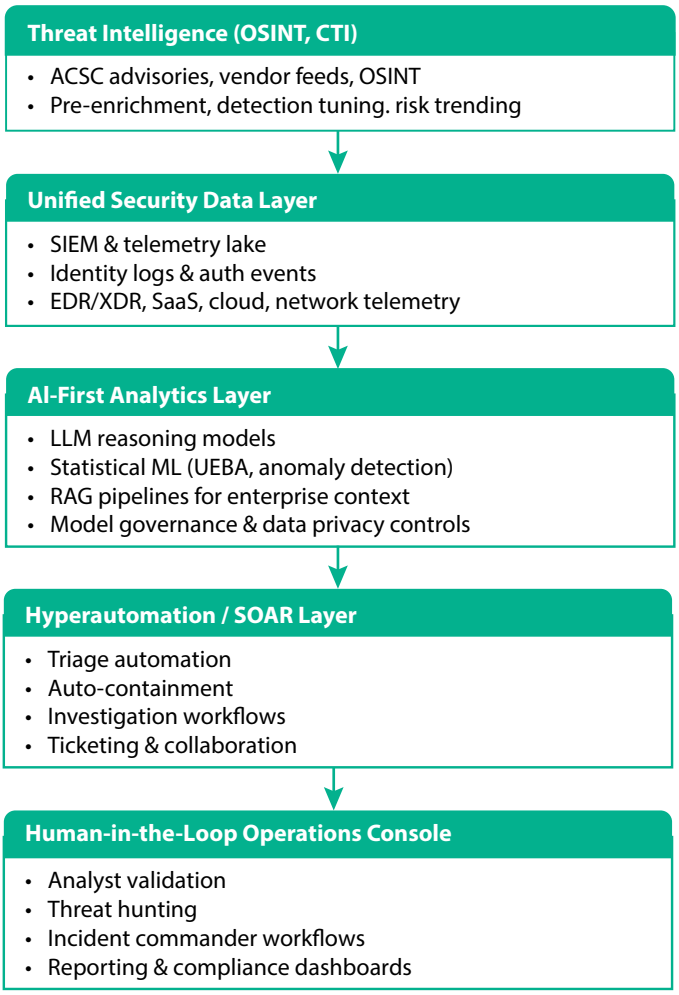

Best practices for building an AI-first SOC

1. Build around adversary centric frameworks

AI-first SOCs should align their detection and response capabilities to adversary behaviour frameworks. This ensures automation is mapped to real attacker tradecraft rather than abstract controls or compliance checklists. When AI is grounded in observed adversary behaviour, it produces more relevant detections and more defensible response actions.

Key frameworks include:

MITRE ATT&CK for detection logic, behaviour analytics, and threat hunting: ATT&CK provides a common language for mapping telemetry to attacker tactics and techniques, enabling consistent detection engineering and structured threat hunting [10].

MITRE D3FEND for defensive control selection and automation logic: D3FEND helps organisations link specific defensive actions to adversary behaviours, supporting explainable automation and measurable control effectiveness [11].

Cyber Kill Chain for structuring automated responses across the attack lifecycle: Mapping response workflows to Kill Chain stages ensures containment and remediation actions are timely and proportionate to attacker progress.

This alignment ensures that AI and automation are targeting real, observed adversary techniques, not theoretical compliance requirements.

2. Implement strong model governance and data controls

AI systems must operate within well-defined governance boundaries to avoid introducing new operational, privacy, or security risks.

Core practices include:

Restricting sensitive logs through policy-driven access controls: Not all telemetry should be available to all models. Access should be governed by data sensitivity, regulatory requirements, and operational need[12].

Using private or containerised LLMs to maintain data sovereignty: For many organisations, particularly in regulated or critical infrastructure sectors, private deployments are necessary to control data residency and exposure[12].

Prevent contamination of defensive models with unverified or attacker crafted data: Defensive AI must be protected from poisoning through strict validation of training and inference data sources[12].

Logging, auditing, and versioning all model inferences and automation decisions: Every AI-driven action should be traceable. This supports forensic analysis, regulatory review, and continuous improvement[12].

Strong governance ensures AI enhances security outcomes without undermining trust or accountability.[12]

3. Deploy reasoning models only where they add strategic value

Not all SOC activities require deep reasoning. Reasoning capable LLMs should be applied selectively, where context synthesis and judgement materially improve outcomes.

Reasoning capable LLMs should be reserved for tasks where depth, context, and synthesis matter:

- Multi-signal correlation across ATT&CK techniques
- Kill Chain stage prediction and escalation
- Production of high-quality incident summaries
- Analyst decision support during active investigations

Lightweight models and automations should be used for high-volume triage to optimise cost and performance.

4. Use retrieval augmented generation (RAG) to ground AI in your environment

RAG is critical to ensuring AI operates with environment-specific context rather than generic training knowledge. Without grounding, AI outputs risk being inaccurate or operationally irrelevant.

Best practices include:

Indexing detection logic mapped to ATT&CK
This allows AI to reason using existing detection strategies and known coverage gaps.[10]

Connecting the LLM to asset inventories, identities, topology, and SOPs
Environmental context ensures responses align with business priorities and operational constraints.

Pulling live enrichment from threat intelligence feeds and telemetry stores
Real-time inputs ensure decisions reflect current conditions, not historical assumptions.

Grounded AI produces outputs that are relevant, defensible, and operationally actionable.

5. Automate repetitive work before introducing AI

Automation is a prerequisite for effective AI adoption. AI amplifies existing processes but cannot compensate for poorly defined or unreliable workflows.

Priority areas include:

Workflow automation for common tasks
Consistent processes create a stable foundation for AI-assisted decisioning.

Auto-containment for low-impact assets
Pre-approved actions reduce response times without increasing risk.

Automated evidence collection for investigations
This improves speed and ensures consistency across incidents.

Clean and well-structured SOC telemetry pipelines
AI results are only as good as the data provided.

Drafted ticketing and communications
Automation reduces administrative overhead and improves reporting quality.

AI delivers value only once foundational automation is reliable and repeatable.

6. Evolve SOC roles into AI enhanced specialisations

An AI-first SOC does not reduce the importance of human expertise. It changes where that expertise is applied. Analysts shift from reactive alert handling to higher-order, intelligence-driven roles.

Key skill areas include:

Threat hunting aligned to ATT&CK
Focused on hypothesis driven exploration rather than alert chasing.

Engineering countermeasures aligned to D3FEND
Designing and refining controls that directly disrupt adversary behaviour [11].

Authoring and maintaining automation playbooks
Ensuring workflows remain accurate, auditable, and effective.

Incident command and strategic decision-making
Coordinating response across technical, legal, and business stakeholders.

Interpreting and validating AI-generated recommendations
Maintaining human accountability for security outcomes.

This evolution creates a human–machine collaborative SOC capable of scaling beyond traditional operating limits.

Conclusion

Attackers are accelerating their operations through automation and AI, fundamentally changing the economics of cyber offence. Traditional SOC models, built on manual analysis and fragmented tooling, cannot keep pace with this shift.

The SOC of the future must operate on a different foundation. Hyperautomation, AI-driven reasoning, and unified telemetry are no longer optional enhancements. They are core requirements for maintaining visibility, response speed, and governance in modern environments. Just as importantly, human expertise must be elevated, not removed, with analysts focusing on decision making, oversight, and adversary focused defence rather than repetitive execution.

Organisations that fail to adapt are not standing still. They are accepting growing detection gaps, slower response times, and increasing regulatory exposure. In an environment where adversaries can deploy autonomous AI agents at scale, evolving the SOC operating model is no longer a future consideration, it is a present-day necessity.

References

1. CrowdStrike. (2024). 2024 global threat report. https://www.crowdstrike.com/global-threat-report/
2. Mandiant. (2024). M-Trends 2024. Google Cloud. https://www.mandiant.com/resources/m-trends
3. Anthropic. (2024). Disrupting AI espionage. https://www.anthropic.com/news/disrupting-ai-espionage
4. Australian Cyber Security Centre. (2023). Essential Eight mitigation strategies. Australian Signals Directorate. https://www.cyber.gov.au/acsc/view-all-content/essential-eight
5. International Organization for Standardization. (2022). ISO/IEC 27001:2022 Information security management systems.
6. National Institute of Standards and Technology. (2024). Cybersecurity Framework (CSF) 2.0. https://www.nist.gov/cyberframework
7. Gartner. (2023). Market guide for security operations platforms.
8. U.S. Securities and Exchange Commission. (2023). Cybersecurity risk management, strategy, governance, and incident disclosure. https://www.sec.gov/news/press-release/2023-139
9. Australian Government Attorney-General's Department. (2023). Privacy Act review and reform. https://www.ag.gov.au/rights-and-protections/privacy/privacy-act-review
10. MITRE Corporation. (2024). MITRE ATT&CK® knowledge base. https://attack.mitre.org/
11. MITRE Corporation. (2024). MITRE D3FEND™ knowledge base. https://d3fend.mitre.org/
12. National Institute of Standards and Technology. (2023). AI Risk Management Framework (AI RMF 1.0). https://www.nist.gov/itl/ai-risk-management-framework

About the Author

Tim Niblett

Head of Security Operations

Tim leads the Security Operations function at The Missing Link with a team providing 24/7 Detection, Response and Engineering services. He has a 30-year career in IT, working in complex industries such as Defence, Law Enforcement and Finance in both client-side and supplier-side roles and major transformational projects. His career includes global experience in the UK, France, Spain and Australia as well as managing delivery teams in the Americas and Asia.

Contact The Missing Link

E-mail: contactus@themissinglink.com.au

Website: www.themissinglink.com.au