GSDC
Global Skill Development Council

GENERATIVE AI CYBERSECURITY

Real World Case Studies

www.gsdcouncil.org

USE CASE 1

How Google Cut Vulnerability Response Time from Days to Minutes

The Problem

Security teams at large organizations face an overwhelming volume of vulnerability disclosures daily. For Google's Project Zero team, the manual analysis process — reading disclosures, understanding exploit mechanisms, assessing affected systems, and generating remediation plans — was consuming days of expert analyst time per vulnerability, creating dangerous windows between disclosure and remediation.

How Gen AI Was Used

Google integrated Gemini AI models into its internal vulnerability management workflow. When a new CVE is published, the AI automatically reads the full technical disclosure, generates a plain-English explanation, maps affected systems across Google's infrastructure, cross-references active exploitation evidence, and produces a prioritized remediation recommendation — all before a human analyst begins manual review.

Results

- Vulnerability triage time reduced from days to under an hour for standard CVEs
- AI correctly identified exploitability status ahead of manual analysis in the majority of cases
- Significantly higher disclosure volumes processed without increasing headcount
- AI-assisted research identified zero-day vulnerabilities in real-world code — a landmark Gen AI demonstration

Triage First

Gen AI's most immediate value is compressing the time from CVE publication to informed remediation decision.

Human + AI

The combination of AI speed and human expert judgment — not AI alone — produced the best outcomes.

Window Advantage

AI triage workflows create a measurable window advantage over attackers exploiting the gap between CVE publication and enterprise patching.

USE CASE 2

How Microsoft Used Security Copilot to Reduce SOC Investigation Time by 40%

The Problem

Microsoft's own SOC — one of the world's most sophisticated — saw experienced analysts spending 60–70% of investigation time on data gathering and correlation tasks rather than analysis. Junior analysts faced an even steeper challenge, often spending hours on investigations senior analysts completed in thirty minutes due to experience gaps.

How Gen AI Was Used

Microsoft deployed Security Copilot internally across Sentinel, Defender, and internal threat intelligence platforms. When an alert fires, Copilot auto-generates a complete incident summary — correlated events, affected assets, TTPs, and relevant intelligence — before the analyst opens the ticket. Analysts investigate through natural language conversation rather than navigating multiple tool interfaces.

Results

- Investigation Time Reduced by 40%
- 60% Faster Report Generation
- 25% More Volume, Same Team

Investigation Is Where AI Wins

The biggest SOC efficiency gain from AI is not in detection — it is in the investigation phase where human time is most consumed.

Narrows the Experience Gap

AI elevates junior analyst output to senior analyst quality — a critical insight given the global cybersecurity talent shortage.

Measure the Right Metric

Organizations that only track detection metrics miss where AI delivers the most operational value — investigation time.

USE CASE 3

How Mastercard Used AI to Detect Fraud Across 143 Billion Transactions

The Problem

Mastercard processes over 143 billion transactions annually — a volume making meaningful human review impossible. Traditional rule-based fraud detection generated unacceptably high false positive rates while missing sophisticated fraud patterns. Organized fraud rings were actively probing detection boundaries to identify rules they could evade.

How Gen AI Was Used

Mastercard deployed Decision Intelligence Pro — an AI system using recurrent neural networks and generative AI techniques that analyzes each transaction against a cardholder's complete behavioral history. The system generates a dynamic behavioral model per cardholder and scores each transaction in real time in under 50 milliseconds. Gen AI components synthesize novel fraud pattern variants for training before they appear in live data.

Key Results

- **300%** Improvement in fraud detection vs. prior rule-based systems
- **50ms** Real-time scoring speed — every transaction evaluated without impacting customer experience
- **210+** Countries now protected by the deployed system globally
- **1.3T** Data points processed annually across the cardholder base

Behavioral AI vs. Rules

Behavioral AI outperforms rule-based systems because it personalizes detection to each individual rather than applying universal thresholds.

Proactive Training

Using Gen AI to synthesize novel fraud patterns for training before they appear in real transactions is a model for proactive security AI development.

False Positives Matter

A system that blocks too many legitimate transactions creates business damage of its own — false positive reduction is as important as detection rate.

USE CASE 4

How Darktrace Detected a Novel Supply Chain Attack at a European Manufacturer

The Problem

An attacker compromised a trusted software vendor's update server and used the legitimate update mechanism to push malicious code. Because the attack arrived through an authenticated, trusted channel, every traditional control failed — signature-based AV saw clean files, firewall rules permitted the connection, and the attack was entirely invisible to rule-based defenses.

How Gen AI Was Used

Darktrace's self-learning AI had established behavioral baselines for every device including the specific update process patterns. When the compromised update executed, it deviated subtly — making network connections to infrastructure it had never contacted, accessing directories outside its normal scope, and spawning inconsistent child processes. Darktrace detected these deviations within minutes. Antigena autonomously slowed suspicious connections, buying time for human investigation.

Results

- **Attack detected within minutes:** Before any data was exfiltrated or lateral movement occurred — every traditional control in the environment had failed.
- **No signature existed:** Post-incident forensic analysis confirmed this was a previously unknown supply chain technique — no signature could have existed.
- **Multi-week dwell time avoided:** Forensic analysis determined the manufacturer avoided what could have been weeks of attacker persistence under traditional detection.

Rapid Detection

The attack was detected within minutes, preventing data exfiltration or lateral movement before traditional controls could react.

Novel Attack Protection

Behavioral AI successfully identified a previously unknown supply chain technique that signature-based tools could not detect.

Resilience to Supply Chain Threats

This case highlights how behavioral AI provides crucial detection against supply chain attacks, exploiting trust relationships that traditional rule-based security often respects.

USE CASE 5

How JPMorgan Chase Generated 3.5 Billion Cybersecurity Rules with Gen AI

The Problem

JPMorgan Chase operates one of the world's most complex and heavily targeted technology environments. Protecting it requires an enormous volume of detection rules, security policies, and monitoring configurations — far exceeding what human security engineers can write, maintain, and update manually. As the threat landscape evolved, the gap between the rules the organization had and the rules it needed was growing continuously.

How Gen AI Was Used

JPMorgan Chase developed an internal Gen AI system — built on LLMs trained specifically on cybersecurity data — to generate, validate, and maintain security rules at scale. The system analyzes threat intelligence feeds, observes patterns in internal telemetry, and automatically generates SIEM detection rules, firewall policies, and monitoring configurations that engineers review and approve. It also identifies outdated or redundant rules and proposes deprecations.

Results

- **3.5 Billion** cybersecurity rules generated using AI — a volume that is orders of magnitude beyond human authorship capacity.
- **Improved Detection Coverage:** AI-generated rules keep pace with novel attack techniques in near-real-time.
- **Human Role Shift:** Engineers moved from rule creation to rule review, focusing on higher-value activities.
- **Enhanced Consistency:** AI-generated rules follow standards more consistently than rules authored by large engineering teams.

Scale Humans Cannot Reach

The value of AI in security is not doing what humans do faster — it is doing things at a scale humans cannot reach at all.

Review Over Creation

Humans add more value validating AI output than generating first drafts — an important organizational design insight.

Sustainable Coverage

Organizations using AI to maintain rule libraries will have broader, more current coverage than those relying on human-authored rules alone.

USE CASE 6

How a Global Bank Used Gen AI to Automate AML Transaction Monitoring

The Problem

A major global bank was processing millions of transactions daily through a traditional rule-based AML monitoring system generating tens of thousands of suspicious activity alerts daily. Over 95% were false positives. Analysts were spending the overwhelming majority of their time investigating legitimate transactions while potentially missing genuinely suspicious patterns buried in the noise.

How Gen AI Was Used

The bank deployed a Gen AI-powered AML system replacing static transaction rules with dynamic behavioral modeling. The AI builds behavioral profiles per account — understanding normal transaction amounts, frequencies, counterparties, geographies, and timing — then scores deviations for genuine risk. Gen AI components generate narrative explanations of why each flagged transaction is suspicious, in language compliance analysts can directly use in SAR filings.

Results

- **60% Fewer False Positives:** Analysts reviewing genuinely suspicious patterns rather than noise.
- **2x Analyst Capacity:** Same team handling twice the genuine investigation volume.
- Improved regulatory examination outcomes due to higher quality and consistency in SAR documentation.
- Avoided significant regulatory fines that peers operating less effective AML programs received.

False Positives Are a Systemic Failure

The false positive problem in AML is not a minor inconvenience — it causes genuine suspicious activity to be missed because analysts are overwhelmed with noise.

Narratives Over Scores

Gen AI's ability to generate narrative explanations — not just scores — is uniquely valuable in compliance contexts where documentation quality is a regulatory requirement.

Highest ROI Use Case in Banking

AML is one of the highest-ROI cybersecurity AI use cases due to direct regulatory cost of failures and the severity of the false positive problem with traditional approaches.

USE CASE 7

How CrowdStrike Used Gen AI to Stop a Nation-State Attack in Real Time

The Problem

A major technology company was targeted by a nation-state actor using living-off-the-land techniques — abusing legitimate Windows tools (PowerShell, WMI, PsExec) rather than deploying custom malware, making detection through traditional signature-based endpoint tools essentially impossible. The attack began with a successful spear phishing compromise of a developer's workstation.

How Gen AI Was Used

CrowdStrike Falcon's behavioral AI detected anomalous combinations and sequences of legitimate tool usage inconsistent with normal administrative behavior — despite zero malicious files. Charlotte AI provided SOC analysts a real-time natural language investigation interface. Within the first hour, the AI identified the attacker's behavioral pattern as consistent with a specific named nation-state group, drawing on CrowdStrike's threat intelligence.

Results

- **Real-time Detection:** Attack detected by behavioral AI within 2 hours, before lateral movement to sensitive systems.
- **Nation-state Attribution:** AI identified the attacker's behavioral pattern as a specific nation-state group within 2 hours+.
- **Zero Malware Deployment:** The attack was stopped without any malicious files being deployed.
- **Rapid Containment:** Full containment executed within 4 hours, significantly faster than the 200+ day industry average dwell time.

Living-off-the-Land Is the New Normal

Sophisticated attackers use legitimate tools specifically to defeat signature-based detection — behavioral AI is the primary defense.

Attribution Improves Response

Knowing who is attacking you enables security teams to apply intelligence about that actor's known objectives and next likely moves.

4 Hours vs. 200 Days

AI-accelerated response against a nation-state actor — a stark illustration of the gap between AI-assisted and traditional SOC models.

USE CASE 8

How a Healthcare System Used Gen AI to Prevent Ransomware Before Encryption Began

The Problem

A large US healthcare system operating dozens of hospitals detected suspicious activity. The SOC team, overwhelmed by alert volume, had not yet begun investigating the specific alerts that were early indicators of an active ransomware campaign. In healthcare, ransomware that encrypts clinical systems can delay treatment and has contributed to patient harm in documented cases.

How Gen AI Was Used

The organization's AI-powered security operations platform continuously monitored all network and endpoint telemetry. The AI correlated 7 separate low-priority alerts — an unusual authentication event, shadow copy service access, unusual lateral movement between clinical workstations — into a single high-priority incident with the AI-generated assessment that these behaviors were consistent with pre-ransomware staging. Automated containment isolated the staging systems within 4 minutes of detection.

Results

- **No Encryption:** Ransomware staging contained before a single file was encrypted.
- **No Disruption:** Hospital operations and patient care continued without interruption.
- **Rapid Containment:** Automated containment executed in 4 minutes, faster than any human response could have achieved.
- **Pre-emptive Action:** Prevented encryption that would have begun within 6–8 hours during overnight low-staffing, when response would have been slowest.

Correlation Capability is Key

AI's ability to connect multiple low-priority alerts into a single critical incident is a valuable and often underappreciated capability.

AI Surpasses Human Analysis

Human analysts reviewing alerts in sequence often miss patterns that AI sees across the full dataset simultaneously.

Pre-encryption Detection

Behavioral AI is crucial for preventing ransomware before encryption begins; by the time it's active, damage is already in progress.

USE CASE 9

How Abnormal Security Stopped a $36 Million BEC Attack with Behavioral AI

The Problem

A large financial services organization was targeted by a sophisticated Business Email Compromise attack. The attacker spent weeks in reconnaissance — studying email patterns, mapping vendor relationships, identifying wire transfer authorizers. They crafted a highly convincing thread impersonating a senior executive and known vendor, requesting a $36 million wire transfer. The email was grammatically perfect, contextually accurate, and contained zero traditional phishing indicators — no malicious links, no attachments, no suspicious domains.

How Gen AI Was Used

Abnormal Security's behavioral AI analyzed the email against multiple dimensions simultaneously: the executive had never used this communication style for financial requests; the wire transfer initiation pattern was inconsistent with established norms; the vendor's communication pattern deviated from historical behavior; urgency language was statistically anomalous; and the request bypassed the organization's normal multi-party approval process. The AI flagged the email before any human read it, with a detailed explanation of each behavioral anomaly.

Results

- **$36 Million Saved:** Wire transfer blocked before execution, preventing any financial loss.
- **No Compromise:** Zero traditional technical indicators of compromise detected by traditional tools.
- **Traditional Tools Ineffective:** Content-based email security produced zero detections; behavioral AI was the only mechanism that worked.
- **Pre-delivery Detection:** The attack was detected and eliminated before delivery, removing the human judgment point entirely.

Behavioral AI is Critical

Traditional content-based security tools missed all indicators; behavioral AI was the only effective defense against this sophisticated BEC attack.

Eliminates Human Judgment

Detecting and blocking sophisticated BEC attacks before delivery eliminates the need for human judgment, which is often unreliable against such convincing fakes.

Resilient to Reconnaissance

Attackers invest heavily in reconnaissance, but behavioral AI, profiled at the relationship level, remains resilient even against these advanced preparation tactics.

USE CASE 10

How the US Department of Defense Deployed Gen AI for Autonomous Cyber Defense

The Problem

The DoD operates one of the world's most targeted network environments — spanning classified and unclassified systems, OT, weapons platforms, and logistics networks across hundreds of global locations. Nation-state adversaries were conducting attacks that moved through networks faster than human incident response could track. The DoD needed a defense capability that operated at machine speed against machine-speed attacks.

How Gen AI Was Used

DARPA and DoD cyber agencies developed AI-powered autonomous cyber defense systems through programs including the Cyber Grand Challenge and subsequent operational deployments. These systems continuously monitor network behavior, autonomously identify attack patterns, generate and deploy defensive countermeasures in real time, and adapt to novel techniques without human intervention. CISA also deployed AI tools for continuous monitoring of critical infrastructure with AI generating prioritized alerts and recommended responses.

Results

- **Seconds, Not Hours:** Autonomous AI detects and responds to threats at speeds measured in seconds, versus hours or days for human-staffed responses.
- **Coverage Gaps Closed:** AI-powered continuous monitoring identified critical exposures that manual, periodic assessments missed.
- **Cyber Grand Challenge:** Demonstrated AI systems autonomously finding, patching, and exploiting vulnerabilities, now becoming operational reality.
- **Core Strategy:** DoD established AI as a core component of its cyber defense strategy with ongoing investment across all service branches.

Machine-Speed Defense is Essential

Human-staffed defense operating at human speed is insufficient against sophisticated, machine-speed nation-state adversaries.

AI is a Defense Prerequisite

AI enables offense at scale, making AI-powered defense a necessity, not just an enhancement, for modern cyber defense.

Autonomous with Human Oversight

The future security operations model will shift dramatically towards AI autonomy with human oversight, requiring early transition.

GSDC
Global Skill Development Council

CERTIFICATION IN GENERATIVE AI IN CYBERSECURITY

ABOUT GSDC CERTIFICATION

EBOOK

Extensive and exclusive Ebook created by world's experts to help you with understanding core concepts.

LEARNING MATERIALS

Get access to learning materials such as videos, ebooks, templates, and practice exams, which will help you clear the certification exam.

CREATED BY EXPERTS

GSDC certifications are created and authored by world's leading experts in the field.

LEARNING OBJECTIVE

- Gain insights into autonomous decision-making processes
- Apply knowledge using ready-to-implement templates
- Demonstrate ability to work with Agentic AI models
- Validate your skills wit

Enroll now with the code **LEARN20** To avail **20%** discount

Enroll Now

www.gsdcouncil.org