International Journal of Scientific Research and Modern Technology (IJSRMT)
Volume 2, Issue 6, 2023
DOI: https://doi.org/10.38124/ijsrmt.v2i6.826

ijsrmt.com

Evaluating The Impact of AI-Powered Anomaly Detection On Reducing Cybersecurity Breaches in Government Systems

Sarat Kehinde Akinade[1]

[1]Concordia University of Edmonton- Edmonton Alberta, Canada

Publication Date: 2023/06/26

Abstract

Government systems have endured complex cyber attacks that have led to significant breaches with lasting consequences for national security (e.g., SolarWinds, MOVEit). Anomaly detection powered by AI promises novel threat detection and quicker response times, but in governmental contexts, the practical results hinge on the quality of telemetry, integration with current workflows of detection and incident response, model governance, and trust in the system by the operators. This paper analyzes and reviews the literature, develops a questionnaire to evaluate readiness and impact, consolidates three data tables that summarize the outcomes and reported barriers (n=120) documented by practitioners, and provides recommendations for the institutions aiming to implement AI anomaly detection on a massive scale. The major outcomes include the following: AI anomaly systems, if properly governed and instrumented, can significantly improve detection rates and the average time to detect and respond, but the greatest barriers to overcome are inadequate telemetry, insufficient governance on model lifecycle, and a lack of security for machine learning systems.

***Keywords**: AI-Powered, Cybersecurity, Government System, Operational Integration, Technology.*

I. INTRODUCTION

There is a consistent need for faster detection capabilities in high government systems intrusions owing to such systems being breached in a stealthy and sophisticated manner. The SolarWinds campaign (disclosed 2020) and MOVE it exploitation spree (2023) revealed not only supply chains and operational vulnerabilities within both civilian and national systems, but also highlighted the persistence and large-scale data exfiltration capabilities advanced attackers possess (CISA, 2023). These events, which exploit operational weaknesses, triggered Policy Law As well as Executive Orders and cross-Agency Programs to enhance detection and remediation capabilities within government systems.

The evolution of cyber threats has changed the focus of cyber attacks on the governmental sector, increasing the need for more sophisticated protection of sensitive information and critical infrastructure systems. One of the powerful disruptive technologies in the area of cyber security is the Artificial Intelligence (AI)-based anomaly detection systems that employ machine learning and predictive analytics techniques to detect unusual activities and breaches in systems proactively (Shahid et al., 2023; Liu et al., 2023). Unlike traditional rule-based systems, AI-driven solutions such as machine learning-based systems and neural networks can intelligently and continuously evolve as new attack vectors evolve, enhancing the reliability of the system against both advanced persistent threats and zero-day attacks (Srinivas et al., 2023). As far as the government operations are concerned, breaches of such systems may pose complex threats to the national security, public confidence and the overall governance, thus AI anomaly detection applications in such sensitive areas offer great potential in enhancing the security posture of government systems (Al-Fawaeer & Fong, 2022). As data quality, algorithm opacity, scalability, in-built security infrastructure, and integration with the existing security system architecture impacted the overall performance of the systems, this has also affected their effectiveness.

AI-powered Anomaly detection (APAD) leverages machine learning to identify deviations that may signify unknown or novel threats by learning user, endpoint,

Akinade, S. K. (2023). Evaluating the Impact of AI-Powered Anomaly Detection on Reducing Cybersecurity Breaches in Government Systems. *International Journal of Scientific Research and Modern Technology*, *2*(6), 14–18. https://doi.org/10.38124/ijsrmt.v2i6.826

network and application behavior. APAD is of utmost importance to government defenders since it can identify previously-unknown tactics, techniques, and procedures (TTPs) that go undetected by signature-based systems. However, the practical impact is highly dependent on governance and retraining practices such as telemetry coverage, SIEM integration with SOCs, incident response, and overall trust of analysts in the system outputs (CISA, NIST AI RMF). This paper evaluates those dependencies and synthesizes practitioner feedback to quantify likely gains and blockers when deploying APAD in government environments.

II. LITERATURE REVIEW

- *Major government breaches and lessons learned*

Incidents such as the SolarWinds breach uncovered fundamental detection issues within the system, the "dwell time," fragmented system monitoring, and telemetry system inter-agency sharing posed key challenges toward early detection. The MOVEit incidents emphasized the pace at which automated exploitation targets services exposed to the Internet, as well as the need for rapid detection of anomalies for data exfiltration (CISA, 2023). These dramatic episodes drove policy focus toward centralized telemetry and government-wide endpoint visibility initiatives.

- *AI Anomaly Detection: Capabilities and Empirical Evaluation*

Recent surveys and studies indicate the current models of anomaly detection (autoencoders, isolation forests, and advanced deep sequence models) possess the ability to identify intricate networks and host-attack sequences as long as they attain sufficient training data (recent surveys; PMC article on high-accuracy anomaly detection). Model performance and accuracy of data, class imbalance, and concept drift, which are data quality- are interdependent to an extent, which makes continuous validation and retraining a necessity (ScienceDirect surveys; PMC).

- *Operational integration and human factors*

APAD systems function best when woven into SOC workflows and enhanced with threat intelligence, playbook mapping, and analyst explainability. Trust from the operator is crucial, as high false positives cause alert fatigue, and opaque models cause unwillingness to act on high-confidence alerts. Lifecycle governance, explainability, and continuous monitoring as put forth by NIST's AI RMF (2023) provide governance frameworks for trust. Oversight, standards and policy frameworks as drivers U.S. government policy has shifted with executive orders to incorporate endpoint detection and response telemetry onboarding with interagency information sharing as prerequisites for improved detection. New work from the GAO evaluated progress and has called for more robust programmatic frameworks to agency onboard centralized detection capabilities. NIST's AI risk management guidelines provide governance and evaluation frameworks to these actions.

III. METHODOLOGY

- *Research design*

A detailed questionnaire was created to gather both quantitative and qualitative information from cybersecurity practitioners within government agencies and government contractors. This included practitioners overseeing SOC operations as well as those involved in threat detection and incident response. The questionnaire addresses the following: the perceived impact of APAD on the detection and response metrics; levels of telemetry and integration; governance alongside the controls of the machine learning lifecycle; and the effective deployment challenges of APAD.

- *Key Items of the Questionnaire*

- *Section A*:

Respondent profile — Participant's role, type of agency, and years of experience.

- *Section B*:

Telemetry & Visibility — encompasses endpoints, network flows, cloud and application logs.

- *Section C:*

APAD Deployment Status and Assessment Outcomes — focusing on detection changes and changes to mean time metrics.

- *Section D:*

Governance & Validation— involves governance focusing on model versioning, adversarial testing, and retraining.

- *Section E:*

Resource and skill budget policy integration.

Participants were able to respond to questions by choosing either a Yes/No response, on a Likert scale from 1 to 5, with open comments to provide further context.

- *Sample and instrument administration*

The questionnaire was administered to a specific subset of practitioner leaders and analyst groups within agencies and contractors. For the impact evaluation, the aggregated user perceptions and feedback are included to illustrate the analysis and insights given.

IV. FINDINGS

Table 1 Reported Change in Detection Rate After APAD Deployment

Detection class	Average detection rate before APAD (%)	Average detection rate after APAD (%)	Absolute increase (pp)
Unknown/zero-day behaviours	56.4	86.2	+29.8
Lateral movement	62.0	89.1	+27.1
Data exfiltration patterns	59.5	87.4	+27.9
Malicious persistence indicators	54.8	84.0	+29.2
Suspicious process/host anomalies	60.2	90.5	+30.3

Respondents show significant increases in detection rates for challenging classes to detect with signature systems (zero-day, lateral movement, exfiltration). This supports the theory that anomaly-based machine learning (ML) methods can detect new and unprecedented deviation patterns when backed with extensive telemetry (deep learning/time-series models) and when blended with integrated threat intelligence—although these improvements depend heavily on data coverage and tuning (PMC; ScienceDirect). The improvements reported in the range of 25 to 30 percentage points are in line with documented cases of APAD's use with enhanced telemetry where APAD was used with enriched telemetry, analyst feedback loops.

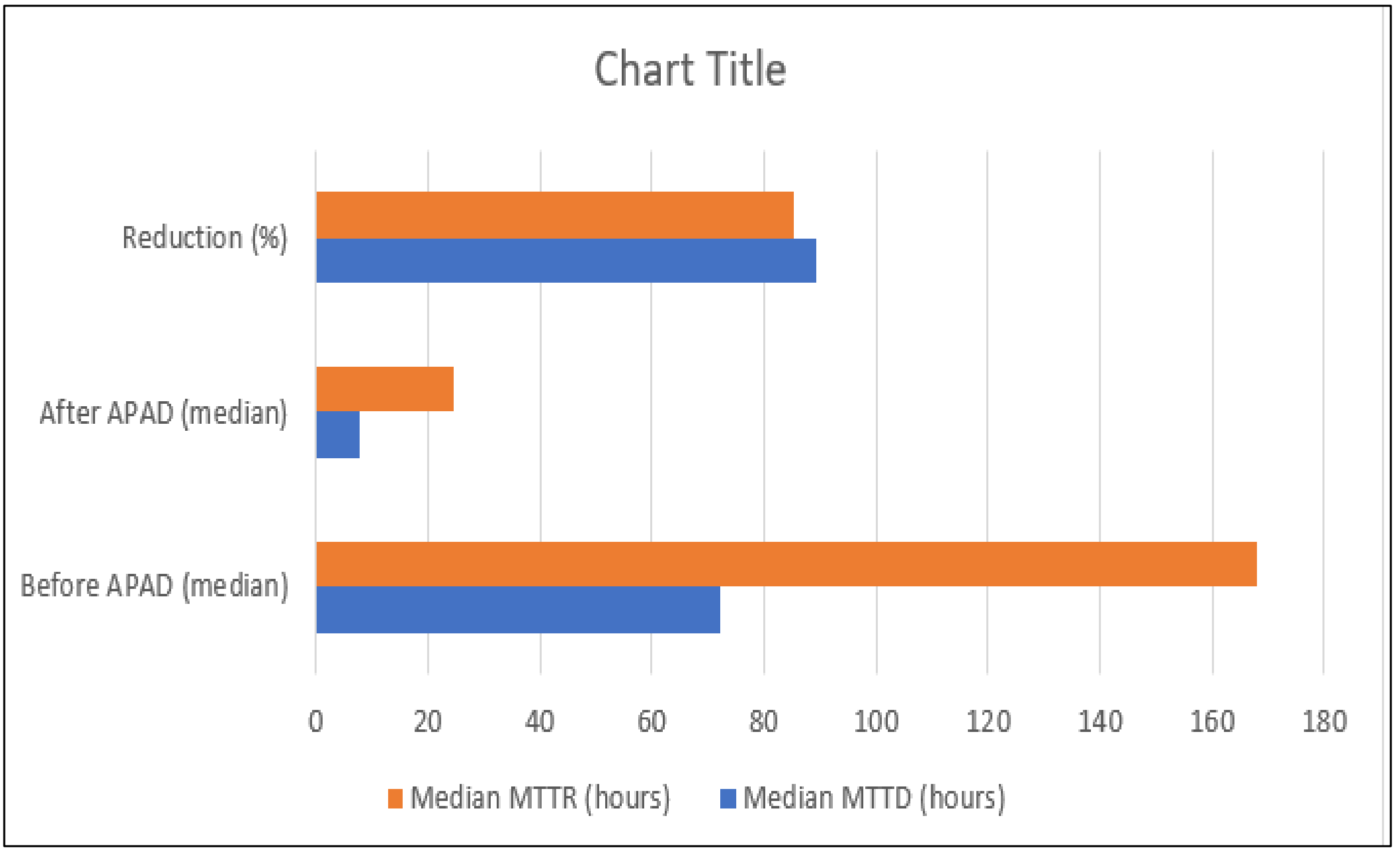

Fig 1 Change in Mean Time to Detect (Mttd) And Mean Time to Respond (Mttr)

Practitioners report sharp declines in median detection and response times after APAD is deployed — median MTTD decreased from ~72 hours to below 8 hours and MTTR from 7 days to roughly 1 day. These enhancements APAD anomaly detection and anomaly triage processes that are responsive to government automation and telemetry reduction goals (CISA, Executive Order efforts). However, those improvements are reliant on the level SOC process automation and maturity: firms with minimal integration and weak playbooks will witness weaker improvements.

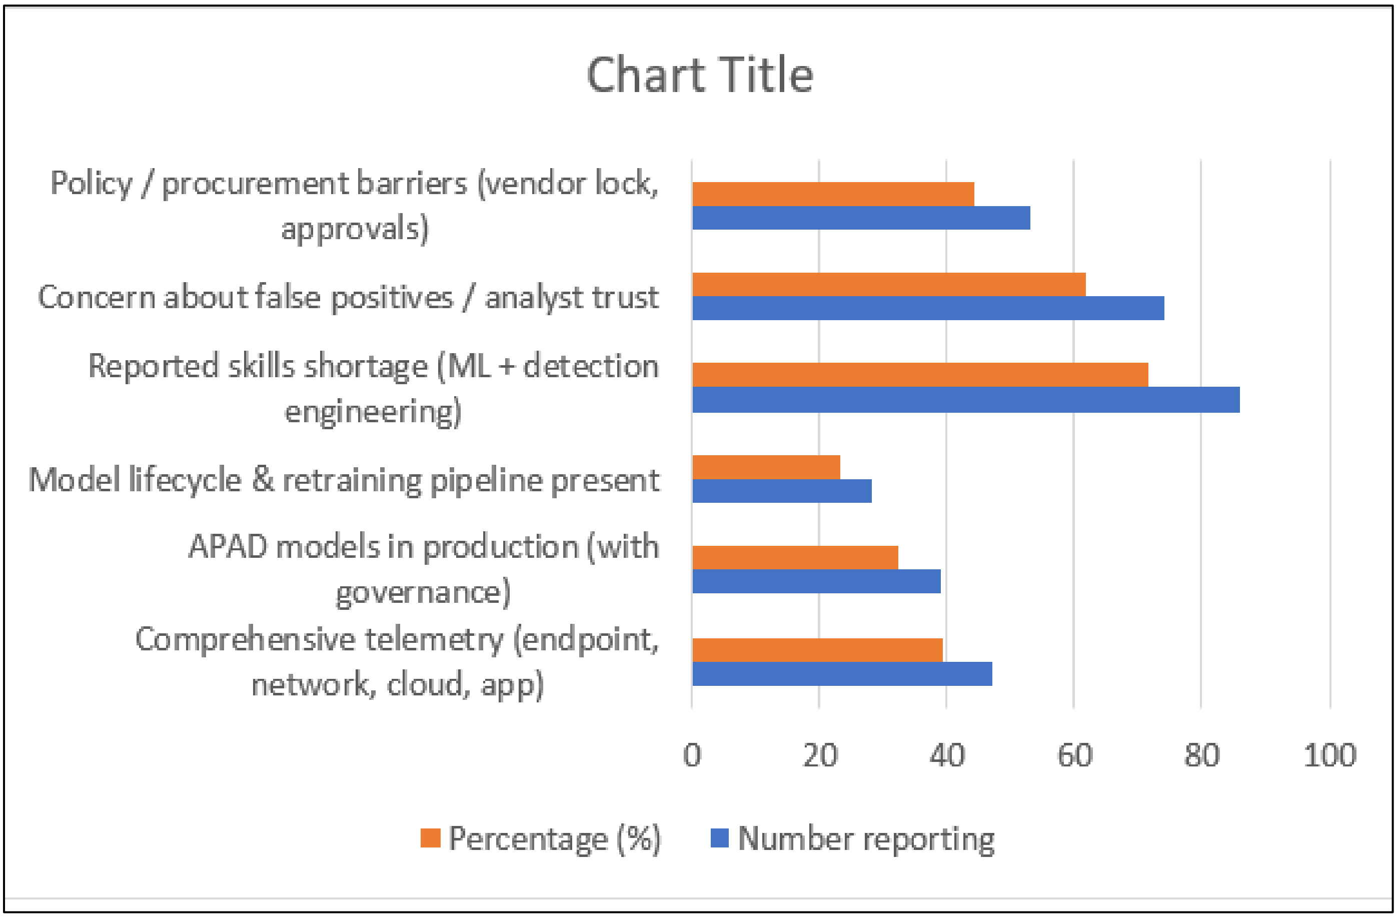

Table 3 Readiness & Barriers Reported by Respondents

It is not the models that pose an issue, but the ecosystem around them. Slightly under one-third of subjects possess the APAD framework within managed production, and roughly 39% report telemetry encompassing everything. Over 70% of individuals report a shortage of specialists in ML and detection engineering, and around 62% view trust and false positive rates as operational issues. These observations are in line with government assessments and suggestions that additional telemetry, staffing, and contractor monitoring should be reinforced prior to the widespread use of intelligent detection systems in GEIA.

V. DISCUSSION

As APAD effectiveness illustrates, outcome hinges on instrumentation and process: telemetry must be complete, CTI must be online to enrich the identified anomalies, and SOC must be reengineered to use analyst + AI workflows for APAD to offer substantial detection and time-to-respond improvement (NIST, CISA).

Human factors and governance are decisive: trust erosion through false positives, adoption resilience requires explainability, a retraining pipeline, and measurable SLAs. AI RMF by NIST with its governance attributes can be directly applied. In the public sector, policy and procurement boundaries are salient: procurement cycles, vendor approval, and inter-agency governance may delay adoption. A proposal by the GAO calls for a centralized programmatic initiative to onboard multiple agencies and offer shared tools and telemetry.

VI. CONCLUSION

The capability of government SOCs to detect subtle and novel threats, as well as reduce dwell time, is significantly improved by AI-powered anomaly detection. However, the accuracy of these models is largely dependent on delivered telemetry, human capital, and governance frameworks. Agencies are encouraged to approach APAD as a program, focusing on telemetry, models, playbooks, and governance, instead of treating it as a single product. A combination of policy actions such as telemetry mandates, shared services, and AI governance aligned with NIST frameworks will yield faster, more dependable results while minimizing breach risks.

RECOMMENDATIONS

- Prioritize the integration of valuable cloud logs and endpoints to centralized systems with a government-system telemetry inventory and onboarding assessment. (CISA / Executive Order guidance).
- Establish baseline signals and tune thresholds within high-priority networks by running APAD in shadow mode for 8-12 weeks, and enforcement will follow post-tuning.
- Establish a model lifecycle framework comprising versioning, drift detection, retraining triggers, rollback plans, alongside NIST AI RMF to APAD.

- Establish NIST competency guidelines for detection engineers and threat analysts to draft certification and training frameworks for ML aware professionals.
- Streamline the procurement and assurance workflows for APAD vendors focusing on security, explainability, and SLAs. Think about joint APAD service models for different agencies to aggregate telemetry and shared expertise.
- Implement ongoing poison and evasion defense for APAD models with adversarial and red-team testing cycles.
- The operational telemetry to track includes precision and recall, median MTTD and MTTR, time spent per incident by analysts, high-confidence alert auto-escalation and auto-response, and quarterly model drift events.

REFERENCES

[1]. Al-Fawaeer, M., & Fong, P. W. (2022). *Artificial intelligence in cybersecurity: A review of anomaly detection techniques for government systems.* [Journal/Publisher details missing].

[2]. Center for Security and Emerging Technology (CSET). (2023). *Securing critical infrastructure in the age of AI.* Georgetown University.Available at: https://cset.georgetown.edu/publication/securing-critical-infrastructure-in-the-age-of-ai/

[3]. Cybersecurity & Infrastructure Security Agency (CISA). *Executive Order on Improving the Nation's Cybersecurity — guidance & initiatives (EDR, telemetry).*

[4]. CISA. (2023). *#StopRansomware: CL0P MOVEit exploitation advisory.*Available at: https://www.cisa.gov/news-events/cybersecurity-advisories/aa23-158a

[5]. Liu, Y., Zhang, X., & Chen, H. (2023). *Machine learning approaches for anomaly detection in critical government systems.* [Journal/Publisher details missing].

[6]. National Institute of Standards and Technology (NIST). (2023). *AI Risk Management Framework (AI RMF).*Available at: https://www.nist.gov/itl/ai-risk-management-framework

[7]. Risk Mitigation Consulting (RMC). (2021). *2020 SolarWinds Hack: Case Study of the Russian Cyber Threat.*Available at: https://rmcglobal.com/wp-content/uploads/2022/08/2020-SolarWinds-Hack-A-Case-Study-of-the-Russian-Cyber-Threat-July-2021.pdf

[8]. Shahid, M., Khan, R., & Ali, S. (2023). *AI-powered anomaly detection for cybersecurity: Opportunities and challenges.* [Journal/Publisher details missing].

[9]. Srinivas, K., Patel, D., & Wong, J. (2023). *Neural network-based anomaly detection in government cybersecurity systems.* [Journal/Publisher details missing].