A Buyer's Guide to AI in Network Security: 5 Essential Features

Introduction

Organizations face numerous challenges to detect and contain sophisticated threats. Three obstacles in particular stand out. First, the threat landscape is becoming increasingly complex. In 2017, NTT Security observed a 24 percent increase in network-based attacks against organizations over the previous year. In its 2018 Risk Determination Report, the Office of Management and Budget found the cyberthreat situation to federal agencies "*untenable*" as agencies lack both the visibility into their networks to determine the occurrence of cybersecurity incidents and the ability to minimize the impact from an intrusion.

Second, organizations are having trouble hiring skilled professionals to protect their systems. In May 2017, Cybersecurity Ventures estimated that the number of unfilled jobs in digital security – the "skills gap" – would increase to 3.5 million by 2021. This is concerning given its earlier prediction that the global cost of digital crime will increase from $3 trillion in 2015 to $6 trillion in the same year.

Last but not least, organizations' systems are becoming more complex. Two major developments are driving this trend. First, organizations are increasingly migrating to the cloud, which makes it more difficult for organizations to inventory and protect their sensitive IT assets and data. Second, organizations are connecting IoT devices to their networks that lack proper security safeguards. As a result, many organizations must contend with defending an even larger attack surface.

AI: A Way Forward

To respond to these challenges, organizations need to look at their information security programs. Any well-structured plan will take into account the Golden Triangle, which consists of a company's people, processes, and technology. Unfortunately, many companies have too few people to properly safeguard the network given the skills gap. Still, businesses need some way to fulfill their information security processes. This leaves technology as a possible way forward.

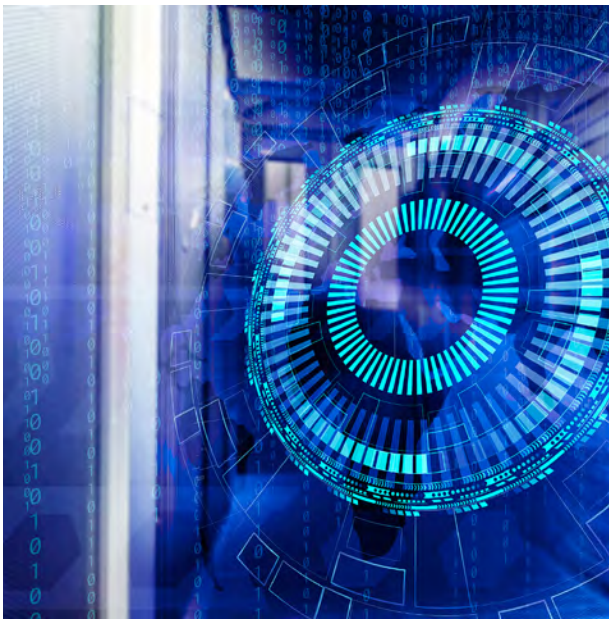

One of the most promising technologies for digital security in today's world is artificial intelligence (AI). SAS Institute defines artificial intelligence as technology that makes it "possible for machines to learn from experience, adjust to new inputs, and perform human-like tasks." Security analysts can take advantage of this functionality to teach AI systems what benign and malicious activity typically looks like. These types of applications could save businesses time and money by using automation to help information security processes function more efficiently. Information security personnel can then focus on responding to the most important threats, instead of slogging through a mountain of alerts looking for clues.

Of course, not every security solution incorporates AI into its toolset. Those that do commonly differ in the types of features that they offer. This complicates the process of purchasing an AI product that provides an adequate level of protection against today's cyberthreats while meeting all of their business requirements.

To help make the buying experience easier, here are five essential features in a network security product utilizing AI:

1. Reduce false positives

A security product utilizing AI will inevitably identify some network behavior as anomalous, since one of the primary benefits of AI is detecting patterns in vast amounts of data. However, this means that most tools will flag anomalous activities as malicious, when in fact they are benign and thereby waste scarce analysts' time with unnecessary investigations. A tool that generates a significant number of false positives, regardless of the potential value of the underlying technology, will quickly become a burden to the security team.

To avoid this, companies need to identify an AI-based product that minimizes the delivery of false positives. One method to reduce false positives is to train the AI using the right dataset. Exposing AI to a very large data set of malicious behaviors enables to the tool to learn both what is "good" and what is "bad", and not generate an overwhelming number of false positives.

Another method to reduce false positives is to provide additional context to the anomalous activity detected to better distinguish between benign and malicious activity. To achieve this, some AI-based security tools correlate anomaly detection with other data sets to highlight any relationships between them.

2. Detect both misuse and anomalous malicious behaviors

The term "anomalous activity" is a commonly used synonym for malicious behavior in vendor marketing material. However, there are actually two types of malicious behaviors that organizations need to detect: Misuse and Anomalies.

- **Misuse:** To detect misuse, the AI creates a model of known malicious activity, and then identifies instances of maliciousness in the data. Using signatures to detect network attacks is one example of a misuse detection approach. The advantage of these approaches is that they are generally very accurate. However, they can only detect malicious behavior that has been seen before.
- **Anomalies:** To detect anomalous behavior, the AI creates a model of what is normal and identifies behaviors that are outside the parameters of normality. The advantage of this approach is that it is possible to identify malicious behavior that has never been seen before. Many vendors of AI-based security tools use this approach, but it is based on two important assumptions: "what is anomalous is malicious and what is malicious will generate an anomaly". Unfortunately, both assumptions are wrong, causing both false positives and false negatives.

By selecting an AI-based tool that can detect both types of malicious behaviors, organizations will significantly increase their ability to protect themselves from advanced threats.

3. Visualize the entire attack chain

Security incidents are complex events that don't consist of a single action. They are usually comprised of a series of smaller events that together form a trail of what happened. Companies need a tool that can pick up on these minor warnings and even more importantly connect these notifications together under a single incident and then combine those incidents to visualize the entire attack chain: From the initial device compromise to command and control communication to network discovery to lateral movement to data harvesting to exfiltration.

Also, security tools and the alerts they generate are often focused on one particular stage in the attack chain, rather than presenting a more comprehensive view of the threat. Security Operations Center (SOC) analysts and Incident Response (IR) teams must try to build a complete picture of the threat based on hundreds or thousands of alerts about discrete events. Many security controls are very good at identifying events within a limited focus and cannot "connect the dots" to show the end-to-end attack chain activity.

In addition, advanced threats use increasingly sophisticated evasion techniques to fly under the radar. They are engineered to bypass "next-generation" tools like firewalls and sandboxes when attempting to enter a network. And, once past perimeter controls, the intrusions are difficult to detect due to the high volume of alerts security teams have to analyze. Attackers are also able to hijack legitimate IT tools and services to hide their malicious activity as they move laterally in the network.

Accordingly, organizations need the complete attack chain visibility, beginning with the initial compromise. Relying on an AI-based tool to detect a threat as it begins to move laterally carries significant risk, due to false positives and evasive behavior.

4. Automate threat response

Security teams face a number of challenges when it comes to the quality of the alerts they have to analyze. In addition to receiving too many warning messages, many of those alerts yield only a low-fidelity assessment of the scope of the threat. They need a tool that can generate the detailed information they need to accelerate and simplify their ability to respond to those few significant events buried in the daily flood of alerts.

Otherwise, SOC and IR teams must use time-consuming manual steps to investigate suspicious activity to ensure the alerts are legitimate, understand the alerts' context in an attack chain, and initiate response workflows. Unfortunately, many AI-based tools simply cannot generate the accurate alerts security teams need because they lack the ability to detect and fully analyze malicious behaviors. The lack of high-fidelity alerts causes the security team to not trust the detection capabilities of their tools and consequently not automate workflows (such as blocking of malicious activity as it begins to move across the network). The result is slower, more labor-intensive analysis and delayed response, increasing the potential of a successful breach.

5. Integrate with existing tools

Too often a security analyst needs to perform "swivel chair analysis" because of the lack of integration among security controls. These siloed tools force an analyst to move between different consoles to investigate potentially malicious behavior.

Organizations get the most value out of any security tool, including an AI-based security tool, if it integrates with the security controls they already use. For instance, many AI-based tools analyze network traffic for lateral movement in a network and generate alerts which are not easily integrated into existing workflows (resulting in the "swivel chair analysis" described above). Instead, if that tool were able to share data automatically with other controls, such as endpoint security and firewalls, it would be much more effective in preventing additional system compromises.

Summary

Sophisticated security threats are too numerous, diverse, and evasive for a limited number of skilled security professionals to see and recognize on their own. AI provides a way forward for companies to accelerate and improve their information security processes and strengthen their cyber resilience. But not any solution will do.

Organizations need an AI-based solution that provides high fidelity insights into complex network attacks. It should able to "connect the dots" of malicious network activity to create a complete picture of the entire attack chain, to accelerate and simplify threat response. Using the tips above, companies can find a solution that meets these criteria.

Introducing Lastline Defender

Lastline Defender™, the industry's most accurate Network Detection and Response (NDR) platform, delivers the key capabilities described above. Its AI-powered detection learns from both network traffic and malicious behaviors to deliver the highest fidelity insights possible into malicious activity across your entire network, on-premises and the cloud.

Unique Approach

The Lastline Defender platform uses a unique combination of three complementary technologies to significantly reduce false positives and detect both misuse and anomalous behavior:

- First, it leverages the knowledge in the Lastline Global Threat Intelligence Network to scan traffic metadata and payloads for variants of known threats
- Second, it applies unsupervised machine learning (ML) to network traffic to detect protocol and traffic anomalies
- Third, it uses supervised ML to automatically create classifiers that recognize malicious network behaviors and previously unknown malware

Most AI-based network security products only implement a subset of these technologies. These probabilistic approaches lead to many false positives – after all, not all anomalies are malicious and not all malicious activity is anomalous. Applying AI techniques to network traffic will inevitably find anomalous patterns of behavior within the network traffic, because that's what AI is designed to do. And, because it is virtually impossible for other AI-based tools to understand if the detected anomaly is malicious or benign, it falls upon the security team to have to manually investigate the alerts.

The Lastline Defender platform's AI learns from analyzing both network traffic and malicious behaviors to eliminate most false positives. It delivers the highest fidelity insights possible into threats entering or operating within your network, including compromised personal devices and rogue IoT devices.

Accelerate Response

Your existing security team and security controls are more effective on day one with the Lastline Defender platform. Its deterministic alerts significantly reduce the high volume of false positives and generic alerts that other tools generate. Your team will finally have the confidence to automate many of your threat response workflows. This means better enterprise security with fewer resources.

The Lastline Defender platform connects threats identified in emails and web content with intrusion activity on the network, such as initial compromise, command & control communication, privilege escalation, lateral movement, and data exfiltration. This complete visibility of the entire attack chain provides the critical context that other technologies lack.

The native cloud architecture secures your workloads regardless of where they reside. It combines detected malicious activity from your on-premises and cloud networks into a single console, simplifying your response.

It also consolidates multiple events into a single incident and multiple incidents into a single intrusion. This consolidation makes it easier for your team to understand the significance of the malicious activity without additional investigation.

Armed with deterministic, high fidelity alerts, your security team can automate incident workflows by blocking advanced threats entering or operating within your network:

- Lastline Defender sensors can actively block malicious behavior at the perimeter or in internal traffic
- Lastline Defender also automates protection by integrating with third-party products, incident response workflows, and custom applications an organization has in place, whether on-premises or in the cloud. Existing security controls can automatically send unknown objects and websites to Lastline Defender for analysis and receive actionable threat intelligence to incorporate into threat response workflows.

We Understand What's at Stake

Most organizations struggle to see and recognize sophisticated cyberthreats. The Lastline Defender platform enables you to detect and contain sophisticated threats before they disrupt your business. It delivers the cybersecurity industry's highest fidelity insights into advanced threats entering or operating in your entire network, enabling your security team to respond faster and more effectively to threats.

Experience the Lastline Advantage

For more information please visit www.lastline.com

Lastline Corporate Headquarters
203 Redwood Shores Parkway, Suite 500
Redwood City, CA 94065

+1 877 671 3239
info@lastline.com
www.lastline.com