SCLUTIONS

TECHNOLOGY RESEARCH REPORT

Artificial Intelligence

A relative reality

By Esther Shein

Introduction

We have all heard the expression, "Go big or go home." Today's conventional wisdom says if vendors are not considering implementing artificial intelligence (AI) into their security software, they might as well not bother innovating. The reality is there is no sliver bullet that solves all of the CISO's security needs.

Pundits today tell CISOs that pretty much everything from endpoint protection and perimeter monitoring to security information and event management (SIEM) and malware protection will include AI components, which are touted as the panacea for stopping attacks, protecting your network from advanced persistent threats, enhancing your threat intelligence and making all your security technology smarter, faster and more effective. That's a big promise, but can AI deliver?

SC Media set out to differentiate what is marketing jargon from what CISOs actually can do today with AI-enhanded security products. To determine the veracity of the claims, we spoke to CISOs and other security leaders and analysts to find out which products embedded with AI deliver on their promises.

Then we asked vendors whose products were selected in five categories to describe briefly how their tools actually use AI. Those categories include SIEM, Endpoint Protection, Network Performance Monitoring/Intrusion Detection, Antimalware, and Antiphishing Software. We also asked industry experts questions such as what should CISOs expect from AI offerings, how do they ensure they are getting what they expect, and what should they know before they buy?

OUR EXPERTS

Doug Barbin, principal and cybersecurity practice leader, Schellman & Company, LLC
Paul Hill, senior consultant, SystemExperts Corp.
Mark Horvath, senior research director, Gartner
Chris Kissel, research director of worldwide security products, IDC
Taylor Lehman, CISO, Athena Health
Gary Miller, senior director of information security, TaskUs
Fernando Montenegro, senior industry analyst, 451 Research
Eric Ogren, senior security analyst, 451 Research

AI today is still a work in progress

Experts caution that AI is no silver bullet, but implemented correctly it can counter many of the advances made by cybercriminals

There is no doubt Artificial Intelligence (AI) deployments are on the rise, but opinions remain mixed about what this term means in practice and how viable it is at present, especially when it comes to security. Research firm Deloitte broke out AI adoption into four distinct categories in its 2018 *State of AI in the Enterprise* report: machine learning (ML), deep learning, natural language processing, and computer vision. Fifty-nine percent of respondents reported using enterprise software with AI. Additionally, 44 percent said the benefits of AI are that it enhances current products.

The same report found that "executives are commonly concerned about the safety and reliability of AI systems as well."

Specifically, a little over half of the respondents said they are concerned about the cybersecurity vulnerabilities of AI, while 43 percent of respondents rated "making the wrong strategic decisions based on AI/cognitive recommendations" in their top-three concerns.

Thirty-nine percent cited failure of an AI system in a mission-critical or life-or-death situation.

The promise of AI

Ultimately there appears to be no single definition of AI, as marketers and analysts tend to use AI and ML interchangably and ofen with slightly different connotations.

"Placing strategic decisions or mission-critical actions entirely in the hands of an AI system would certainly entail special risks," the report says. "Entrusting AI systems with such responsibilities remains rare today, however."

> "A lot of end user companies ... go buy a product and say, 'hey, it has AI' and turn it on, but does that customer understand that some data may be sent off to that vendor?"
>
> *Paul Hill, senior consultant, SystemExperts Corp.*

Yet vendors continue to promote their AI capabilities. Some of the decriptions vendors use for their products include: "AI baked in," "unlock the potential of your data with AI" and "zero false positives."

Many tout the use of AI to address the dearth of cybersecurity professionals with the promise that AI can replace Level 1 and 2 tech support engineers — and perhaps it can.

"Both artificial intelligence and machine learning are being applied to cybersecurity use cases to address fundamentally the same problem: the cybersecurity workforce shortage," notes an IDC Market Perspective *Artificial Intelligence and Machine Learning in Cybersecurity: Creating Meaning with the Terms* from February 2019. "One is looking to make security professionals more effective; the other is looking to improve the efficiency."

AI is also increasingly being viewed as another tool in the arsenal for combating the increasing sophistication of cybercriminals since current security controls are not doing enough to defend networks, as evidenced by the rise in breaches.

At the same time, though, it can be hard to know which products offer the promise of AI and ML and which are hype. "Unfortunately, the recent uptick in marketing hyperbole around the generic terms 'machine learning' and 'artificial intelligence' often gets in the way of understanding the benefit-risk trade-offs," observes Gartner, in the 2018 report *Lift the Veil on AI's Never-Ending Promises of a Better Tomorrow for Endpoint Protection*.

What you should know

Unlike a decade ago when an organization bought a security information and event management (SIEM) application and their data stayed on premises, today, a large number of systems are cloud-based, and cybersecurity vendors need larger datasets

AI today is still a work in progress continued

to implement effective AI. The caveat "buyer beware" is a good reminder as organizations begin to deploy more cloud-based security offerings embedded with AI, says Paul Hill, a senior consultant at SystemExperts Corp., a Sudbury, Mass-based security and compliance consultancy.

Customers should be asking their vendors how they are collecting data if they are planning to use it to feed machine learning algorithms, he advises.

Another question they need to ask is whether their data will be anonymized and filtered so that no sensitive data will be sent to a third party. "It's not talked about much," Hill observes. "And a lot of end user companies … go buy a product and say, 'hey, it has AI' and turn it on, but does that customer understand that some data may be sent off to that vendor?"

In terms of the vendors, Hill wonders whether they are proactively explaining to customers how they are collecting their data and whether customers understand that their data might be used.

There is no denying that AI means different things to different people. "I'm a little skeptical any time I hear 'AI baked in,' and if [vendors] can't tell me what that means fairly quickly, I've moved on," says Doug Barbin, principal and cybersecurity practice leader of Schellman & Company, LLC, an independent security and privacy compliance assessor in Sacramento, Calif. "AI is a very, very broad term. It's like saying 'cloud,'" he adds. "Is it machine learning, is it natural language processing, is it robot process automation?"

How to ensure your AI product is performing as promised

Marketing teams tend to create FUD, or fear, uncertainty, and doubt, says IDC Research Director of Worldwide Security Products Chris Kissel. There are two simple metrics an AI system should be able explain: the mean time to detect an incident and the mean time to respond to it, says Kissel.

Conducting a proof of concept pilot can help set expectations appropriately, observers say.

"The most important thing someone can do is when you go to buy a tool, get three to four vendors and tell them specifically what your terms are," he says. Terms might include: "We want to be able to find an anomaly in this part of the network." Often, a vendor will tell a security operations center (SOC) team what they should be looking at in their environment, "so a buyer has to own the process," Kissel says.

> "So there is a place for machine learning in [identifying] insider threats and anomaly detection, but it's not the silver bullet people think it is. You'll get a heck of a lot of false positives or false negatives depending on how you tune your system."
>
> *Fernando Montenegro, senior industry analyst, 451 Research*

"I always recommend people pilot these things first for six to eight weeks," agrees Mark Horvath, senior research director at Gartner. "It is not enough to just look at the product interface and say, 'It's sophisticated, I'm going to buy it.'"

In some cases, piloting the product might mean having to invest in staff or additional infrastructure to get the most out of it, he adds.

Users say proof of concepts are key. "We test it," says Gary Miller, senior director of information security at global outsourcing provider TaskUs, headquartered in Santa Monica, Calif. "In a demo, I would definitely want to sync up Active Directory to the tool to continuously inform it of current credentials and in testing that in a proof of concept, we'd provide [the vendor] with fake credentials and have them put it out there and see how it searches."

And Miller adds, "I'd also look to them to prove to me how it works."

The promise of AI is that it enables automation and more effective decision making, he says, but admits it is a "major challenge" to boil out the false promises from the reality. "A lot of vendors want to explain this advanced algorithm they have that they can't allow you to see — but trust them, it's very effective," he notes.

Kissel maintains that there is no "artificial intelligence truly in cybersecurity" right now, citing the Turing Test, which states that true AI must be equivalent to, or indistinguishable from human performance.

Gary Miller, senior director of information security, TaskUs

"It's early days, the stuff is evolving and evolve rapidly," he says. "The incident response piece has been very visible, and I think, the most mature in AI in having an influence and impact."

Taylor Lehman, former CISO of Wellforce and Tufts Medical Center and now CISO at Athena Health, says the best proof of concept he did was the time he picked out 13 antivirus platforms "and downloaded every virus known to man, and threw them at each solution and watched them" to see which one caught which virus. "Then you look at what it didn't block," he says. "The proof is in the pudding."

There are other considerations, of course, such as how easy a product is to deploy and update, he adds.

AI today is still a work in progress continued

> Machine learning is a great way of reducing the number of false positives you have to deal with. You can't get it down to zero, but you can reduce it."
>
> *Mark Horvath, senior research director, Gartner*

By getting deep into a security platform, organizations can understand if an AI-embedded feature set not only improves the security controls they have in place, but whether there's an added cost for it, notes Miller.

AI for intrusion detection

There are machine learning techniques that are particularly good at anomaly detection, but the challenge is knowing the underlying assumptions you're making about the data you have, says Fernando Montenegro, a senior industry analyst at 451 Research.

For example, if your security team is monitoring the network and suddenly sees a spike, there could be a number of legitimate reasons why, he says. These might include a new initiative within the business or a change in how often you decide to monitor, or it could be a new employee going through different databases to learn their job.

"There's a whole slap of things that can be anomalous that aren't necessarily malicious," Montenegro says. "So there is a place for machine learning in [identifying] insider threats and anomaly detection, but it's not the silver bullet people think it is. You'll get a heck of a lot of false positives or false negatives depending on how you tune your system."

AI for Phishing

Anomaly/fraud detection and insider threats are good use cases for AI, says Horvath. "Machine learning is a great way of reducing the number of false positives you have to deal with," he says. "You can't get it down to zero, but you can reduce it."

Surveys indicate that traditional virus attacks are down, but phishing is way up, which begs the question: How effective can AI be at spot phishing attacks?

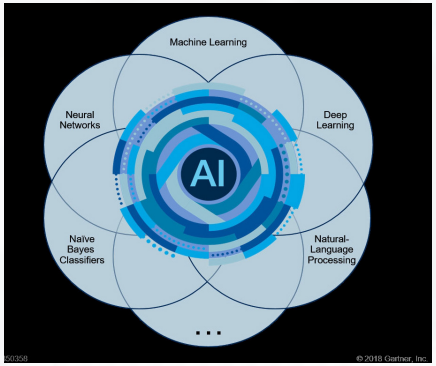

While building specific products that use AI components is a tricky, time-intensive activity, almost all commercial EPPs (endpoint protection platforms) make use of some of these technologies for malware detection in easy-to-use forms that require little to no understanding of or interaction with their AI components. Source: *Gartner Lift the Veil on AI's Never-Ending Promises of a Better Tomorrow for Endpoint Protection*

AI-based antiphishing software looks for patterns either in the email or within the malware, says Horvath. "One of nice things about AI is it can examine polymorphic malware," which looks at very specific parts of the virus itself that it has learned to recognize. So it recognizes little bits and pieces humans would not see. AI replaces looking at signatures with looking at even smaller aspects of the code that it recognizes as malware, but would be invisible to a human and wouldn't show up in a signature scan, he notes.

The problem is that as effective as AI is in reducing the current generation of malware, "the people writing the malware are also writing new malware to get at this, so it's a continuous cycle," he says.

Research is being done using natural language processing to determine if a sample of text was written by a person or generated by a machine, says Barbin, but he isn't sure whether that should be called AI or language heuristics.

Echoing Horvath, Barbin says "the same AI tools you can throw at that problem could also be leveraged by attackers to improve the phishing attempts."

This is especially true when an attacker is able to gather samples of emails to and from particular individuals, which helps them replicate language effectively, he says.

As always, it comes down to the human element as phishing attempts become more sophisticated, Barbin says.

"It is truly tricky stuff to figure out what is a bot-driven email and what is legitimate," agrees Eric Ogren, a senior security analyst at 451 Research.

Doug Barbin, principal and cybersecurity practice leader, Schellman & Company, LLC

TaskUs' Miller says he believes "phishing is one of those wins for AI that's very visible to me."

AI for Endpoint Protection

Hospitals have thousands of endpoints and anything that can improve anomaly detection — especially in a life or death situation. So far, most security categories with an AI play are in the infancy stage while endpoints are in the "toddler years," according to Lehman, who has been using AI to look at behavior on endpoints as well as on the network.

AI on the endpoint is "the furthest along and most proven in the sense that it works and meets the objectives it sets out for itself," he says.

The next step for Wellforce is to look at an individual's identity, the identity of the

AI today is still a work in progress continued

computer being used, and how the system behaves, he explains. That might mean three people on the network all exhibit similar behaviors and then a fourth person appears whose actions are atypical for that type of employee. "The goal of machine learning is to figure out the patterns of data and the volume of data," and determining what is acceptable behavior when more people share data and what is anomalous, Lehman explains.

When Wellforce began using AI-based endpoint protection, for monitoring "we almost immediately got two full-time people back," he adds. Previously, "It was like there was a ghost in the system and once we stuck a product in that was prevention focused … we stopped it and started seeing a decreased volume of users calling in." With AI, about 15 percent more of Wellforce's devices received higher level of security controls than they previously had, he says.

For all the skepticism Schellman's Barbin has about broad claims proclaiming the benefits of ML and AI in security, he also finds the technology to be very useful for endpoint protection because of the ability to apply clustering techniques around samples.

"Where I think it doesn't help is when people start assuming [AI] can do more than it can do," he says, noting that "There is always a trade-off between false positives and false negatives. So there's always going to be an incident."

For example, a SOC operator might get an alert from an endpoint saying it has malware. "If it's true, great, if false, you as an organization have to spend resources to go in there and check out the false positive" and either ask the user to re-image their laptop or confiscate it while you investigate, he says.

> If you think about it, perimeter threat-facing filters whiff on many advanced threats so the only way to catch active threats is to analyze traffic for signs of threat behaviors. It's hugely effective."
>
> *Eric Ogren, senior security analyst, 451 Research*

"What this means is there's a cost to the organization," says Barbin. "Not only are you paying for the time for the security team to handle this incident, you're also dealing with lost productivity for the employee."

On the flip side, that SOC operator might catch an attacker doing something else that they were not looking for. AI works as part of a layered defense, he says.

AI can deal with immense volumes of data, which is just not feasible for a human to look at, Barbin says. There is just too much data coming in and too many malware samples. "Humans can augment the process," he says, "but you need AI to do some triage on the data."

AI for SIEM and more

Most SIEMs now come with analytics to detect problems, assemble threat timelines and consolidate alerts, says Ogren. "It is hugely effective."

He says he has seen a "resurgence in network security based totally on AI. "I used to call this 'network traffic analytics' but believe NVDR (network visibility, detection and response) speaks better to its value. The network sees everything, providing operations teams insight into all devices using the network" using AI to detect threats, and then using AI to help correlate information to accelerate remediation, he says.

It is still the proverbial early days for NVDR, Ogren adds, "but if you think about it, perimeter threat-facing filters whiff on many advanced threats so the only way to catch active threats is to analyze traffic for signs of threat behaviors. It's hugely effective."

Eric Ogren, senior security analyst, 451 Research

The only drawback, he says, is remediation still requires IT people to fix a host somewhere, but security needs the visibility that can only come from NVDR to make good prioritized decisions.

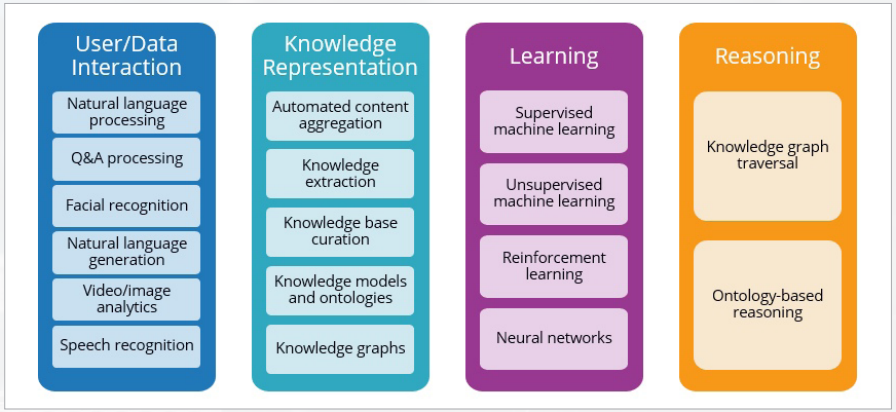

This set of technologies uses natural language processing, machine learning, knowledge graphs, video/image/speech analytics and other technologies to answer questions, discover insights and provide predictions or recommendations. Applications using these technologies hypothesize and formulate possible answers based on available evidence, can be trained through ingestion of vast amounts of data/content and automatically adapt and learn from their mistakes and failures.

Source: IDC Market Perspective Artificial Intelligence and Machine Learning in Cybersecurity: Creating Meaning with the Terms

Artificial Intelligence: A relative reality

AI powers next generation SIEM systems

Here is a look at how AI is being used in select products in five categories from the vendor perspective. The included vendors were identified by analysts, consultants and CISOs as offering the best-of-breed offerings that incorporate artificial intelligence and machine learning.

LogRhythm's SIEM processing platform and cloud-based security analytics offering, CloudAI.

SIEM

IBM — According to research from IBM, security teams without AI technology sift through more than 200,000 security events per day on average, leading to more than 20,000 hours per year wasted chasing false positives. QRadar incorporates the MITRE ATT&CK framework, a compendium of knowledge about hacker techniques and behaviors, and uses it to analyze cybersecurity events and network flows.

— Rueben Rodriguez, product marketing manager, IBM Security

LogRhythm — LogRhythm's SIEM processing platform and cloud-based security analytics offering, CloudAI, can aggregate user accounts and identifiers into a singular user identity. CloudAI then learns models that characterize a user's behavior on the network. When a user identity subsequently changes behavior in a way that could represent a security threat and organizational risk, the platform can notify security analysts of the suspicious activity.

— Phil Villella, chief scientist and co-founder

Sumo Logic — The Sumo Logic platform leverages a variety of machine learning techniques to accelerate actionable insight in large volumes of data, to derive deep insight on comparative patterns in different data sets, and to identify insights across customers. These pragmatic use cases have helped customers gain instant insight into millions of logs to accelerate investigations from hours to minutes and to detect security threats, which evade traditional thresh-holding and correlation techniques.

— Dave Frampton, vice president of Security Solutions

Splunk — Splunk user behavior analytics (UBA) provides insider and advanced threat detection to help organizations find unknown threats and anomalous user behavior. Because Splunk UBA is powered by machine learning, it can determine automatically when users are taking unusual actions and detect threats that are made up of patterns of abnormal behavior.

— Girish Bhat, director of product marketing

Artificial Intelligence: A relative reality

Endpoint protection becomes more intelligent

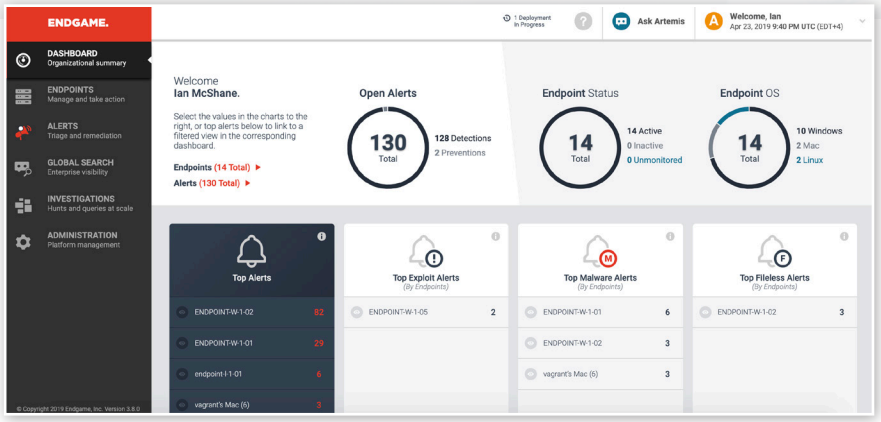

The Endgame Dashboard showing open alerts, status and endpoints.

Endpoint protection

CrowdStrike — CrowdStrike Falcon leverages machine learning to identify new and existing threats and provides coverage in many areas ranging from pre-execution file analysis over host execution behavior to macroscopic behaviors at the network level over long durations. AI models on the endpoint allow rapid detection of threats, even when endpoints are offline. To train AI models, Falcon taps into a comprehensive collections of security data including four million events streaming from 176 countries into the Falcon cloud every second, contextualized in the ThreatGraph graph database by tracking more than six trillion connections.

— Sven Krasser, chief scientist

Blackberry Cylance — The Cylance AI Platform is a unified technology architecture built on continuous integration and continuous delivery (CICD) principles to drive a high velocity of capabilities through native AI endpoint security products and delivers continuous threat prevention, detection, and response regardless of whether the endpoint resides on or off the network. The platform delivers an average 25-month predictive advantage, a measure defined as "the time difference between the creation of the model and the first time a threat is seen by victims and security companies protecting those victims."

— Sasi Murthy, vice president of products and solutions marketing

Endgame — The Endgame platform features Artemis, an AI-powered security mentor that relies on a natural language interface and automates SOC analyst actions to guide users of any skill level to detect and respond to advanced attacks through a simple conversational interface. It also features MalwareScore, our first product leveraging machine learning, which detects malware based on common patterns in the structure of the file. It makes this determination pre-execution, before the process is allowed to execute/be opened.

— Hyrum Anderson, chief data scientist

Sophos — Powered by deep learning technology and an intelligent neural network, Sophos Intercept X Advance with endpoint detection and response (EDR) allows businesses to detect, prioritize, investigate and respond to incidents rapidly. Deep learning enables Sophos to stop never-before-seen malware, unknown exploit variants, stealth attacks and ransomware attacks. Based on threat intelligence from SophosLabs, Sophos Intercept X Advanced with EDR learns the observable threat landscape, processes hundreds of millions of samples, and makes more accurate predictions at a faster rate than traditional machine learning products.

— Seth Geftic, director of product marketing

Artificial Intelligence: A relative reality

Network visibility, detection & response drive the future

Network visibility, detection & response

ExtraHop — We see CISOs investing in machine learning but remaining justifiably skeptical of AI. Right now, AI in security is still mostly artificial and not that intelligent, and, frankly, there are other investment areas with a better return on investment. That said, a key value of our approach is the integration of ML detections with guided investigations, so it's hard to say precisely how much can be attributed exclusively to ML, but users offered up productivity improvements in the 50-95 percent range in the areas of time to detect, time to resolve, and number of resources required to resolve.

— Tom Stitt, senior director product marketing - security

Palo Alto Networks — The company enhanced its ability to detect phishing sites by adding deep learning/image recognition to the product earlier this year. We take an instant snapshot of the millions of web pages every day, breaking the content apart pixel-by-pixel to identify evasive phishing sites. It gives us a high degree of confidence whether a web page is a hidden phishing site, allowing customers to automatically enforce protections when found. This technique has improved our phishing coverage by 61 percent, all while reducing false positives.

— Navneet Singh, product marketing director

ThousandEyes — ThousandEyes' network monitoring product processes collective data from tens of thousands of concurrent global monitoring tests across the internet. Our traffic outage algorithms identify ISP outages that have geographic and network topology scope, automatically correlate those outages to application and user experience issues, present direct root cause indicators

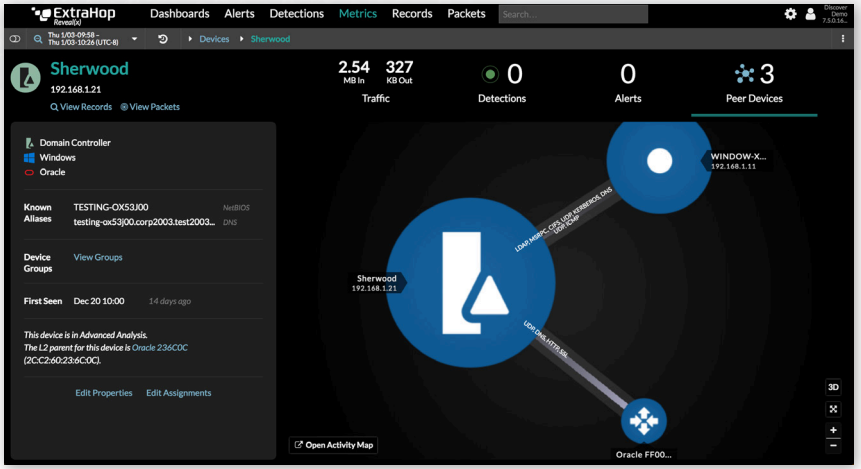

ExtraHop Device Overview Domain Controller provides stats on traffic, detections, alerts and peer devices.

in dashboards as well as detailed read-outs of the nature and scope of the outage into the specific network path visualization context of every impacted monitoring view across our entire base of customers.

— Alex Henthorn-Iwane, vice president, product marketing

Vectra Networks — The company is applying AI to metadata rather than using deep-packet inspection or analyzing logs or NetFlow records. The outcome of using AI is a measurable reduction in the workload and time to detect and respond to hidden threats. We have been measuring this workload reduction in our semi-annual Attacker Behavior Industry Reports based on anonymized metadata that our customers share with us.

— Mike Banic, vice president of marketing

Artificial Intelligence: A relative reality

Kicking antimalware intelligence up a notch

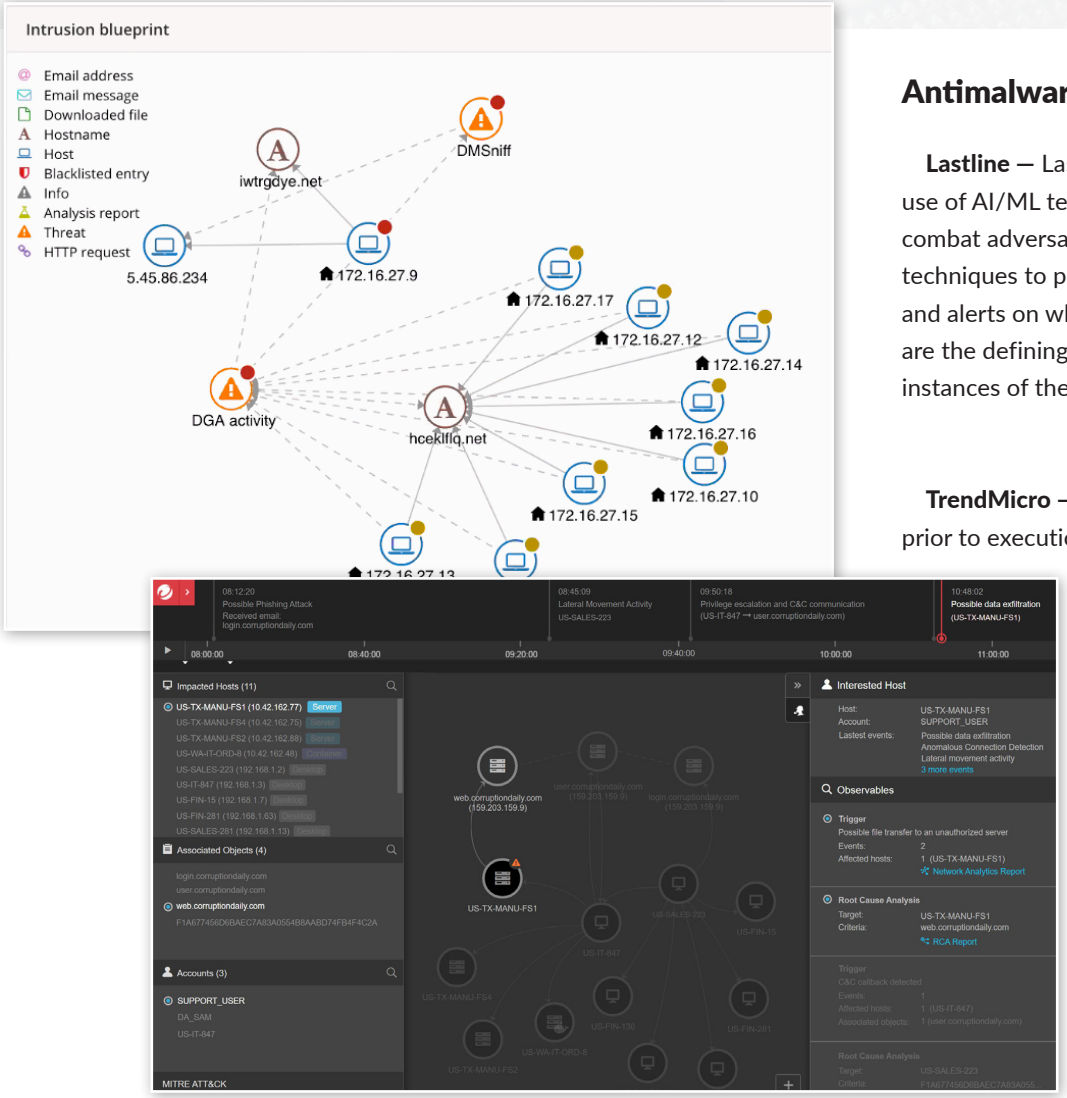

Antimalware intelligence

Lastline — Lastline is focused on adversarial machine learning, which is the use of AI/ML techniques in domains that actively resist the learning process. To combat adversarial machine learning, Lastline Defender uses a combination of AI techniques to perform both anomaly detection, which is to learn what is normal and alerts on what is abnormal, as well as misuse detection, which learns what are the defining characteristics of attacks are and use the knowledge to detect instances of the attacks.

— Giovanni Vigna, co-founder and CTO

TrendMicro — Pre-execution machine learning (applied to file-based samples prior to execution/opening) has increased Trend Micro's detection ability, especially relating to new malware families and some types of malicious scripts embedded in documents. Run-time machine learning is used in parallel with rule-based behavioral detection techniques in order to detect malicious activity, including ransomware behavior and file-less threats abusing Powershell or other scripting tools.

— Eric Skinner, vice president of solution marketing

Screenshot of an intrusion blueprint from Lastline (top) and the TrendMicro Investigation Workbench (bottom), showing a simulated phishing attack with data exfiltration.

AI and ML power new antiphishing analysis and detection tools

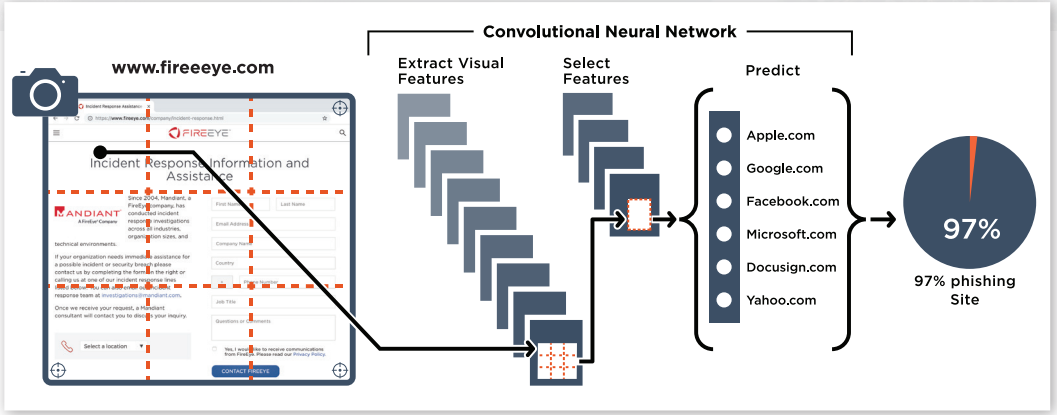

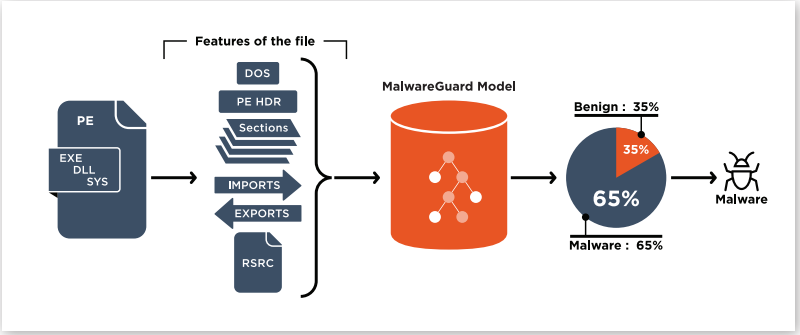

This is how FireEye PhishVision (top) and FireEye MalwareGuard (bottom) work.

Antiphishing analysis and detection

FireEye — By training AI and its subset machine learning, to understand the difference between legitimate benign and malicious email, FireEye has real-time detection and blocking of malicious emails undetectable by traditional engines without updates. Additionally, AI in FireEye Research Labs rapidly identifies threat landscape changes and subsequently deploys detection updates without humans.

— Matt Allen, vice president, FireEye Research Labs

Proofpoint — The Proofpoint Nexus security and compliance platform correlates threat intel from more than five billion daily emails, 200 million social media accounts, and 300,000 daily malware samples, while using a threat scoring model based on type/spread/targeting. Proofpoint's Targeted Attack Protection (TAP) uses multiple machine-learning engines spanning threat classification, composite multistage threats, relationship and communication classification, key employee classification, and evolving page design classification. TAP applies machine learning-led predictive analysis of URLs, detection of malware-free attacks like BEC and credential phishing, and password cracking for zip files.

— Mark Guntrip, group product marketing director

Artificial Intelligence: A relative reality

Up and comers

Here are some companies that our experts believe are worth watching over the next couple of years.

The future of AI?

Gartner's Horvath has industry veterans Securonix Inc. and Exabeam Inc., two AI-embedded SIEM products on his radar. 451 Research's Ogren seconds Exabeam.

Securonix and **Exabeam** got their starts in the user behavior analytics space, at the time offering AI insights and their own data lake to augment SIEMs, Ogren notes. Now their strategy is to provide "full-fledged SIEMs for organizations that understand analytics empowers SIEMs to better detect account take-overs, misconfigurations, consolidate alerts, and confirm corrective actions are in place."

Other newcomers he is eyeing in the NVDR category include DataVisor Inc. (for fraud detection), U.K.-based Netacea and PerimeterX.

DataVisor applies unsupervised AI to user activity, such as within e-commerce web sites. "Since fraud tactics are always changing, static filters are insufficient in reducing business risks of financial losses and customer dissatisfaction. DataVisor analyzes all user activity for accounts acting in unusually coordinated manners – a sign of a fraud campaign in action, he says.

Netacea is a bot management provider that takes on the challenge of separating authorized bots such as search engines and business partners from questionable bots without causing user friction, with AI, Ogren says. "Unlike others the vendor does not use Javascript in trying to fingerprint the source computer," he says.

Miller from TaskUs believes there are lots of CISOs saying they aren't yet using machine learning products who unknowingly are, because not all vendors promote that. Like Kissel, he thinks everyone will be aware they are using AI and machine learning in cybersecurity within five years.

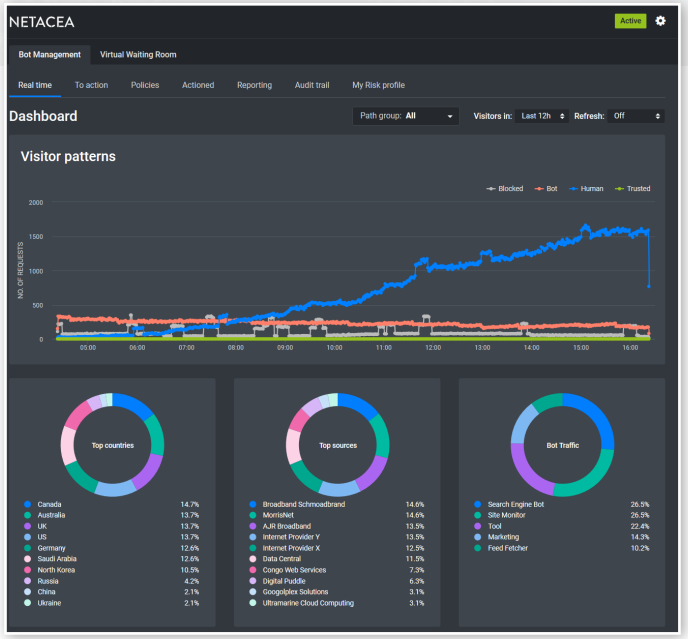

Netacea platform dashboard with at-a-glance visitor patterns, including bot, blocked, human, and trusted activity.

Lehman also believes "we're in our infancy" right now, and that there are "very few models that really work, but the ones that do deliver on the promise." Echoing Kissel concern that AI is not yet ready for mass consumption, he says the next evolution of AI will be when AI models can be trusted to make decisions without human intervention. But then, he muses, "eventually, we'll need AI to manage and monitor the AI."

Sponsor

Rise Above the Noise.

ExtraHop provides cloud-native network detection and response for the hybrid enterprise. Whether you're investigating threats, ensuring the availability of critical applications, or securing your cloud investment, ExtraHop's breakthrough approach helps you rise above the noise so you can protect and accelerate your business.

For more info, visit extrahop.com