AVANAN

WHITE PAPER

Why You Need Artificial Intelligence for Email Security

Executive Summary

- The phishing threat has advanced and stopping these advanced threats today requires a solution that's built from the ground up using Artificial Intelligence (AI) and Machine Learning (ML).
- Yesterday's rules and signature-based solutions are not equipped to prevent these advanced threats. For example, of today's advanced phishing threats, 51% require AI/ML to identify and stop.
- A critical factor determining the effectiveness of an AI/ML algorithm is the dataset on which it is trained. Therefore, it must be trained on the most sophisticated and evasive attacks.
- For an AI/ML model to be successful, it must learn to improve over time

The Importance of Artificial Intelligence

Today's cybersecurity response teams are more crucial than ever. With threats coming from every direction and more data to secure, SOC teams are under immense pressure to detect and respond to threats.

Avanan's research has found that managing the email threat alone takes up 22.9% of the SOC's time. In addition, a survey found that 60% of SOC team members consider leaving their job or even changing careers altogether.

Why is there so much burnout and stress? With overworked and understaffed teams, along with often repetitive work, there's a high likelihood that an important detail will be missed. And as threats get more advanced, the problem compounds.

Phishing is, at the core, a cognitive attack. It attacks vulnerable humans with messages and visuals. Of course, humans can be trained to spot and stop phishing, and indeed, training is an integral part of any cybersecurity program. But Artificial Intelligence (AI) and Machine Learning (ML) are the only technologies that provide a cognitive-like behavior and adapt and learn.

It's easy to say that Artificial Intelligence can solve all your problems. Indeed, many companies come out and say just that. You'll see things such as Unsupervised Machine Learning or AI Email Security. Of course, artificial intelligence is many things. But by using state-of-the-art Natural Language Processing and image recognition, AI is the technology best able to identify phishing attacks at an accuracy that matches the best analysts.

We believe that Artificial Intelligence is an integral part of email security, as we'll detail in this report. What's important to note, however, is that simply applying AI is not enough.

Applying AI that is trained on the best data set, however, can be a game-changer.

With AI that is trained on a rich data set, the real-world impact is profound. The power of Artificial Intelligence comes from identifying unseen patterns in data. Avanan's API-based approach provides access to such data, including years of historical data. This is then compiled and fed into the AI, which allows for detections of security events that no human can match. For example, it can lead to a 99.2% drop in phishing emails reaching the inbox and a 71% decrease in end-user reports reaching the SOC.

Without AI, it is not humanly possible to keep up with the threat landscape, regardless of the size of your organization. Further, good phishing attacks look and read the same as regular emails. Frequently, only small variations in data patterns, unseen to the human eye or too complex for a human to observe, will reveal that it's an attack. Thus, artificial intelligence can work as a tool

to help save your SOC.

Artificial Intelligence can be used to free up time and headspace for your SOC, making their job easier and more enjoyable. Further, AI can be used to reduce threats and give better visibility into the attack landscape.

It is not a silver bullet. But when AI is infused through everything you do—advanced protection, improved response capabilities, better training—the phishing problem can be solved.

How It Works

Artificial Intelligence is a nebulous term that can mean many different things to many different people in many different contexts.

Let's take a non-technical example. This isn't necessarily realistic, but it helps illustrate an important point. Let's say your company hired someone, and their sole job was to read everybody's email, in real-time, across your entire organization to determine whether each email was phishing or not. If an email were clean, they would let it go right into the inbox. If it were phishing, the email would be quarantined.

Before this person starts their job, you would need to train them. First, you have to explain to them what to look for in a phishing email. To do this, you would take one million examples of phishing emails. You'd print them out and have the new hire review every single one. Then, you'd have them indicate why a particular email was phishing. Every email would have a different reason as to why it was declared phishing. In some emails, there would be clear phishing language in the body. In other emails, links would go to a newly registered domain. Some emails would come from individuals the company had never heard of before. In fact, there could be hundreds of reasons in every single email to indicate that it could be phishing.

After a period of time learning what to look for in a phishing email, you'd then test this person with real examples. If they did a better job of detecting phishing emails, with lower false positives, than what was in place beforehand, you'd turn your new "Real-Time Security Analyst" loose. Sometimes, this person would get something wrong. Then, through a feedback loop, you'd understand why and correct it going forward.

While this is a fictional example, this is essentially how AI/ML is used to identify and stop phishing attacks. It also brings us to the most important part of any AI/ML algorithm: The dataset on which the algorithm is trained.

In general, artificial intelligence learns to classify the data on which it is trained. So if it's trained on shoddy data, then there will be shoddy output.

That's why it's important to drill down into what the data set actually is.

For Secure Email Gateways, their training set is less focused, because they are just a first line of defense. They see a single email at a single point in time, giving them a lack of situational awareness. For other API-based companies, a lot of AI-based integrations are roadmap items, not in actual usage today.

For AI to work effectively, it needs to be trained on the best data set. For email security, it must be embedded within the cloud suite via API. Once embedded, the data set of cloud email security solutions is much richer. By being embedded, Avanan understands who the people being emailed are, the social graph, internal email, geo-suspicious login events, and more. Beyond that, as an inline security solution, Avanan's security layers run after Microsoft and Google's default security filters. That means Avanan's AI is trained on the specific attacks not caught by Google or Microsoft.

That posturing is incredibly important because hackers are constantly leveraging vulnerabilities to bypass default security. An example of this is HTML obfuscation methods. Avanan can train its AI on these methods, adding them to an ever-growing list of Attack Indicators.

In our model, we're constantly training and tuning our AI on the specific tenant. We have separate training sets for Office 365 and Google and separate models based on the direction of mail (inbound, outbound, internal). We use best-in-class AI algorithms and put our own inputs into them. By applying custom threat profiles for each organization, we can better tune our AI and keep phishing out. Instead of applying a one-size-fits all

approach, like Microsoft and Google do, we work.

Further, we're constantly learning from end-users. For example, every time an end-user reports an email as phishing, our analysts look at it. This is the supervised part of our solution. This allows us to learn across all of our customers and input that information into the algorithms. This feedback loop is how the AI/ML improves over time. This feedback comes in many forms, including emails being recategorized automatically or done manually by an administrator or analyst. Every time more data is acquired, it goes into improving the AI/ML.

In our dashboard, we aim to be as transparent as possible, so you can see which parts of the algorithm detected an email as phishing.

We also go further, showing which individual parts set alarm bells:

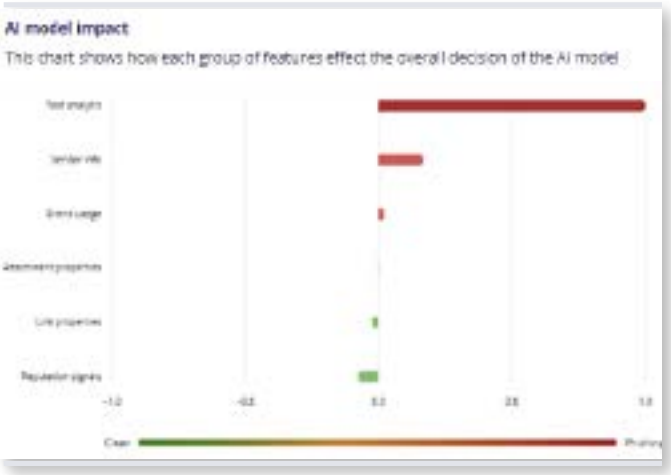

We also dig deep into the text of the email to show what specific words and phrases tipped off our AI:

By continually learning and understanding and auto-correcting as we go, our model busts false-positives and provides the highest catch rate possible.

How It Impacts Email Security

Today's email attacks are more sophisticated than ever. When relying on static filters alone, traditional email security misses 50% of attacks. Furthermore, static filters lead to a lot of false positives because these filters, by definition, don't exhibit the "intelligence" needed to stop today's advanced attacks. AI can see through the noise to better determine whether an email is phishing or not.

By utilizing advanced AI, Avanan prevents malicious emails from reaching the inbox. Take the below email as an example. This received an SCL score of 1 from Microsoft—about as clean as it gets.

By applying numerous Indicators of Attacks, Avanan was able to see this was phishing:

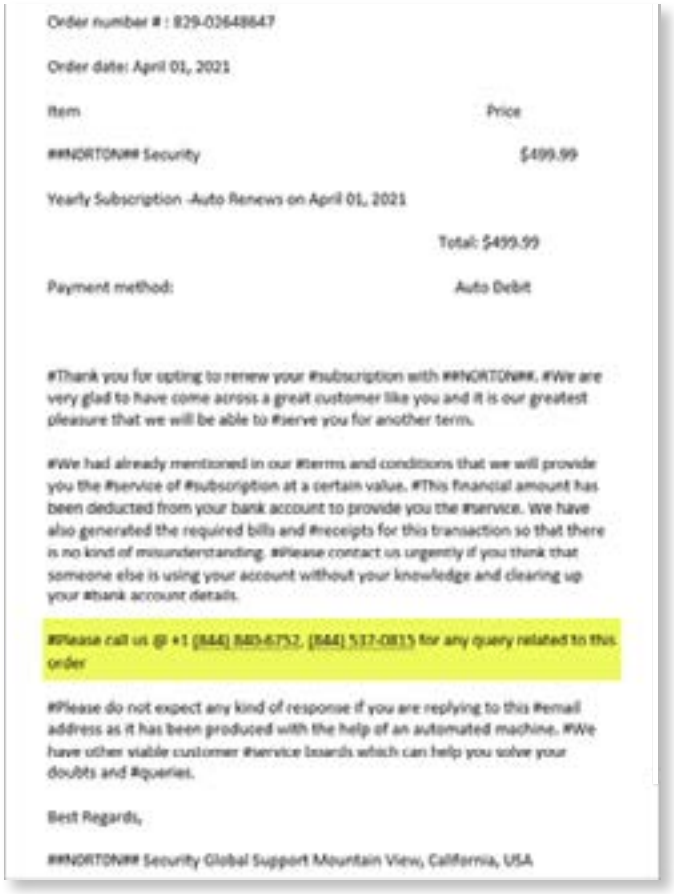

Order number # : 829-02648847

Order date: April 01, 2021

Item — Price

##NORTON## Security — $499.99

Yearly Subscription -Auto Renews on April 01, 2021

Total: $499.99

Payment method: — Auto Debit

#Thank you for opting to renew your #subscription with ##NORTON##. #We are very glad to have come across a great customer like you and it is our greatest pleasure that we will be able to #serve you for another term.

#We had already mentioned in our #terms and conditions that we will provide you the #service of #subscription at a certain value. #This financial amount has been deducted from your bank account to provide you the #service. We have also generated the required bills and #receipts for this transaction so that there is no kind of misunderstanding. #Please contact us urgently if you think that someone else is using your account without your knowledge and clearing up your #bank account details.

#Please call us @ +1 (844) 865-6752, (844) 537-0815 for any query related to this order

#Please do not expect any kind of response if you are replying to this #email address as it has been produced with the help of an automated machine. #We have other viable customer #service boards which can help you solve your doubts and #queries.

Best Regards,

##NORTON## Security Global Support Mountain View, California, USA

The social graph saw a low reputation sender score, a red flag for an email containing financial information. (Avanan's AI was able to suss out that the email was financial in nature.) Even though the email had no links, Avanan analyzed the email text and noted that the 844 toll-free number is typically associated with financial scams.

The Area Code 844 Scam: What It Is and How to Protect Yourself From It

Calls from 844 area code numbers are almost always dangerous phone scams

Another example of this is a recent attack that Avanan stopped. Traditional SEGs scan emails through filters that check URLs in emails against various static lists to determine how to treat each URL. Based on the decision made by an SEG's URL scan engine, the URL will either be allowed, wrapped, redacted, or kept from the user's inbox. Unfortunately, this approach has an overarching problem: Zero-day links, which are not found on any of these lists.

Further, there are a few ways around these static lists beyond zero days. One easy way is using CAPTCHAs like Google's reCAPTCHA. Since the main task of a reCAPTCHA challenge is to make content inaccessible to crawlers, the malicious nature of the target websites will not be apparent until the CAPTCHA challenge is solved. Beyond that, the reCAPTCHA service connects to IP addresses that belong to Google and are already in an SEG's allow list. Therefore, a phishing email gets through. Once the reCAPTCHA is solved, the end-user is directed to this phishing page:

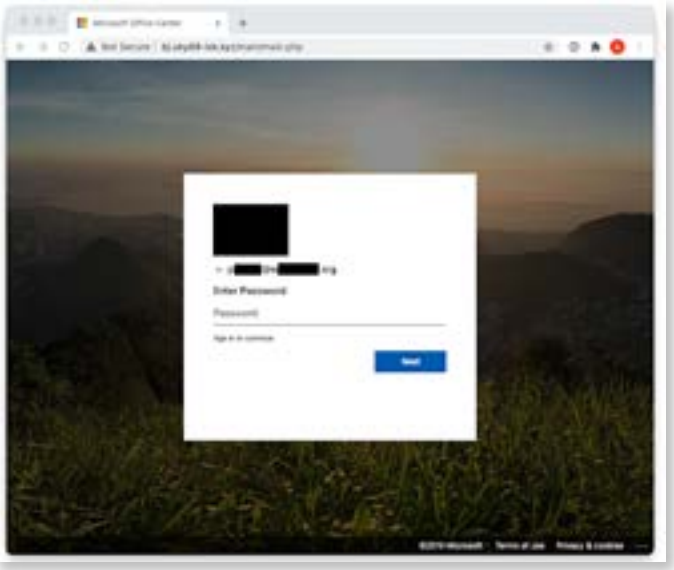

With advanced AI, Avanan was able to detect this as phishing for these reasons:

- Links with suspicious patterns
- URL and UTF encoding in the email's headers and attachments
- Low Sender Reputation

Consider, as well, the attack form known as Business Email Compromise. This is a fast-growing trend that has caused major damage. Gartner found that BECs increased by nearly 100% in 2019 and through 2023, predicts that BEC attacks will continue to double each year, at the cost of over $5 billion to its victims. Further, the average BEC payment nearly doubled between the first and second quarters of 2020. It's now at $80,183, on

average. And the FBI has noted that, between 2014-2019, they saw claims of over $2.1 billion in losses from BECs. Avanan's research has found that BECs make up over 20% of all attacks.

BEC attacks are popular because they are simple to execute and incredibly destructive. A successful attack often results in the attack having full access to a company's entire cloud. On the other hand, spoofing a trusted user requires no malware or malicious URL to convince a recipient to share valuable information or send significant amounts of money.

However, in order to stop these, what's needed, according to Gartner, are the following:

- Deploy inside the cloud email server, typically via API
- Offer internal email protection, between users
- Use advanced machine learning for internal email context
- Offer account takeover protection

Without AI, when the solution sees an email from the 'CEO to the CFO,' it will be the first time it has seen such a conversation, giving it no context to know if it's out of the ordinary or not.

Avanan approaches BECs differently. Within hours of the first deployment, Avanan's AI scans a year's worth of email conversation to build a reputation network, the type of internal context that alerts the AI to something suspicious.

Phishing is evolving every year. Hackers will continue to find vulnerabilities in existing security systems. They will continue to find new and innovative ways to bypass this protection and get into the inbox. And once they get into the inbox, they find unique and sophisticated ways to get the user to click and enter the information they're looking for. Finally, hackers will continue to focus on industries that don't have proper defenses, preferring to gather up the lowest hanging fruit at a higher rate than hoping to bring in the big whale.

As email continues to migrate to the cloud, and as the market continues to shift away from SEGs towards inline API security, leveraging that positioning by training AI on the most sophisticated attacks is the best way forward.

THE AVANAN DIFFERENCE

Avanan differentiates itself with its AI detection methodology. Avanan implements over 300 (and growing) Indicators of Attack (IoA) to determine whether an email is phishing. A non-exhaustive list of some of the indicators Avanan uses:

- Phishing language in an email's subject and body
- Encoded content, such as scripts to encode or decode Base64, Morse, etc
- HTML obfuscation methods, such as ZeroFont and baseStriker
- Existence of a Crypto wallet address

Avanan leverages multiple types of AI and ML. It starts with combining machine learning and static analysis, to determine if an email is phishing and, if so, what type of phishing it is..

To do so, Avanan employs several techniques, including, but not limited to:

- Social graph and sender reputation
- Language Processing
- Anomaly Detection
- Anti-Impersonation/Conversation

What further separates Avanan's AI analysis is the fact that it lies inline. That means it sits behind default security but before the inbox. Avanan's AI, then, is trained on phishing emails that got past Microsoft EOP or ATP, Google, and any other SEG a client may have installed.
By doing so, Avanan is better equipped to stop the newest, most sophisticated, and evasive attacks in the wild today.

This patented approach has been a boon for customers. Avanan customers see a 99.2% reduction in phishing attacks reaching the inbox. In addition, the SOC is less busy, with a 71% reduction in phishing-related alerts sent to that department.

An Avanan analysis found that 50% of all attacks would be missed without AI.

Additionally, in a study of 360 million emails, Avanan was 15x more effective than legacy gateways like Proofpoint, Mimecast, and Barracuda, and 18x more effective than Microsoft ATP.

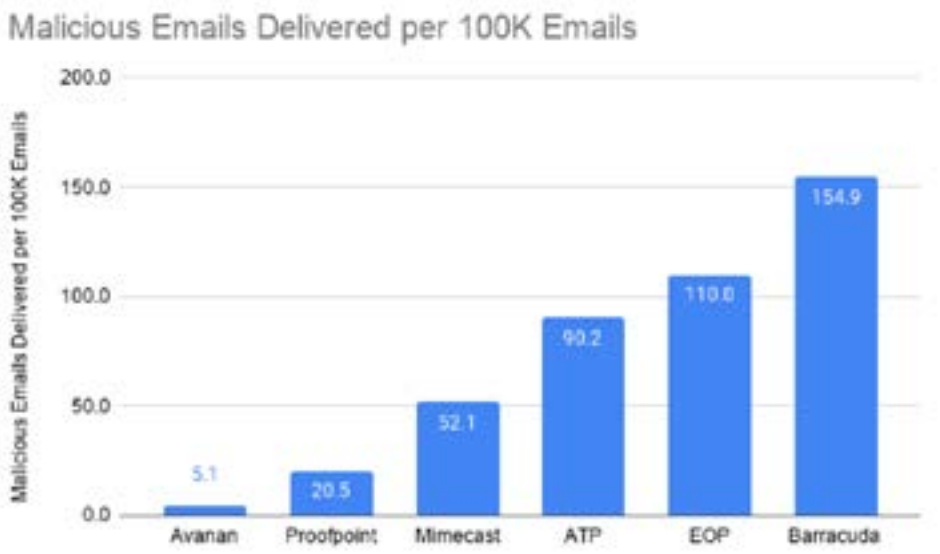

Conclusion

Today's email threat and SOC overloaded can be unbearable. Adding advanced Artificial Intelligence into the mix can be a real difference-maker. Advanced threats are coming from every direction. The only way to stop these threats is through AI.

However, it is only a difference-maker if it is the right AI, trained on the best dataset and learns/improves over time. That sort of AI, which Avanan employs, can then be used to improve security, stop phishing attacks and take a significant load off the SOC.

With Avanan, over 50% more advanced threats can be blocked. And every day, through our feedback loop, the AI is getting better and better.

The world of advanced threats can be scary. Having those threats reach your inbox can be even scarier. It's why it's more important than ever to deploy the right AI for your organization.

The right AI is the Avanan AI. See the difference it can make today.

AVANAN

Avanan is a cloud email security platform that pioneered and patented a new approach to prevent sophisticated attacks. It uses APIs to block phishing, malware, and data leakage in the line of communications traffic. This means Avanan catches threats missed by Microsoft while adding a transparent layer of security for the entire suite that also protects other collaboration tools like Slack. The solution has been recognized as the top-rated cloud email security solution by customers and can replace the need for multiple tools that surround email and file sharing.

See how Avanan can transform your security avanan.com/demo