

Applying Artificial Intelligence to Cybersecurity Beyond the Hype

Table of Contents

Executive Overview	3
Introduction	5
Where Does the Artificial Intelligence Reside?	
Global Threat Research Labs	6
Centralized Customer Data Lake	8
Distributed AI Throughout the Organization	10
Where in the Cyber Kill Chain Is AI Applied	12
Conclusion	13

Executive Overview

Artificial Intelligence (AI) is a “hot” technology that holds great promise in many fields, including cybersecurity. In fact, many security vendor threat research labs—rich in big data and security expertise—have been using it to keep pace with a constantly evolving cyber-threat landscape. More recently, such technologies are being applied to an organization’s own cybersecurity information to detect threats that may have bypassed traditional cyber defenses. In certain cases, AI has even been distributed to the existing inspection points for real-time blocking of threats without global threat intelligence.

Assessing the true value of an AI solution requires cutting through the hype and ambiguity. Key considerations include:

- **Deployment location and model:** Is the AI deployed in a global threat lab, organization-specific data lake, or as part of in-line security controls? Which AI models are used to generate threat intelligence?
- **AI output:** Does the threat intelligence block a threat when received, trigger an alert, or trigger an immediate blocking action?
- **Threat coverage:** Which threat classes and cyber kill chain stages are covered by the AI model?

No single location, output, or cyber kill chain stage is best to apply AI technologies. Each organization should determine the right mix based on factors such as their risk appetite, budget, staffing, and security maturity level.



“The battleground of the future is digital, and AI is the undisputed weapon of choice.”¹

Introduction

To keep up with the volume, velocity, and sophistication of today's (and tomorrow's) cyber-threat landscape, organizations must utilize automation, and ultimately, AI. However, as Gartner notes,² “Artificial intelligence (AI) is ‘hot’ and hyped. CIOs, AI, data and analytics leaders across many industries are seeking breakthroughs, which will come in the long run. For now, though, they should focus on finding practical uses for AI that will have immediate impact.”

This is especially true for cybersecurity, where hype abounds. Identifying those practical uses, in order to stay ahead of the increasing volume, velocity, and sophistication of cyber threats, requires consideration of three factors:

1. Where the artificial learning resides, and thus, the data to which it can be applied
2. The security output received and how it can be used
3. The threat classes and cyber kill chain stage(s) to which it can be applied

Where Does the Artificial Intelligence Reside?

Global Threat Research Labs

Security vendors' threat research labs have led the way for the development of AI in cybersecurity. With global views and reach, many of these groups were among the first to realize that human-developed threat analysis would not be able to keep up with the rapidly evolving threat landscape.

At the same time, tasked (and indeed financially incentivized) with protecting hundreds of thousands of customers, the business case for an investment in advanced analytics was strong. Today, it is common that machine learning (ML) in particular is used to speed the identification of new cyberattacks, and more importantly, the indicators of compromise (IOCs) associated with them. These IOCs make up a large part of the threat-intelligence updates delivered to security products and services that protect a security vendor's customers.

Key benefits to AI-powered threat intelligence from global threat research labs include:

1. A huge dataset, across multiple threat classes and their full threat life cycles, to initially develop and constantly refine AI models
2. Massively scalable processing power whose cost can be shared across a large customer base
3. Some of the foremost security experts continually validating the fidelity of the output

Where Does the Artificial Intelligence Reside? (cont.)

Global Threat Research Labs (cont.)

However, there are also significant limitations to this ultra-centralized approach, such as:

1. The individual customer only receives IOCs, and those IOCs relate only to threats that reach the global lab
2. The process for validating and delivering those IOCs typically takes hours or more
3. In such complex, multidisciplinary, and distant research locations, customers have limited visibility into how AI is really being used by the security vendor

Customers of global threat research labs receive AI outputs in the form of threat-intelligence updates designed to help protect against the latest cyber threats. This may include updates to security products deployed within the organization, such as antivirus (AV) updates to a next-generation firewall (NGFW) or endpoint protection platform, or a subscription threat feed, providing a raw list of known malicious IP addresses or similar IOCs.

“Threat intelligence that tips your organization off to an impending cyberattack is timely. Putting together the indications that an attack was coming after it already happened is not.”³

Where Does the Artificial Intelligence Reside? (cont.)

Centralized Customer Data Lake

A popular contemporary approach is to bring AI in-house, to a centralized location for the individual organization via a “data lake.” This approach compensates for the limitations of global threat intelligence by combining AI models similar to those used in the research labs with data specific to the organization. As a result:

1. The protected organization identifies IOCs specific to the cyber threats to which it is exposed
2. Those IOCs, plus related insight about the threat campaign and its stages, are often available to staff after the AI analysis is applied
3. The organization knows exactly the type (model) and scope (threat class) of AI applied

However, there are also drawbacks to this organization-specific approach, such as:

1. Lost visibility into global threat activity, which may reach them in the future
2. Expensive infrastructure, such as storage, processing, space, and power, required for a central data lake
3. Time required for data collection, normalization, and analysis delaying results

This last point is one of the most significant limitations since it means that AI can only be deployed in a detective capacity. By the time analysis is complete, the attack has already occurred and requires remediation.

Specifically, data lakes and AI analysis, output in the form of alerts, must be prioritized, investigated, and confirmed or invalidated. IOCs and other threat intelligence are identified by the organization’s security staff, rather than global threat researchers, and must be added to the organization’s security controls, either manually or using automation. And of course staff must clean up the compromised systems.



“There is no doubt that possession of data can confer competitive advantage, but it must be timely data and relevant to current challenges and market opportunities. Having more data for the sake of it doesn’t deliver beneficial business outcomes. It creates liability.”⁴ — Gartner

Where Does the Artificial Intelligence Reside? (cont.)

Distributed AI Throughout the Organization

A third approach to applying AI to cybersecurity involves deploying the AI models where the data (files, IP addresses, system activity, etc.) to be analyzed passes or resides. Typically, this includes traffic ingress, egress, and internal inspection points, server and end-user host devices, and on-premises or cloud application delivery.

This approach to AI has several advantages, such as:

- The ability to apply AI for prevention as well as detection
- AI identifying the specific threats facing an organization
- Full understanding of the types of AI in use and the threat classes they address

However, this approach does have its limitations, including:

- A focus on threats at the organizational, rather than global, level
- Limited processing power for distributed AI
- Focus only on the threat classes passing a given inspection point

Output from distributed AI models can either be used to prevent a potential threat or to generate an alert for further investigation and response. The optimal configuration depends on deployment specifics, such as the exact deployment location, data analyzed, duration of analysis, desired configuration, and similar factors.

For example, at the endpoint, ML that inspects file characteristics or early behaviors associated with exploit attempts can often be deployed as a prevention mechanism. In contrast, a lightweight sensor that transmits host system activity to a cloud for ML analytics will typically generate an alert for further investigation.



“Teams that are using [AI] to augment their existing analysts ... are more effective than their peers and even SOC teams with more than 10 members who are not using AI.”⁵

Where in the Cyber Kill Chain Is AI Applied

In addition to the physical location of the analytics and nature of the output, it is important to consider the cyber-attack phase to which it applies. Lockheed Martin created the cyber kill chain, outlining seven common stages of a cyber threat,⁶ all of which must be successful for a cyber criminal to achieve their ultimate objective. These are reconnaissance, weaponization, delivery, exploitation, installation, command and control, and action on objectives. Successfully thwarting the cyber criminal at any single stage denies the attacker their ultimate objective.

The stage in the cyber kill chain where the organization acts determines whether an attack is prevented before impact or requires more costly detection and response. Threat intelligence generated in global threat research labs is often useful for prevention. IP reputation databases may detect reconnaissance efforts, and threat intelligence provided to deployed security systems are designed to interrupt the delivery, exploitation, and installation stages.

However, a prevention-only strategy is not always effective, as evidenced in the steady stream of data breach headlines. Most organizations are investing in threat-detection capabilities, which Gartner calls out as “still heavily weighted toward the end of the cyber kill chain.”⁷

The data lake approach often serves as the foundation for AI-based endpoint detection and response, user and entity behavior analytics (UEBA), and other detection methods. These controls are applied after the installation stage to identify anomalous activity often associated with command and control or the action on objectives, which is often data exfiltration. At this stage, AI is geared toward identifying an attack in progress, before final-stage objectives such as data exfiltration occur.

A distributed approach to AI holds great promise, since it enables the application of organization-specific AI models early in the cyber kill chain. By targeting the delivery, exploitation, and installation phases, an organization decreases the probability that a costly response will be required.

Conclusion

An organization can apply AI in different locations, whether at global research labs or within the organization, and at different stages of the cyber kill chain. The output of AI can also be applied for either detection or prevention. Each approach has its advantages, as well as limitations, and organizations must cut through the hype and hyperbole to identify the mix of approaches that is optimal for their situation.

This should include a mix of each type of AI. An organization's security vendor should provide broad coverage of the threat intelligence types and threat classes. This enables deployed security controls to detect and block threats early in the cyber kill chain.

This threat intelligence should be complemented with security products, such as next-gen AV, web application firewalls (WAFs), secure email gateways (SEGs), and sandboxes, with integrated AI. This built-in AI often enables prevention of threats specific to an organization or early response to global outbreaks.

When possible, an organization should also utilize advanced detection and response systems, such as endpoint detection and response (EDR), security information and event management (SIEM), and UEBA. These complement controls targeting early kill chain stages, enabling comprehensive detection and response for attacks that evade preventative controls. However, it is essential that security teams are appropriately staffed and skilled to effectively respond.

¹ William Dixon and Nicole Eagan, "[3 ways AI will change the nature of cyber attacks](#)," World Economic Forum, June 19, 2019.

² Kenneth Brant, et al., "[Hype Cycle for Artificial Intelligence, 2019](#)," Gartner, July 25, 2019.

³ Zane Pokorny, "[3 Key Elements of Threat Intelligence Management](#)," Recorded Future, August 8, 2018.

⁴ Nick Heudecker and Adam Ronthal. "[How to Avoid Data Lake Failures](#)," Gartner, August 10, 2018.

⁵ Zeljka Zorz, "[AI is key to speeding up threat detection and response](#)," Help Net Security, August 14, 2017.

⁶ "[The Cyber Kill Chain](#)," Lockheed Martin, accessed March 23, 2020.

⁷ Craig Lawson, et al., "[Market Guide for Managed Detection and Response Services](#)," Gartner, July 15, 2019.



www.fortinet.com

Copyright © 2020 Fortinet, Inc. All rights reserved. Fortinet®, FortiGate®, FortiCare® and FortiGuard®, and certain other marks are registered trademarks of Fortinet, Inc., and other Fortinet names herein may also be registered and/or common law trademarks of Fortinet. All other product or company names may be trademarks of their respective owners. Performance and other metrics contained herein were attained in internal lab tests under ideal conditions, and actual performance and other results may vary. Network variables, different network environments and other conditions may affect performance results. Nothing herein represents any binding commitment by Fortinet, and Fortinet disclaims all warranties, whether express or implied, except to the extent Fortinet enters a binding written contract, signed by Fortinet's General Counsel, with a purchaser that expressly warrants that the identified product will perform according to certain expressly-identified performance metrics and, in such event, only the specific performance metrics expressly identified in such binding written contract shall be binding on Fortinet. For absolute clarity, any such warranty will be limited to performance in the same ideal conditions as in Fortinet's internal lab tests. Fortinet disclaims in full any covenants, representations, and guarantees pursuant hereto, whether express or implied. Fortinet reserves the right to change, modify, transfer, or otherwise revise this publication without notice, and the most current version of the publication shall be applicable. Fortinet disclaims in full any covenants, representations, and guarantees pursuant hereto, whether express or implied. Fortinet reserves the right to change, modify, transfer, or otherwise revise this publication without notice, and the most current version of the publication shall be applicable.