WHITE PAPER

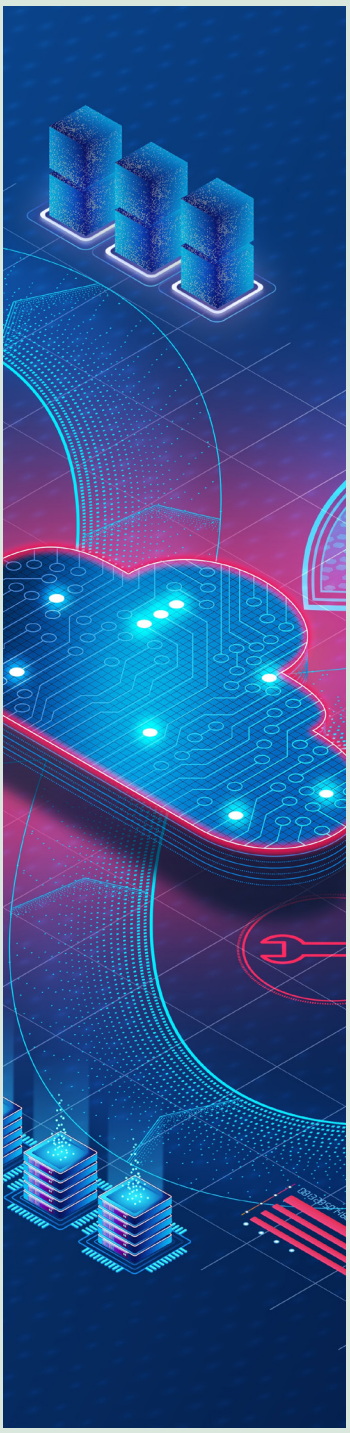

AI-AUGMENTED SECURITY OPERATIONS CENTERS: ENABLING INTELLIGENT DETECTION AND RESPONSE AT SCALE

Infosys®
Navigate your next

Executive Summary

Security Operations Centers (SOCs) act as the nerve center of an organization's cyber defense. They monitor, detect, analyze, contain, and respond to cybersecurity incidents in real time — protecting business continuity, customer trust, and regulatory compliance. SOCs act as a centralized fulcrum, bringing together teams across IT and business functions to drive unified cybersecurity defense.

As threat volumes rise and attack vectors evolve, traditional rule-based approaches are no longer enough to protect today's digital enterprises. To keep pace, organizations must shift toward AI-augmented SOC models that blend automation, intelligence, and human judgment along with Asset Centric, Attacker Centric and Threat Centric models.

This whitepaper draws from our experience in leading this evolution by integrating artificial intelligence and machine learning (AI/ML) into core SOC operations. Drawing from real-world use cases and measurable outcomes within Infosys CyberSecurity, we demonstrate how this transformation has helped reduce mean time to detect/respond (MTTD/MTTR), optimize analyst productivity, and align threat detection with business risk priorities and regulatory compliance.

Moreover, this whitepaper provides a strategic and operational roadmap for building AI-Augmented Security Operations Centers — designed to help businesses enhance threat detection, streamline compliance, and drive cyber resilience at scale. It reflects our experience in enabling Cyber Defense Centers, security teams, and transformation leaders to align security outcomes with evolving business priorities.

Current State of SOCs

Most traditional SOCs operate with rigid, rule-based detection systems, limited context sharing, and high analyst fatigue. These centers often rely heavily on signature-based detections, resulting in:

- **High false positive rates**
- **Long MTTD (Mean Time to Detect) & MTTR (Mean Time to Respond)**
- **Reactive vs. predictive postures**
- **Low analyst productivity** due to repetitive triage tasks
- **Limited prioritization** of alerts based on business context

As per SANS monitoring security across a growing and changing attack surface was cited as the most significant security operations challenge, AI is having an upper hand here to assess and monitor large attack surface.

Modern threat actors now use polymorphic malware, living-off-the-land techniques, and AI-enabled attack methods. Static use cases and manual correlation can no longer keep pace.

Need of the hour:

A transformation toward predictive, adaptive, intelligence-driven SOCs (AI-first SOC) — built to ingest behavior analytics, automate triage, and prioritize risks based on impact and exposure.

The Role of Artificial Intelligence in Modern SOCs

AI enhances every stage of the SOC lifecycle. Here is the structure of SOC integrated with Generative AI solutions.

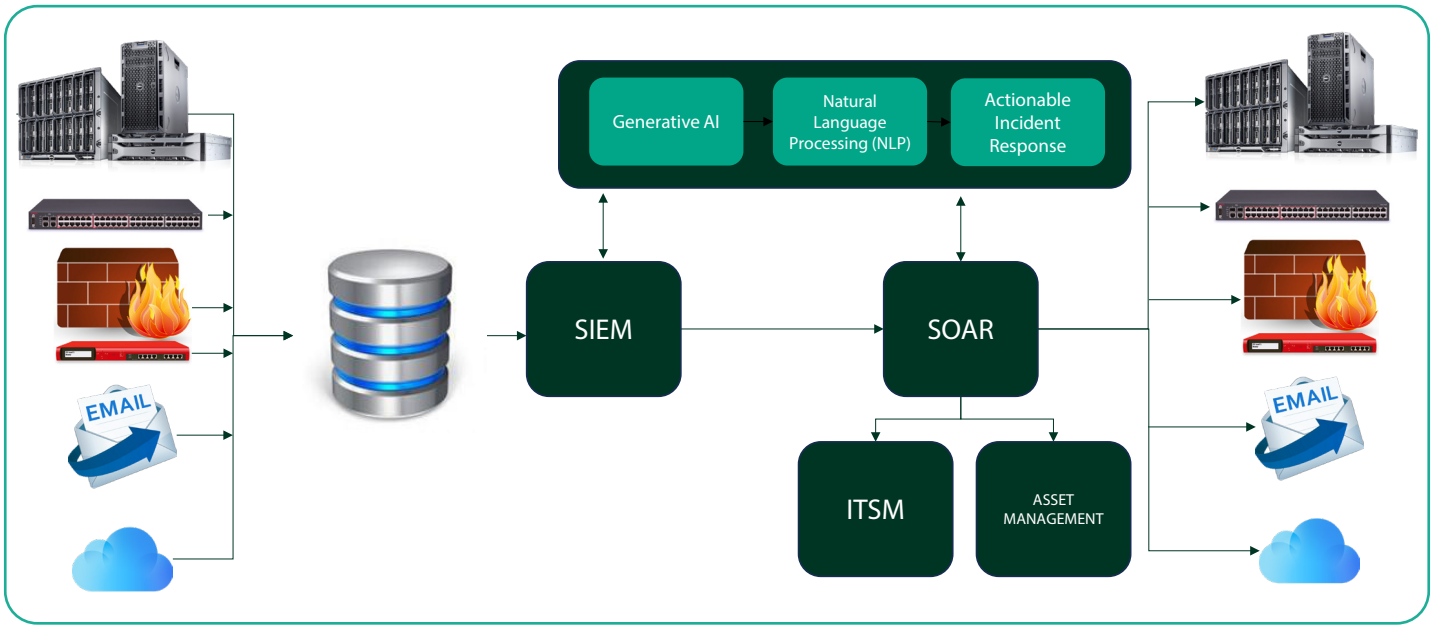

Fig. 1 Generative AI SIEM | SOAR Integration

Area	AI-Augmented Capability
Detection	Behavior-based anomaly detection using ML models, integration with Cyber Threat Intelligence (CTI) and identification of sophisticated attacks which may otherwise go undetected
Triage	Auto-prioritization based on risks, confidence scoring & business impact
Response	Integration with SOAR platforms for automated playbooks
Investigation	AI-guided analysis with pattern recognition across logs
Learning	Continuous feedback loop to retrain detection models

AI Use Case Development Lifecycle

This lifecycle guides the transformation of security problems into AI-driven solutions.

Use Case Idea

use case to be created for the specific log source/ framework

Data Collection

gathers relevant logs, telemetry, and contextual data across environments

ML Model Training

uses structured features to train AI/ML algorithms to detect patterns

Validation with Analyst

the analysts review and refine detection logic, ensuring operational relevance

Integration into SIEM/ SOAR

deploys model into live SOC pipeline for real-time alerts/automation

Continuous Tuning

incorporates feedback and threat shifts to improve accuracy and reduce drift This lifecycle ensures that detection models are not only technically effective but operationally relevant — continuously refined through feedback and real-world applicability.

Different AI Models

To operationalize AI within SOC environments, it has adopted engineering models that integrate AI capabilities into the detection, triage, response, and learning processes. These models enable scale, precision, and continuous improvement while keeping analysts at the center.

1. Threat Scoring Model (Risk = Likelihood × Impact)

A practical risk-based formula that powers AI-driven prioritization in the SOC. This model prioritizes alerts based on the calculated threat severity.

Threat Score = (Anomaly Confidence × Asset Criticality × Threat Context Weight)

Anomaly Confidence	assesses how unusual the behavior is
Asset Criticality	evaluates how valuable the targeted asset is
Threat Context Weight	adds intelligence from past attacks or known threat sources
Threat Score	is derived by multiplying anomaly confidence, asset criticality, and threat context weight — offering a quantifiable risk metric to prioritize alerts.

This model helps focus analyst attention on **what matters most,** improving signal-to-noise ratio in high-volume environments.

Parameter	Value (0–1 scale)	Example
Anomaly Confidence	0.9	High match to malicious pattern
Asset Criticality	0.8	Payment server
Threat Context Weight	0.7	External IP + Privilege escalation

Final Threat Score = 0.9 × 0.8 × 0.7 = 0.504

(Above threshold → Alert promoted for review)

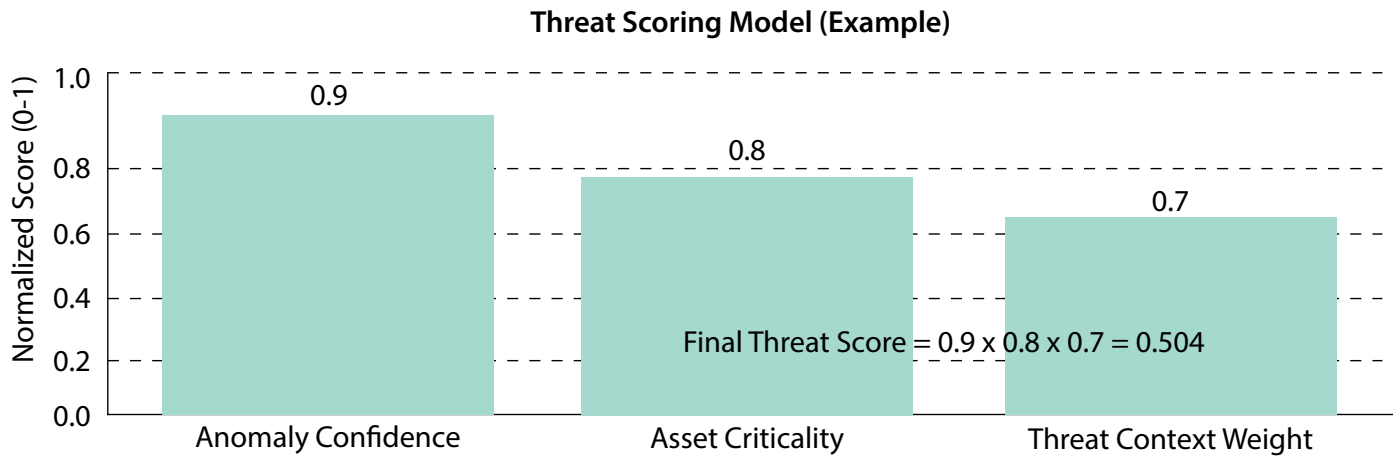

Fig. 2 Threat Scoring Model

2. Feedback Loop Model – AI Learning Lifecycle in SOC

Log Ingestion	from various log sources
Feature Extraction	turns raw logs into meaningful attributes
Anomaly Detection	AI detects unusual behavior using extracted features
Analyst Feedback	helps refine what's normal or risky
Model Retraining	makes AI smarter next time

This continuous loop enhances detection precision over time, reduces false positives, and builds model resilience against evolving attack patterns. It transforms static detection logic into a living system that learns and improves with each incident.

3. Use Case Coverage Heatmap (AI vs. Rule-Based)

This comparative model illustrates how traditional rule-based detections are extended by AI's behavioral and contextual capabilities.

Identify Critical Use Cases	for different scenarios/ attack vectors
Assess Existing Rule-Based Coverage	evaluate logic gaps and alert noise
Design AI-Augmented Logic	develop behavior- or anomaly-based models
Compare Detection Depth	visualize improvement across use cases
Prioritize SOC Investment	focus on areas with biggest detection gaps and business risk

By mapping use cases such as phishing, lateral movement, insider threats, and DNS tunnelling, this model highlights how AI provides **behavioral visibility** that traditional rule-based systems often miss — making the SOC more adaptive to advanced threats.

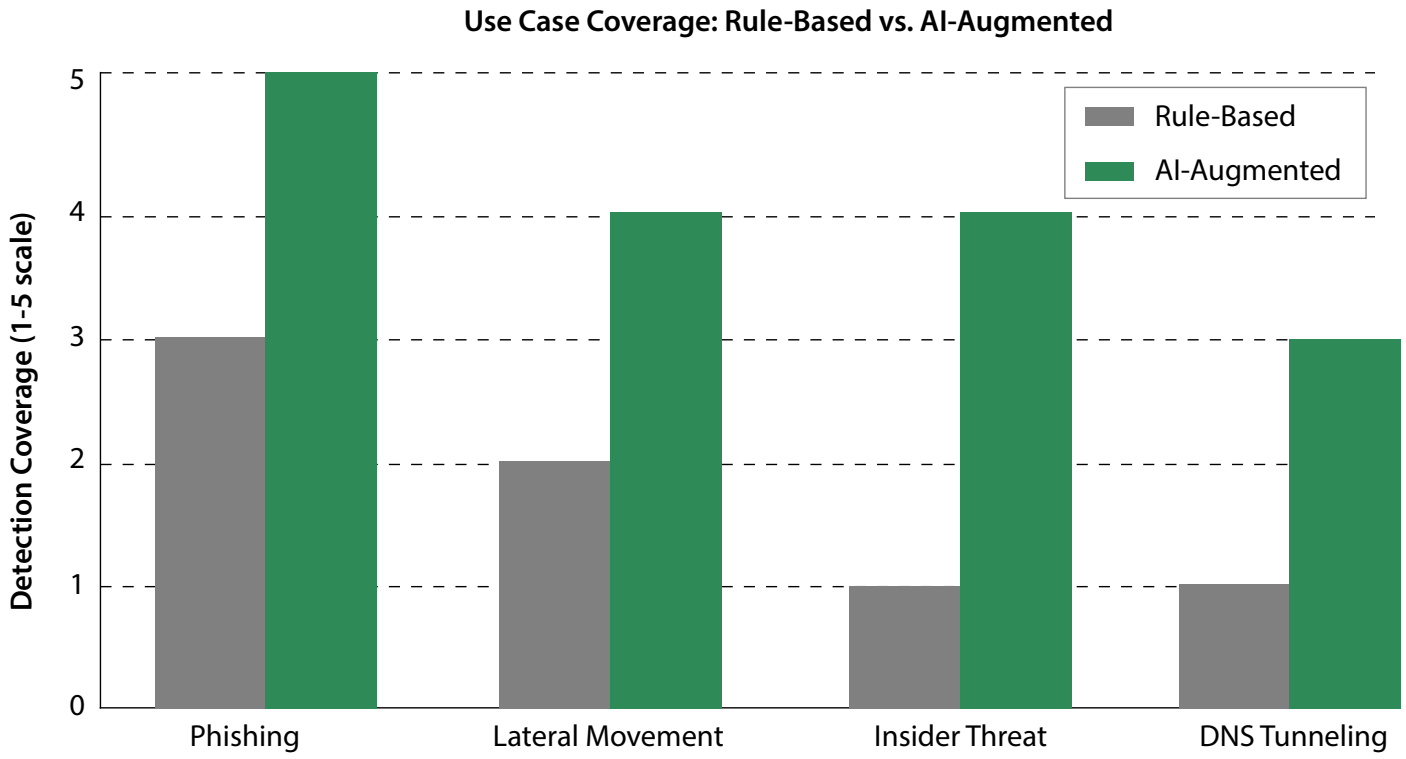

Fig. 3 Use Case Coverage

The use case coverage helps security teams identify gaps and prioritize investment in AI-driven coverage.

4.Human-AI Interaction in SOC Workflow

Alert Ingested	from various log sources
AI Filters & Tags	scores, enriches, and labels alerts based on confidence and context
Analyst Reviews	validates and interprets AI insights for action
SOAR Executes Playbook	automates standard response actions
AI Learns from Outcome	integrates results to improve future detections

This model preserves human judgment in the loop while scaling efficiency — enabling analysts to focus on strategic decisions rather than repetitive triage.

Infosys CyberSecurity has begun leveraging the AI capabilities as part of the AI-first SOC to reduce false positives, enable faster threat validation and remediation, and scale security engineering without linear headcount growth.

Case Insight: Infosys SOC Evolution

Infosys' Global Cyber Defense Centers (CDCs) operate at an immense scale, managing billions of security events and terabytes of telemetry data each month. These operations span a diverse set of businesses across industries including banking, healthcare, utilities, telecom, manufacturing, government, education, food and beverages, personal care and more.

To address the growing complexity and limitations of traditional rule-based detection systems, Infosys CyberSecurity initiated a strategic transition toward an AI-augmented, use case–driven SOC model. This shift is aimed at increasing detection precision, minimizing alert fatigue, and aligning threat identification more closely with business context. The journey toward building this AI-first SOC unfolded through a series of structured innovations and operational transformations.

Key Enhancements Made:

1. Use Case Engineering

- AI-aligned detection use cases for business-critical assets
- Mapped alerts to MITRE ATT&CK, TTPs, and risk impact

2. AI in Triage & Enrichment

- Deployed machine learning models to de-duplicate events, prioritize based on threat score + risk score + asset value
- Integrated User and Entity Behavior Analytics (UEBA) to contextualize insider risks

3. Automated Response with SOAR – Managed Protection, Detection and Response

- Reduced manual triage time by using workflow-based response playbooks
- Implemented self-healing scripts for repeatable containment actions

4. Threat Intelligence & Behavior Modeling

- Shifted from IOC to behavior-based detections
- Enabled threat hunting powered by statistical anomaly detection

5. Team Capability Uplift

- Trained SOC engineers and analysts in detection engineering and AI-assisted investigation
- Created dedicated roles for ML-assisted Threat Analysts, Detection Engineering and Use Case developers

Value Delivered

The transformation to an AI-augmented SOC at Infosys CyberSecurity has delivered measurable and strategic outcomes:

Operational Efficiency

- **MTTD/ MTTR** reduced significantly improving detection-response cycles.
- **Analyst effort optimized** through auto-triage, freeing time for deep-dive investigations.
- **Alert fatigue reduced** via AI-based correlation.

Strategic Business Alignment

- Use cases mapped to **business-critical services and SLAs,** increasing alignment with business priorities.
- Enabled **risk-weighted alerting**, helping prioritize threats that could impact compliance or service uptime.

Client Experience

- Enhanced visibility through **custom dashboards and executive reports,** integrating operational + business risk KPIs.
- Reduced incident backlogs and provided **faster resolution paths,** improving stakeholder confidence.

Talent Shift & Training

With AI enabling faster operations and scaled responses, the roles within the SOC evolved significantly.

From Operators to Engineers:

- Tier 1 analysts were upskilled in detection logic, threat modeling, and automation design
- Introduced roles such as:
    - o Detection Engineers
    - o Use Case Developers
    - o AI-Assisted Threat Analysts

Structured Training Programs:

- AI & SOAR training modules rolled out to SOC engineers and analysts
- Internal knowledge portals created to share detection engineering patterns, use case playbooks, and threat analysis examples

This focus on **capability maturity** transformed the Cyber Defense Centers (CDCs) into a true **Center of Intelligence**, not just alert handling.

Strategic Recommendations

The following SOC Maturity Model outlines the path to an AI-first SOC.

SOC Maturity Model chart

Level 1 – Reactive	Alerts handled manually, limited visibility, reactive posture
Level 2 – Defined	Rule-based correlation with static logic and limited enrichment
Level 3 – Enriched	Integration of threat intelligence, asset context, and basic prioritization
Level 4 – Adaptive	AI-assisted scoring, playbook automation (SOAR), and behavior-based insights
Level 5 – Predictive	Autonomous detection using behavior modeling, predictive analytics, and proactive threat hunting

The above will be aligned with Secure by Design, Secure by Scale and Secure the Future of Cybersecurity model.

What This Means for Businesses

This paper isn't just about technology evolution — it's a blueprint for **how Infosys delivers cyber defense outcomes in the real world.** For security leaders, CISOs, and transformation executives, the takeaways are clear:

Accelerated Detection & Response:

Improvement in MTTD and MTTR using AI-powered use case logic and automation.

Reduced Analyst Load:

AI filters the noise, allowing SOC engineers to focus on true threats and high-value detection engineering — not tickets

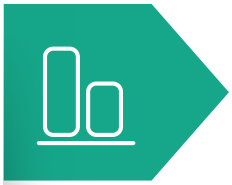

Business-Aligned Alerts:

Threats are scored and prioritized based on your unique business risks.

Scalable SOC Models:

From BFSI to utilities to other sectors, the AI-augmented SOC model is repeatable, predictive, adaptive, and business-specific.

Transparent Metrics:

Real-time visibility into SOC performance, risk reduction, and value delivered to your board.

Conclusion

The journey toward AI-Augmented SOCs isn't just about adopting new technology — it's about fundamentally changing how cybersecurity operations scale, think, and deliver value. Infosys' internal evolution demonstrates that with the right mix of automation, intelligence, and skilled talent, AI-first SOCs can go beyond detection to become engines of business resilience and trust.

At Infosys, we don't just modernize SOCs — we partner with enterprises to reimagine cyber resilience for the digital age. Our global Cyber Defense Centers (CDCs), detection engineering teams, and AI-augmented playbooks are already powering next-gen threat response for businesses across various sectors.

References and Further Reading:

1. Gartner, "Market Guide for MDR Services," 2023.
2. MITRE Corporation, "ATT&CK Framework," https://attack.mitre.org/
3. NIST SP 800-61 Rev. 2, "Computer Security Incident Handling Guide," 2023.
4. SANS Institute, "2023 SOC Survey: Evolving the Analyst Role," https://www.sans.org
5. CISA, "Understanding SOAR Integration," www.cisa.gov/ai-soc
6. The Future of SOC Operations: Autonomous Cyber Defense with AI and Machine Learning, Noman Ali, Gabrielle Wallace
7. SOC Optimization Through AI-Powered Automation and Blockchain Integration, Abhishek Sharma, Rudolf Spunda
8. The Rise of Cognitive SOCs: A Systematic Literature Review on AI Approaches, Farid Binbeshr1 , Muhammad Imam 1,2, Mustafa Ghaleb 1 , Mosab Hamdan1,4 (Senior Member, IEEE), Mussadiq Abdul Rahim 1 AND Mohammad Hammoudeh 3 (Senior Member, IEEE)
9. AI-Enhanced SOC Operations: From Threat Intelligence to Automated Mitigation, Subhan Gill, Alexander Noah

About the Author

Mahesh Gopalakrishnan

Principal Consultant – Cybersecurity, Infosys Americas

Mahesh is an award-winning cybersecurity leader with over 18 years of experience spanning banking, healthcare, utilities, telecom, and IT services. He specializes in building and transforming Security Operations Centers (SOCs), with a strong emphasis on integrating artificial intelligence to enhance threat detection, triage, and operational efficiency. A certified CISSP, CCSP, CISM, CRISC, CGRC, CEH, ISO 27001 Lead Auditor, and a global beta cohort recipient of ISACA's Advanced in AI Security Management (AAISM) certification, Mahesh brings deep expertise in aligning cybersecurity with business resilience. He has successfully led initiatives aligned with MITRE ATT&CK, NIST, PCI-DSS, GDPR, and RBI frameworks, and is passionate about advancing AI-driven SOC capabilities. As the author of Book Cyber *Resilience for Entrepreneurs* and the whitepaper *The Language of Information Security*, Mahesh contributes to global cybersecurity thought leadership through writing and speaking engagements, including ISACA chapter events, CSA Bangalore, awareness programs, and certification training sessions. His approach blends technical depth with strategic vision, making him a valuable voice in discussions on AI adoption, workforce readiness, and governance in modern SOC environments.

For more information, contact askus@infosys.com

Infosys®
Navigate your next