adlumin™

The Ultimate Guide to Using Cybersecurity AI

Get Your SOC Ready for the Future

Contributing Authors:
Mark Sangster, Adlumin's VP, Chief of Strategy
Arijit Dutta, Adlumin's Director of Data Science
Zach Swartz, Adlumin's Senior Data Scientist
Brittany Holmes, Adlumin's Corporate Communications Manager

Table of Contents

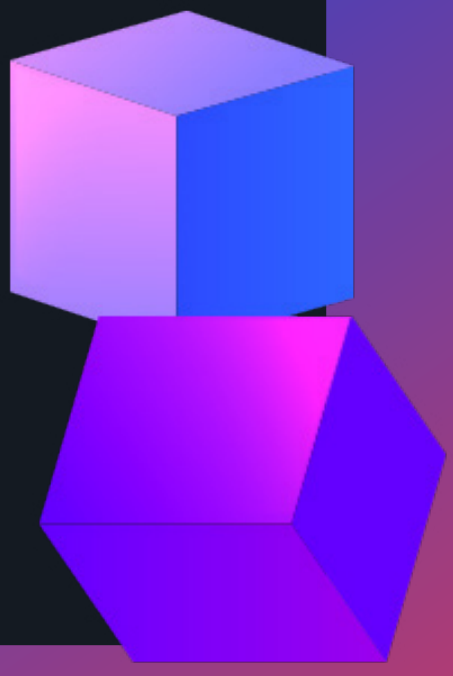

Introduction

The rise of artificial intelligence (AI) has transformed the approach of IT security experts towards cybersecurity. Advanced AI-powered tools and systems now offer enhanced data protection against threats by identifying behavior patterns, automating tasks, and pinpointing abnormalities.

AI plays a crucial role in protecting against cyber threats by actively monitoring, analyzing, and responding to potential risks in real-time. By separating more important insights from vast amounts of data, AI algorithms can detect patterns associated with malicious activities and identify vulnerabilities across entire networks. AI can effectively pinpoint unusual behaviors and prevent unauthorized system access by analyzing behavioral patterns and establishing baselines.

Additionally, AI aids in risk prioritization, enabling the instant identification of malware and intrusions before they compromise security. By automating repetitive tasks and minimizing human involvement, AI drives security automation, optimizing resource allocation and mitigating human errors.

How is AI in Cybersecurity Different?

The human element does not entirely replace the need for human cybersecurity professionals, it plays a crucial role in enhancing their capabilities and effectiveness. AI in cybersecurity involves using machine learning algorithms to analyze enormous amounts of data, identify patterns, and generate valuable insights.

These insights can then detect and prevent cyber threats in real-time by alerting security analysts of high-priority threats. In contrast to traditional security processes, which often involve manual analysis and can take hours or weeks, AI can perform these tasks at extraordinary speeds and with great accuracy.

This guide discusses the benefits of AI within your cybersecurity strategy, incorporating AI within the four stages of threat detection and response, how AI detects incidents, and recommendations for implementation.

How AI Enhances Cybersecurity

Three key benefits of using AI-driven cybersecurity tools are the ability to quickly analyze large amounts of data, detect anomalies and vulnerabilities, automate repetitive processes, and ultimately make cybersecurity faster and smarter.

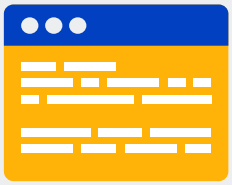

Quick analysis of large amounts of data:

Traditional cybersecurity methods rely on manual analysis, which can be time-consuming and prone to human error. AI-driven cybersecurity tools, on the other hand, leverage machine learning algorithms to process and analyze large volumes of data. These tools can rapidly scan through logs, network traffic, and other sources of information to identify potential cyber threats. By automating data analysis, AI systems enable quicker response times and proactive threat detection.

Detection of anomalies and vulnerabilities:

AI-driven cybersecurity tools excel at identifying anomalies and vulnerabilities that might go unnoticed by human analysts. These tools are trained to recognize patterns and extrapolate historical data to identify deviations or suspicious activities. They can detect subtle changes in network behavior, unusual access patterns, or unrecognized threats that may indicate a potential breach or attack. By continuously monitoring the network and applications, these tools can proactively identify and respond to new and evolving threats, enhancing the organization's overall security posture.

Automation of repetitive processes:

Cybersecurity professionals often spend significant time on repetitive tasks such as patch management, log analysis, and incident response. AI-driven cybersecurity tools can automate these processes, freeing up valuable time for cybersecurity professionals to focus on more complex tasks. AI systems can automatically apply security patches, analyze logs for potential security incidents, and even respond to low-level threats without human intervention. This automation increases the efficiency and effectiveness of cybersecurity operations, reduces the risk of human error, and improves response times.

The Four Stages of AI Advancements in Cybersecurity

AI enhances your Security Operations Center (SOC) by seamlessly integrating into four key stages of threat detection and response: detection, triage, investigation, and response. By harnessing the power of AI at each stage, you can unlock a multitude of benefits, including unmatched threat detection accuracy, rapid incident response, reduced manual effort, improved anomaly detection capabilities, and heightened overall efficiency.

Stage 1: Detection

AI plays a crucial role in the threat detection stage by addressing several challenges traditional detection methods face. One of the main challenges is the evolving nature of cyber threats. Attack techniques are constantly changing and becoming more sophisticated, making it difficult for traditional rule-based systems to keep up. On the other hand, AI algorithms can learn from large amounts of data and adjust their detection capabilities to new attack patterns.

In addition to tackling dynamic threats, AI-driven detection systems can also handle the vast amount of data your organization generates. Traditional methods often struggle to analyze and make sense of large datasets in real-time. AI algorithms, such as machine learning and deep learning, excel at processing and analyzing huge volumes of data, enabling them to identify patterns and anomalies within the data that may indicate malicious activity.

Detection Solution: User Entity and Behavior Analytics (UEBA)

UEBA is a machine learning cybersecurity process and analytical tool usually included with security operation platforms. It is the process of gathering insight into users' daily activities. Activity is flagged if any abnormal behavior is detected or if there are deviations from an employee's normal activity patterns. For example, if a user usually downloads four megabytes of assets weekly and then suddenly downloads 15 gigabytes of data in one day, your team would immediately be alerted because this is abnormal behavior.

The foundation of UEBA can be quite simple. A cybercriminal could easily steal the credentials of one of your employees and gain access, but it is much more difficult for them to convey that employee's daily behavior in order to go unseen. Without UEBA, an organization cannot tell if there was an attack since the cybercriminals have the employee's credentials. Having a dedicated Managed Detection and Response (MDR) team to alert you can give an organization visibility beyond its boundaries.

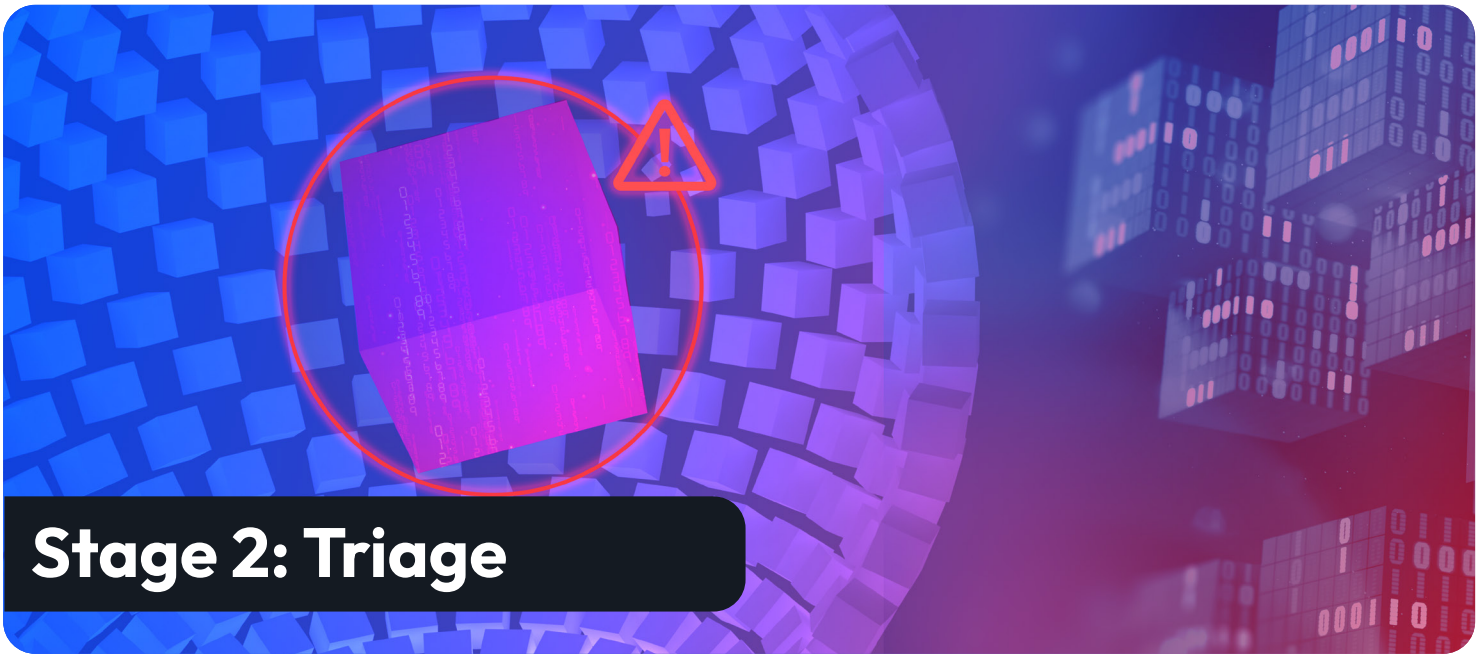

Stage 2: Triage

One of the main challenges is the high volume of false positives that can overwhelm security analysts. Traditional triage methods often generate many alerts that turn out to be false alarms, wasting valuable time and resources. Additionally, the increasing complexity and sophistication of cyber threats make it difficult for security analysts to keep up with the pace at which attacks are evolving.

AI within the triage stage offers significant benefits in improving detection fidelity. Machine learning algorithms can be trained on vast amounts of historical data and patterns to identify potential threats accurately. By analyzing and correlating multiple indicators and variables, AI can effectively separate true threats from false positives, reducing the workload for security analysts. This allows them to focus their attention on investigating genuine threats, thereby improving the accuracy and efficiency of the triage process. AI-powered tools can also continuously learn and adapt to new threats, ensuring that detection capabilities remain up-to-date and effective.

The application of AI within the triage stage is versatile and can be used across various industries to enhance threat detection and incident response. For example, AI can help identify fraudulent transactions and potential cyberattacks on customer accounts within financial services, enabling prompt action to mitigate risks. Another example is in healthcare, where AI can assist in detecting and responding to unauthorized access to patient information, ensuring patient privacy and compliance with regulatory requirements. AI can be used in critical infrastructure sectors such as energy and transportation to detect and respond to cyber threats that could have significant real-world consequences. By leveraging AI within the triage stage, organizations can improve their security posture, save time and resources, and enhance the overall effectiveness of their threat detection and response capabilities.

Stage 3: Investigate

One of the main challenges in the cybersecurity investigation stage for threats is the overwhelming amount of data that needs to be analyzed. With the increasing complexity and frequency of cyber threats, security teams often struggle to identify and prioritize the most critical events quickly. This can lead to delays in response decisions and an increased risk of damage. Additionally, the shortage of skilled cybersecurity professionals adds to these challenges, as their expertise is needed to investigate and mitigate threats effectively.

Security teams can efficiently focus on priority events by leveraging AI algorithms and machine learning capabilities. AI tools can analyze vast amounts of data in real-time and identify patterns, anomalies, and potential threats. This allows security professionals to prioritize their efforts and allocate resources where they are most needed, ultimately reducing response time and minimizing the impact of cyberattacks. Additionally, AI can enhance in-action customer support by continuously monitoring and alerting users about potential threats, enhancing the overall security posture of an organization.

The application of AI in the cybersecurity investigation stage is varied and encompasses several areas. For instance, AI-powered Security Operations Platforms can automatically gather and analyze intelligence from various sources, providing MDR teams with real-time insights into emerging threats. AI can also analyze network traffic to identify suspicious activities and potential breaches. Additionally, AI algorithms can aid in investigating security incidents by automatically correlating data from different sources, such as logs, network traffic, and user behavior.

This not only speeds up the investigation process but also makes it more thorough and accurate. Overall, the application of AI in the cybersecurity investigation stage enhances the efficiency and effectiveness of threat identification and response, ultimately strengthening an organization's overall security defenses.

Stage 4: Respond

One of the challenges of using AI within the response stage of cybersecurity incidents is the complexity of the threat landscape. Cybersecurity threats are constantly evolving, and traditional rule-based systems may struggle to keep up with the sophistication of these attacks. However, by leveraging AI algorithms within a Security Operations Platform, organizations can respond quickly and effectively, mitigating the impact of an attack by identifying patterns and anomalies.

AI in the cybersecurity incident response stage helps build validated response activities. By automatically analyzing and correlating data, AI algorithms can identify and prioritize potential threats and recommend response actions. This significantly increases the speed and accuracy of incident response, allowing organizations to quickly address and contain security breaches. Additionally, AI can provide real-time reporting on the status of the incident, giving organizations a clear and up-to-date view of the situation. This enables them to make well-informed decisions and allocate resources effectively.

Response Solution: Security Orchestration, Automation, and Response (SOAR)

SOAR is a form of pure automation that immediately stops a threat even before a security analyst reviews an alert, greatly reducing an organization's risk. These tools are used for the following operation tasks:

- To document and implement processes
- To support security incident management
- To apply machine-based assistance to human security analysts and operators
- To better operationalize the use of threat intelligence

SOAR takes how IT teams respond to alerts to the next level. When teams are tasked with hundreds, sometimes thousands, of alerts daily, there is no room for human error when evaluating which one should be prioritized as high-risk. SOAR automates the organization's MDR teams' response and alert processes and systematically orchestrates them. SOAR functions can initiate and disable accounts in machine time to contain the threat and reduce the amount of damage done. This can occur before an analyst even has eyes on it.

AI In Action: Incident Detection to Response

AI employs various techniques to detect incidents and can identify deviations from normal behavior by continuously monitoring data streams from multiple sources, such as network logs, system metrics, user activities, and cybersecurity tools. These deviations may indicate potential incidents such as security breaches, system failures, or performance issues. AI algorithms can also learn from historical incident data to improve their detection capabilities over time.

DETECT

Discover novel attacks
Identify dangerous behaviours
Develop early-stage detections

TRIAGE

Eliminate false-positives
Improve detection fidelity
Improve SOC productivity

INVESTIGATE

Focus on priority events
Accelerate response decisions
Provide in-action customer support

RESPOND

Build validated actions
Provide real-time reporting
Coordinate client-only responses

Incident Use Case: Lateral Movement

When a cybercriminal first gains access to a network, they often move laterally to different machines until they find what they're looking for. This movement can usually be traced using Windows access-event logs. However, detecting this behavior automatically amongst the enormous corpus of access events can be extremely difficult, especially when it involves privileged users with network-wide access.

A machine-learned algorithm assigns an anomaly score to each successful logon event that occurs on an employee's machine. Anomalous logon events are then aggregated to form access graphs, which are subsequently assessed for attack signatures. For example, many anomalous logon events from a single user on a single machine may indicate that an attacker is attempting to gain access to network share drives or performing scanning-like behavior to gain a foothold elsewhere on the network. Figure 1 shows an example of a penetration test detected on a customer network where the user attempted to access 15 separate machines.

This "one-to-many" behavior is a common aggressive tactic that AI can target explicitly and can be essential in discovering and mitigating network breaches.

Prompt Detection and Response

A customer received multiple rule-based alerts concerning SSH and FTP connections to a foreign country, along with three lateral movement detections associated with an admin account exhibiting those mentioned above "one-to-many" behavior. The attacker accessed roughly 30 hosts across the network and copied several documents. It appears that data exfiltration was attempted but unsuccessful, since they were quick to contain the threat, damage was minimized. The combination of rule-based and AI-powered detections allowed the MDR team and the customer to promptly address the issue as much as possible.

Figure 1: Penetration Test Example

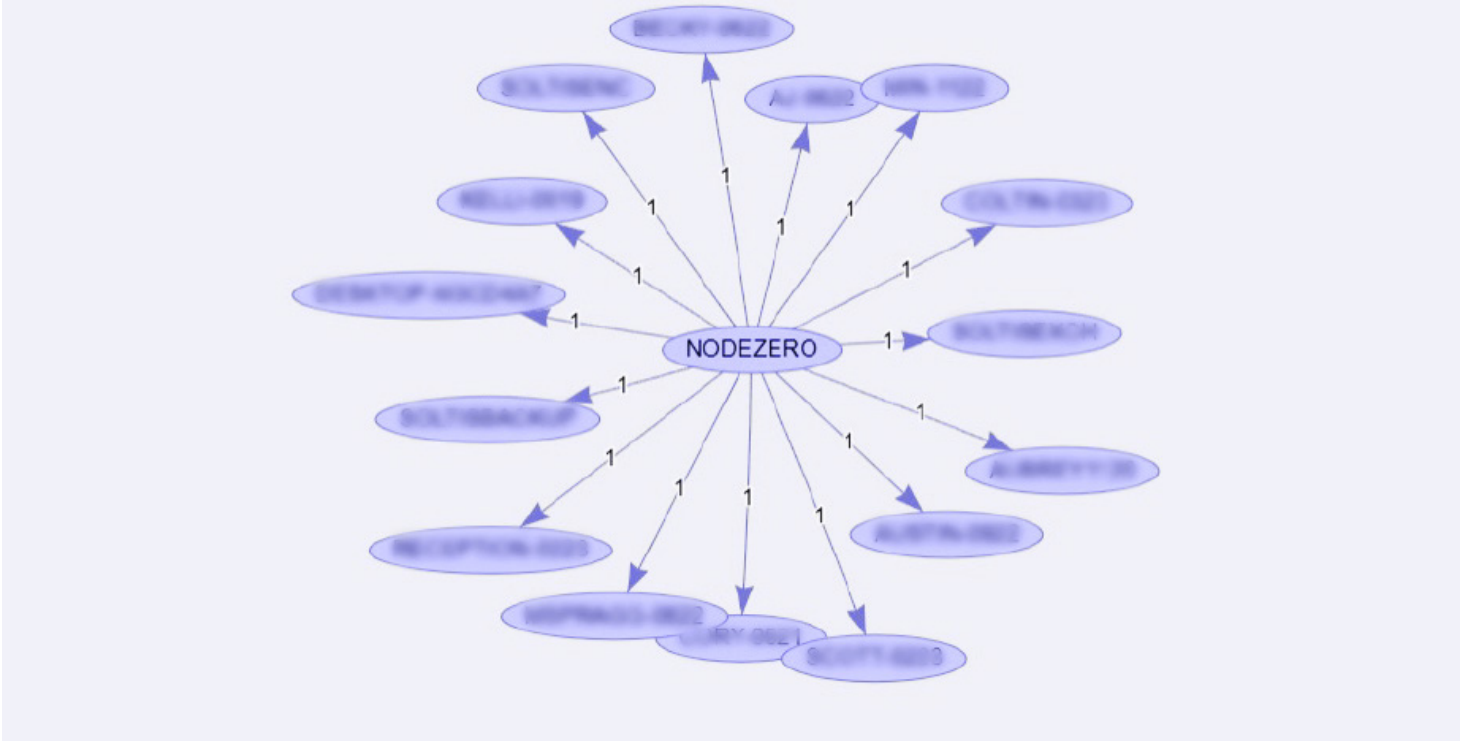

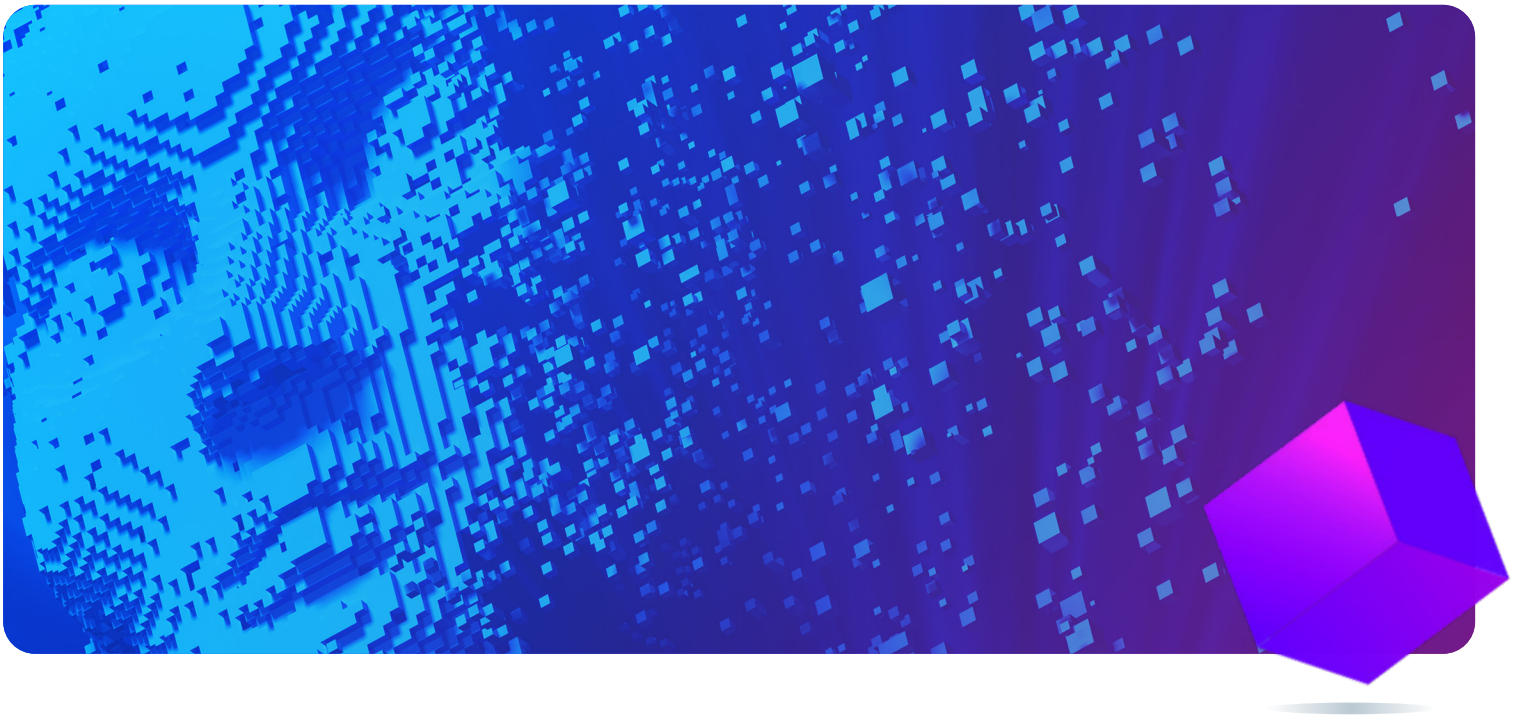

AI Technology: A Look into the Present and Future

As in most industries, AI technology is indispensable in organizations today for distilling actionable intelligence from the massive amounts of data being ingested from customers and generated by employees. Organizations can choose from various available data mining and AI methods depending on desired outcomes and data availability. For example, if the goal is to evaluate each customer for digital marketing suitability for a new product, "supervised" methods such as logistic regression or decision-tree classifier could be trained on customer data.

These use cases require customer data on prior actions, such as historical responses to marketing emails. For a customer segmentation problem, "unsupervised" methods such as density-based clustering algorithm (DBSCAN clustering) or principal component analysis (PCA) dimensionality reduction are called for, where we don't impose prior observations on specific customer actions but group customers according to machine-learned similarity measurements.

More advanced methods, such as Artificial Neural Networks, are deployed when the use case depends on learning complex interactions among numerous factors, such as customer service call volume and outcome evaluation or even the customer classification and clustering problems mentioned earlier. The data volume, frequency, and compute capacity requirements are typically heavier for artificial neural networks (ANNs) than for other Machine Learning techniques.

The most visible near-term evolution in the field is the spread of Large Language Models (LLM) or Generative AI, such as ChatGPT. The underlying methods behind these emergent AI technologies are also based on the ANNs – only with hugely complicated neural network architectures and computationally expensive learning algorithms. Adaptation and adoption of these methods for customer classification, segmentation, and interaction-facilitation problems will be a trend to follow in the years ahead.

Adlumin's AI Recommendations

Establish clear AI policies

It is crucial to establish clear AI policies that outline the ethical guidelines and principles that the systems should follow. This includes defining what is considered acceptable behavior and what actions should be prohibited. Creating a governance council composed of experts from various fields, such as ethics, law, and technology, can help set these policies and provide guidance.

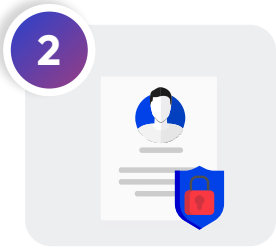

Avoid including confidential data

Sensitive data, such as personal information, should not be included in the training data sets. Removing sensitive data helps prevent privacy violations and potential harm if this data is misused or falls into the wrong hands. Proper anonymization techniques and data protection measures should be employed to safeguard user privacy.

AI systems should not have unrestricted control over publishing content

Automatic publishing can lead to misinformation or harmful content dissemination without proper human oversight. Therefore, it is important to strictly prohibit automated publishing and ensure that any content generated by the GenAI systems is reviewed and approved by human experts before being made public.

Employ human oversight

It is essential to ensure that the actions and policies of AI systems are aligned with ethical standards. Human experts should review and vet the decisions made by the AI system to ensure that they are fair, in compliance, and unbiased. It is important to have independent mechanisms in place to address any concerns or complaints regarding the AI system's actions or decisions. This could involve establishing an independent panel or specific person to review and respond to feedback or concerns raised by users or affected parties.

Organizations Embracing AI

Organizations need to build a cybersecurity infrastructure embracing the power of AI, deep learning, and machine learning to handle the scale of analysis and data. AI has emerged as a required technology for cybersecurity teams, on top of being one of the most used buzzwords in recent years. People can no longer scale to protect the complex attack surfaces of organizations by themselves. So, when evaluating Security Operations Platforms, organizations need to know how AI can help identify and prioritize risk and help instantly spot intrusions before they start.

About Adlumin

Adlumin Inc. provides the enterprise-grade security operations platform and managed detection and response (MDR) services that keep mid-market organizations secure. With one license and one platform, its patented technology gives organizations and solution providers everything they need for effective threat hunting, incident response, vulnerability management, darknet exposure monitoring, compliance support, and much more.

The Adlumin platform is feature-rich enough for organizations to operate independently yet explicitly built to amplify the skills and capabilities of managed service providers who use it to deliver cutting-edge security that can scale to meet the needs of any operating environment. With full access to the platform regardless of whether they are running it themselves or relying on Adlumin's MDR services or expert partners, Adlumin gives organizations unparalleled visibility into their security posture through access to alerts, investigation data, threat intelligence, compliance reporting, and everything else – all in real time.

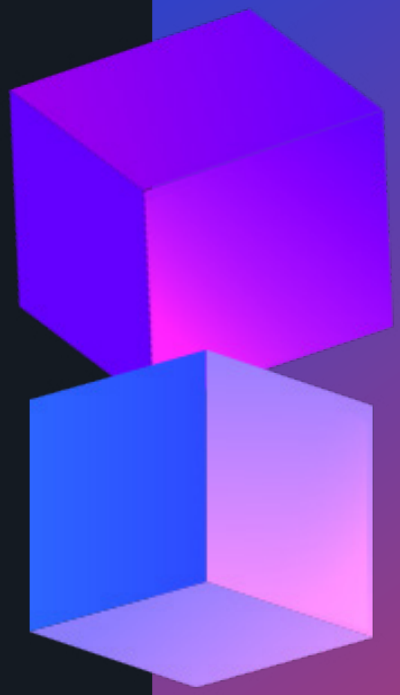

Ready to Demo?

Schedule a briefing and live demo of Adlumin's SIEM platform and learn more about key features designed for your specific industry

Learn More

adlumin.com