5 Categories of Questions for Evaluating AI Driven Security Solutions

Artificial intelligence (AI) has become a security industry buzzword so broadly applied as to become almost meaningless. When every product boasts AI capabilities, security decision makers may quickly become cynical, even in the face of the most exciting innovation shaping cybersecurity today.

To effectively evaluate AI based security technologies, it is first important to understand the meaning of AI and machine learning in the context of cybersecurity:

- **AI** is the broader concept of machines being able to carry out tasks in a way that humans consider intelligent
- **Machine learning** is a more specific application of AI that is based upon the principle that machines can perform assigned tasks intelligently if they are given access to data sets and allowed to learn for themselves — this process is often referred to as "training"

These definitions may raise more questions than they answer when you begin to apply them to how technology vendors are incorporating these capabilities into their products.

To further discern the AI messaging signal from the noise, here are five categories of questions you should pose to any security vendor when evaluating AI based security solutions.

1 Why does your security product include AI capabilities?

Vendors generally add capabilities to their solutions when they have discovered a new, better way to protect a computer or when they get pressure to expand their capabilities to meet market demand. The inclusion of AI is no different, so it's important to understand the vendor's motivation behind incorporating AI into their technology.

- Why does the product have AI?
- Is the AI performing a new capability or automating an existing capability?
- If a new capability, what is the goal of the AI in the product?
- How does including AI improve a product over similar, non-AI offerings?
- Does your AI replace older security capabilities in your product or is this additive?

2 How does your AI benefit my organization?

It is not uncommon for vendors to add capabilities into their product for reasons other than customer benefit, especially for solutions that may have been on the market for a number of years. It is important that you understand how each vendor's implementation of AI will improve your overall security.

- How will your use of AI specifically benefit my organization?
- Will the results show up in my bottom line and in employee productivity?
- How does the incorporation of AI impact the performance of the product and its use of endpoint computing resources?

3 How smart is your AI?

AI can be simple or complex. Simple AI is good at making decisions based on known information, like picking chess moves given the current state of a chessboard. It weighs existing data to determine an optimal result and can repeat this behavior through multiple iterations. It has no memory of the past and no great ability to anticipate the future.

Complex AI requires massive training data sets, a neural-net architecture, and considerable time to train appropriately. It excels at pattern matching and predictive tasks. Complex AI does not return quantitative answers (e.g. make chess move X), but instead returns qualitative answers (e.g. 89% chance this object is the same as other objects).

It's important to understand what type of AI the security vendor is using so that you have the right expectations of results.

4 How is the AI maintained?

The maintenance required to keep AI well-trained and relevant all depends on how the AI is being used. For instance, if the vendor is using AI to automate signature creation for new threats, the AI is typically maintained by the vendor and enables more frequent signature updates. This may not actually benefit the organization as it may result in more updates to the endpoints. Alternatively, if the AI is trained in the cloud and then deployed to the endpoints to make real-time decisions on threats without constant updates, the organization can benefit from consistent prevention with minimal maintenance.

- Where does your AI reside? Is the AI running in your cloud or running locally on the endpoint?
- How is the AI specifically used? Is the AI used to automate signature creation? Is the AI used to make real-time decisions on threats?
- When and where is the AI trained? Is it at the endpoint, or prior to deployment to the endpoint?
- How much maintenance, including employee training and active attention, does your AI solution require?
- How often is the AI retrained?

5 Can you provide a demonstration in our environment?

The true test of any security solution should be how well it performs for your organization. Any company selling a security product should be happy to demonstrate its performance within your infrastructure. Be wary of companies who only offer internal test results and bold assurances because the levels of aggressiveness in testing impact effectiveness in production environments and put the ownership on the end-user to adjust. This means the training of the model for the endpoint is incomplete.

- Does the AI provide levels of aggressiveness?
- What cloud dependencies does the AI rely upon to be effective? Can the AI be as effective offline as it can be online?
- Can the AI prevent never seen before malware on the endpoint without connectivity to the cloud?
- Can the AI prevent malware that its training set has never seen before?
- Has the AI been tested by a third party that confirms its ability to detect and/or prevent malware that did not exist when the AI model was trained?

Training an AI/ML Model

**AI
Math Model**

**Extract DNA
of Files**

**Transform,
Vectorize, and Train**

**Classify and Cluster
Good vs. Bad Binaries**

**Update AI
Math Model**