

Tuesday, June 2

2:20 pm – 2:40 pm

Advancing Security Operations Through Agentic AI: Balancing Automation with Human Oversight

Mark Maglin

VP DoW Cybersecurity Strategy

Everforth ECS Federal

Abstract:

As organizations face exponentially growing alert volumes and increasingly sophisticated cyber threats, the traditional SOC model has reached its practical limits. This Innovation Showcase demonstrates how ECS Federal successfully deployed agentic AI to transform our security operations while maintaining critical human oversight and decision-making authority.

ECS Federal's 24/7, multi-tenant SOC analysts were processing more than 2,500 alerts per analyst per month, making a hire-to-keep-up model operationally and economically unfeasible. The organization faced three critical operational challenges:

- **Limits of Existing Automation:** ECS had maximized the practical limits of its SOAR technology. While effective for enrichment and response actions, SOAR lacked the ability to reason through complex investigations or safely auto-close benign alerts
- **Alert Overload and Analyst Fatigue:** Benign alerts consumed disproportionate analyst time, displacing focus from genuine threats, leading to fatigue and operational inefficiency
- **Inconsistent Manual Triage:** Human-driven triage produced variable outcomes requiring standardized, auditable processes

The Agentic AI Solution

Rather than replacing human analysts, ECS Federal implemented Agentic AI agents as an autonomous investigative layer designed to provide analysts with leverage. The solution

conducts deep, contextual investigations while delivering decision-ready investigation reports with clear conclusions and supporting evidence.

ECS integrated AI Agents directly into its existing workflows moving alerts from detection systems through SOAR, to the AI Agent for autonomous investigation, with outcomes fed back into ServiceNow for centralized case management and auditability.

Defending Against AI-Enabled Offensive Attacks

AI is increasingly being leveraged by adversaries to craft more sophisticated and automated attacks, including highly convincing phishing campaigns, advanced malware, and even autonomously executed cyber warfare.

Agentic AI Defenses address these threats by:

- Matching speed and scale
- Deep contextual Investigation which is essential when AI powered attacks adapt their behavior to avoid detection.
- Consistent, Auditable Analysis

Human-in-the-Loop Safeguards

Critical to this implementation is the preservation of human decision-making authority. The AI Agent system enhances analyst effectiveness by standardizing triage in a repeatable, auditable manner while ensuring transparency. Each alert arrives with clear conclusions, supporting evidence, and rich context, improving trust and decision quality

Results

Tier 1 Alert Auto-Closure Rate:70% of alerts automatically closed safely without analyst involvement

Investigation Response Time:100% of alerts investigated immediately upon creation

Scalability Achievement: Broke the SOC scalability ceiling

Headcount Efficiency: Scaled SOC operations 20% without adding headcount

Qualitative Outcomes

- Scalable solution to combat AI generated attacks.
- Improved Analyst Effectiveness and Retention: No negative impact to SLA. With queues cleared and workloads more manageable, analysts experienced less burnout and greater job satisfaction

- Enhanced Strategic Focus: focus on higher-value work like threat hunting, detection engineering
- Investigation Depth: Context and evidence gathering exceeded what analysts would typically perform during initial triage

Conclusion

The ECS Federal case demonstrates that agentic AI, when implemented thoughtfully with human-in-the-loop controls, can dramatically enhance security operations without sacrificing the judgment and accountability that only human analysts can provide. This phased approach—from virtual agent implementation through autonomous task execution—positions organizations to achieve significant operational improvements while maintaining essential human oversight