Cyber AI: Federal Use Cases, Securing Data and AI-Driven Insights

Dr. Nnamdi Osia

Created with https://labs.google/fx/tools/image-fx.

Introduction

In today's evolving landscape, artificial intelligence (AI) enables processing of information at speeds far beyond human capability – and speed is everything in cybersecurity. Cyber threats are faster and more frequent because bad actors use AI and the latest technological advances to automate attacks. To keep pace, federal agencies must not only adopt AI and emerging technologies but also strategically integrate them across their respective enterprises.

Cyber AI refers to the use of AI to carry out cybersecurity tasks and functions such as protecting, detecting, responding, and recovering systems from cyber threats and attacks. Cyber AI uses machine learning (ML), pattern recognition, and large language models (LLMs) to automate and allow quick decisions in areas that traditionally require some form of human input. Government agencies can utilize Cyber AI as a preventive measure through the identification of early threats; it also can serve as a reactive solution in response to cyber incidents. Effective Cyber AI implementation increases operational efficiency and enhances security across federal agencies.

Federal IT, Security Operations Centers (SOCs), and cyber teams are under constant pressure, often working in 24/7 environments with limited resources and facing complex, AI-powered threats. At the same time, internal systems must be modernized for continuous updates, patching, and monitoring. In environments where resources are scarce and data must be used for actionable decisions at a moment's notice, AI can make all the difference.

To optimize AI's potential, it is important to incorporate structure and insight in implementation. Agencies must move beyond one-offs like proof of concepts and pilots and toward enterprise approaches incorporating:

- **Strategy:** Establishing formal AI policy, roadmaps, and implementation plans that are aligned to mission goals.
- **Governance:** Employing AI councils and governance bodies to guide responsible development, deployment, and operation of AI-related resources.
- **Communities of Practice:** Creating forums to share lessons learned and encourage collaboration among inter-agency and intra-agency bodies (e.g., conferences, working groups, tiger teams, and integrated project teams).

Cyber AI plays a critical role in safeguarding sensitive data across complex digital environments. Through effective AI implementation, agencies can strengthen data protection and unlock enhanced capabilities ranging from real-time SOC operations to predictive analytics. With mandates like Office of Management and Budget (OMB) Memorandum 25-21, there's never been a more urgent need to align AI with both mission objectives and compliance requirements.

This whitepaper will explore ways to operationalize AI in agencies and provide real-world insights into how Cyber AI can secure data, reduce risk, and reduce the burden on an already stretched cybersecurity workforce.

Created with https://labs.google/fx/tools/image-fx.

Operationalizing Cyber AI: Federal Agency Use Cases

Federal agencies are implementing AI across a broad range of cybersecurity operations, compliance workflows, and mission-support activities. The use cases below highlight some of these implementations and demonstrate how AI can streamline operations, reduce workforce burden, and hasten response times while complying with Federal Information Security Management Act (FISMA) of 2002 requirements, National Institute of Standards and Technology (NIST) Risk Management Framework (RMF), and other NIST best practices. This list is not exhaustive; rather, it provides hand-picked examples of AI applications in federal cybersecurity environments.

Use Cases

Evaluating Dynamic Security Controls: An Organizational Defined Parameter (ODP) is a variable security control defined by each agency to customize security requirements based on a specific need. For example, ODPs can enable agencies to define key variables such as the number of unsuccessful logins prior to locking or disabling an account. Agencies can now use AI to ingest structured data, create a recommended set of ODPs across different modalities to reduce configuration errors, and ensure compliance with security controls.

Defining RMF Requirements: Continuous monitoring guidance is a result of FISMA, which mandated that NIST develop security standards and guidelines for federal agencies. As a result, NIST created the RMF, which is a structured process for implementing FISMA's requirements. Today, agencies are using AI to define RMF requirements such as continuous monitoring requirements that align with NIST RMF mandates.

Streamlining ATO and Moving Toward cATO: An Authority To Operate (ATO) means that a system's security controls, risks and internal controls, and documentation have been evaluated and approved for operation. Obtaining an ATO is a critical step in the RMF process and demonstrates compliance with FISMA. More recently, agencies are moving toward Continuous Authorization to Operate (cATO) environments to better implement continuous monitoring of system risks. Agencies are using AI to support multi-cloud, real-time compliance, infrastructure mapping, and automated control assessments to transition toward cATO.

Assessment of Access Controls and Artifact Generation: LLMs are a type of AI that can process, understand, recognize, and generate human language and text. LLMs fall within the generative AI category and focus specifically on language-related tasks. Datasets of text and code help train LLMs. Agencies and IT organizations can use questionnaires to evaluate security posture, including measures in place to control access to systems and sensitive information. When paired with security questionnaires, LLMs can generate

security artifacts and tailor access controls, helping to streamline the documentation process for audits and reduce the burden of evidence collection during the ATO process.

Regulatory Validation, Compliance, and Reporting: The Department of Defense's (DoD) Enterprise Mission Assurance Support Service (eMASS) is a System of Record that contains over 1,000 security controls that facilitate, manage, and track the ATO lifecycle. In addition to the well-known RMF process, NIST developed an AI RMF process to better identify, manage, and assess risks associated with AI. Controls and risks that are part of these RMF processes are largely IT in nature and do not account for governance and privacy concerns. In addition, legislation (e.g., USA PATRIOT Act) that is relevant to security controls may not be covered in RMF requirements. Agencies are using custom Natural Language Processing (NLP) models to streamline reporting and validation and align their efforts with mandates like OMB M-25-21 to demonstrate measurable return on investment (ROI).

Asset Management via Network Scanning: Network scanning helps enable effective asset management through the discovery and identification of all devices connected to a network. As outlined in The NIST Cybersecurity Framework and related publications such as NIST Special Publication 1800-5 (IT Asset Management), network scanning helps agencies better understand their IT environment so they can identify vulnerabilities, track assets, and secure them. Agencies now use AI-driven tools to scan federal networks in real time to identify hardware and software assets that align with NIST frameworks and cybersecurity best practices.

Modernizing Human Resources (HR) Processes: Beyond traditional cyber operations, AI supports federal HR departments by rewriting position descriptions (PDs) for cybersecurity roles. A PD for an open requisition serves as the foundation for recruitment and performance management within federal agencies. PDs also help agencies ensure that positions are classified to the appropriate General Schedule pay scale and demonstrate compliance with the Fair Labor Standards Act (FLSA). By utilizing AI, agencies can create draft PDs that reflect evolving skillsets that can help accelerate the hiring process and improve workforce alignment.

Using AI to Secure Data

In today's federal environment, protecting data requires much more than traditional defense mechanisms, audits, and internal controls. Systems must protect data at rest and in transit across multi-cloud, edge, and on-prem environments. Agencies must not only be reactive across these environments, but also proactive in identifying, classifying, and continuously monitoring sensitive information while adapting to evolving threats and new requirements. Cyber AI can play a role in securing data through data classification, control of sensitive information, adaptive responses to threats detected in real time, automation of repetitive tasks, and post-incident recovery.

Data discovery and classification help agencies understand what data exists, where it resides, and how sensitive it is so they can apply the correct security controls. As stated previously, one AI use case involves scanning networks to better manage assets. Once assets such as files, records, and datasets have been detected, agencies can further the AI use case by tagging the assets with the appropriate metadata and sensitivity labels. When implementing Attribute-Based Access Control (ABAC), AI-driven tagging allows access policies to be defined so that entitlements can be included in access tokens for decision-making within Policy Decision Points (PDPs).

AI also can be used to analyze existing user populations and behaviors and offer suggestions or changes to access privileges through role or entitlement updates. As previously mentioned, NIST RMF requires continuous monitoring. AI can assist with real-time access monitoring through the analysis of behavioral patterns (e.g., anomalies in access attempts or unusual account types attempting to access sensitive data). Regardless of how entitlements are generated and assigned, it is not wise to assume the entitlements are correct. Recurring entitlement validation is crucial in preventing unauthorized access and enforcing the principle of least privilege. Agencies can enhance data security by using AI to push out catalogs of entitlements for validation and auditing across the enterprise.

In SOCs, AI is emerging as a powerful tool to transform how fast cyber personnel can detect and respond to threats. To help with incident response planning, agencies are using AI to show the consequences of cyber attacks, an effort that allows better decision-making in developing effective incident response plans. Once cyber personnel establish a plan, SOCs can utilize behavior baselines as a reference point for users' normal activities, devices, and applications within a network. AI can help augment anomaly detection in an SOC by comparing against the detected baseline to flag deviations, identify false positives, and reduce alert fatigue. In the event of a breach, AI can ingest threat intelligence, trigger containment workflows, and automate recovery tasks such as responding to ransomware, addressing unauthorized data exfiltration, isolating endpoints, or generating compliance reports.

In addition to augmenting the SOC with real-time threat detection and recovery, AI can help secure the infrastructure and defend the network. By assessing traffic patterns, Cyber AI models can determine if there is resource congestion or an issue with routing configuration. In the instance of bandwidth constraints, AI models can recommend resource distribution to maintain quality of service and prevent denial-of-service vulnerabilities. However, as AI is adopted, it also can be targeted; bad actors are constantly devising new techniques to defeat AI technology. Thus, it is important for agencies to implement protection to maintain the integrity, confidentiality, and availability of AI models.

While AI can be used as a reliable tool to enhance security and automate tasks, human oversight is still needed. Generated AI data may not always be accurate and will need to be validated, and models will need to be updated and trained over time. AI should be used to accompany or augment access control policy and decision making. However, if strategically integrated with a feedback loop, AI can secure data and transform how agencies defend their missions.

AI-Driven Insights: What We've Learned

Perhaps the most important insight is the ongoing need for workforce training and education to facilitate culture change and AI adoption. Whether implemented in eMASS or SOC workflows, AI requires personnel who can interpret outputs and understand points of vulnerability in its implementation and deployment. Data outputs like AI performance evaluations, data incidents, and audit reports are great resources that the workforce can use for additional insights or integration into a feedback loop. In some pilots, IT personnel have integrated AI into Security Information and Event Management (SIEM) systems and used AI to map gaps in the MITRE ATT&CK framework. Systems then gain the capability to generate threat likelihood scores and help prioritize alerts, enabling more focused responses to hijacking attempts or insider threats.

A recent paper presents a layered framework approach to achieve a stronger cybersecurity posture known as "The Cyber Cake." Much like baking an actual cake, a Cyber Cake requires using certain ingredients, following a defined process, and introducing variations that can yield a desired, or undesired, outcome. AI and ML are cake toppings that interact with other layers for a more informed, resilient cybersecurity posture.

Agencies are increasingly adopting multi-cloud environments in which they rely on many cloud providers to help modernize IT infrastructure, enhance security, and avoid vendor lock-in. A landing zone, which is a preconfigured, well-architected cloud foundation, must be considered for AI services and combined with DevSecOps principles. This requirement adds to the complexity of building scalable, secure AI models, environments, and architectures. Here, automated AI baselines can help maintain visibility and compliance in spite of changing models, environments, and architectures.

One AI challenge is that publicly available tools such as Neural Processing Units (NPUs) are also available to bad actors. NPUs can help accelerate AI and ML workloads much more efficiently than traditional Central Processing Units (CPUs) or Graphics Processing Units (GPUs). Adversaries also can use AI to create more sophisticated, targeted ransomware attacks that are difficult to detect. In addition, AI can be used to break into systems and code and even anonymize digital signatures. This raises urgent questions

such as: How do we protect the models we rely on? How do we track model drift or overfitting in critical systems like LLMs? And, how do we validate AI-generated decisions in systems that touch sensitive federal data?

These challenges underscore the importance of the AI RMF, which agencies must adopt as a structured way to secure AI throughout their software development lifecycles. Even with RMF adoption, addressing the knowledge and workforce gap is important to scale progress. With over 100+ vulnerabilities reported each day – and thousands of vulnerabilities awaiting analysis in backlogs – traditional manual triage is unsustainable. Cyber AI can help bridge this gap, but agencies must not only prepare their infrastructure, but also their ecosystem. Doing so requires upskilling talent, modernizing systems, and guiding responsible AI use through governance that aligns AI initiatives, acceptable use policies, and threat-modeling capabilities.

Summary

Agencies are moving beyond pilots toward enterprise-wide adoption of Cyber AI capabilities for threat detection, compliance, and sensitive data protection across mission-critical environments. The highlighted use cases range from automated access control to AI-enhanced RMF documentation and anomaly detection in SOCs. Agencies also are operationalizing AI outside of traditional cybersecurity tasks to broader applications such as supporting HR processes.

In securing data, AI powers an array of safeguarding tasks that range from real-time asset discovery to dynamic policy and access enforcement. By tagging, classifying, and monitoring sensitive data at scale, agencies are creating more responsive and resilient data environments. Such data protection capabilities help bridge gaps in the workforce and reduce manual burden through automation, allowing cybersecurity personnel to focus on what matters most.

Important lessons learned include how to prepare the federal ecosystem, manage risk in AI models, and improve personnel skills continuously to operate successfully in AI-augmented environments. The most important takeaway, however, is that strategy and governance matter. Further, threat modeling must evolve in order to triage a backlog of vulnerabilities. And, as adversaries continue to weaponize AI and identify unknown vulnerabilities, the federal government must not only match that pace but also innovate, creating intelligent defenses guided by frameworks like the AI RMF and AI governance.

Electrosoft Is an Industry Leader in Cyber AI

While AI continues to transform the cybersecurity landscape, its success ultimately hinges on how effectively agencies can align emerging technologies with people and processes.

Electrosoft helps bridge that gap through the deployment of AI-powered solutions that align with training initiatives and that enhance SOC performance to strengthen the workforce.

Cyberspace workforce elements are defined in DoD Directive 8140.01, which establishes the DoD Cyberspace Workforce Framework (DCWF) as the authoritative reference for identifying, tracking, and reporting DoD cyberspace positions and developing baseline qualifications. As the federal workforce rapidly changes, readiness is no longer defined solely by experience. Instead, it's about capabilities that evolve with technology.

Electrosoft offers effective training in the following operational areas:

- **Personalized Learning Paths**: The use of AI to tailor training to individual skill levels and learning styles for faster skill development and higher retention rates.
- **Realistic Simulations:** The use of generative AI to create immersive, scenario-based, cyber-focused training to gain improved incident response skills and confidence through hands-on practice.
- **Optimized Knowledge Base:** The use of generative AI to continuously organize and prioritize Knowledge Articles so that cyber personnel are receiving the most up-to-date and accurate training in highly dynamic environments.

In the SOC, speed and precision are essential, but analysts are often overwhelmed by the sheer amount of data and dashboards. Electrosoft is transforming this environment by creating interactive AI solutions that are integrated into SOC dashboards to reduce the burden on cyber personnel. Electrosoft offers SOC analyst capability enhancement by focusing on the following:

- **Smarter Threat Detection and Prioritization**: Improved accuracy in identifying threats through behavioral analytics and enriched alerts with threat intelligence, reducing false positives.
- **Accelerated Incident Response:** Faster resolution times with automated investigations, streamlined root cause analysis, and rapid containment via Security Orchestration, Automation, and Response (SOAR) and playbook automation.
- **Enhanced Analyst Productivity:** Increased efficiency with automated alert triage, summarized log data, faster querying using Secure LLMs, and improved knowledgebase access.

- **Bridging the Skill Gap:** Elevated analyst performance through virtual assistance, realistic training simulations, and effective knowledge transfer.

Whether the challenge is upskilling personnel under evolving frameworks like DCWF or enabling SOC analysts to respond faster with fewer false positives, our solutions help deliver compelling results and value.

Contact Us

To learn more information about Electrosoft and our capabilities, contact us at info@electrosoft-inc.com.

About Electrosoft

Specializing in cybersecurity, Electrosoft supports federal civilian and defense organizations in advancing cyber resilience, achieving digital transformation, and adopting agile approaches that improve operational efficiency and security. With a focus on innovation and excellence, the award-winning company is recognized for its expertise, top workplace status, and leadership excellence. Electrosoft is headquartered in Reston, Virginia. www.electrosoft-inc.com.