Secureworks®

WHITE PAPER

7 Ways AI Can Automate and Improve Vulnerability Management Operations

Introduction

Many business and consumer products have benefited from the application of artificial intelligence technologies in the past decade, but legacy vulnerability management products have yet to fully embrace the potential of AI to streamline operations, and more importantly, greatly reduce enterprise vulnerability risk. In this paper, we'll discuss 7 ways in which AI can be leveraged by security and IT teams to improve traditional vulnerability management activities.

01 Identification of "Outlier Assets"

A favorite technique of pen testers - and hackers alike - when attacking an enterprise network is to locate outlier assets (e.g. laptops, connected devices, servers, routers), or assets that are unique in some way. These assets are attractive to intruders because they know that unique assets can often be soft targets, or indicate enterprise-specific information and patterns, hence exposing valuable information, and are particularly attractive to bad actors in the early stages of an attack. Moreover, when a hacker identifies an outlier asset and successfully exploits it, he or she then searches for similar assets to that which was just successfully exploited, all in an effort to gather as much data and as many credentials as possible while expending the least effort. Because of this, vulnerabilities on these outstanding assets are particularly critical, as they're much more likely to be exploited in the crucial early stages of an attack.

But, with thousands of assets on the typical enterprise network, it's virtually impossible for the average IT analyst to identify outlier assets accurately. Here, AI - specifically machine learning - can be used to filter outlier assets from the thousands on the network. To accomplish this, an asset model is constructed with multiple variables, essentially a mathematical representation of each asset on the network. Then, machine learning techniques from the field of pattern recognition are employed to separate assets that differ from a contextual baseline. By identifying outlier assets without the need for expert human intervention, this process can be revisited frequently, and automatically, as the network evolves, and the results used to rank vulnerabilities in order of remediation priority.

02 Identification of Business Critical Assets

Not all assets on an enterprise network are created equally. Some servers or machines are more important to business operations than others. They might house particularly sensitive data or power applications essential for the day-to-day operation of the business. Vulnerabilities on these "business-critical" assets are often considered higher priority. However, identifying which assets are more critical than others manually is challenging. The sheer number of assets on a typical network makes this task difficult, at best, while the constantly changing nature of the network and the organization compounds the difficulty.

Over a relatively short period of time, these AI-deduced patterns can help draw data-based conclusions as to which assets are most important to the organization, all without

Thus, identifying outlier assets on a network is a key step in the vulnerability management process, particularly vulnerability prioritization.

Here, AI can learn the patterns that make an asset important to an organization by studying the interaction with the asset and then predicting importance and priority of other assets. (The same prediction mechanism can also be applied to assets that might have been forgotten).

the need for tedious manual intervention. As business-critical assets are established, the criticality of vulnerabilities on those assets can be confidently raised.

Put another way, prioritizing on what matters most to the organization depends on context.

03 Exploit Publication Predication

Two vulnerability management realities are:

1. New software vulnerabilities are discovered and published daily, and
2. Most vulnerabilities are never used by bad actors to attack networks.

Confidently predicting which new vulnerabilities are likely to be exploited and which are unlikely to present a threat can be a relevant factor in determining remediation priorities. Based on the characteristics of the new vulnerability, historical data from exploit databases, and other factors, today's AI researchers have developed techniques to predict the probability that an exploit for a newly-published vulnerability will eventually be published.

AI can interpret posts and blend the meanings of thousands to add context to any given vulnerability's practical risk, a risk that can change quickly as new exploits are created and distributed among the ever-growing community of bad actors.

Predicting whether or not a vulnerability is likely to be exploited can de-prioritize a large chunk of an enterprise's vulnerabilities, but with tens or hundreds of thousands of vulnerabilities on the typical enterprise network, even eliminating a large portion of the total will leave thousands to be addressed. Exploit publication prediction is sometimes considered a vulnerability prioritization panacea, but, although an important element of a comprehensive VM program, insufficient in a vacuum to independently prioritize remediation efforts credibly and confidently.

04 Discerning Vulnerability Exploitation Trends

Brand marketing professionals are using AI-based sentiment analysis to evaluate countless posts on social media platforms that reference their products. Collecting this data and applying AI to it can provide insight into how a brand is perceived in the market, and how that changes – for better or worse – over time. Similarly, cyber security chat boards, social media platforms, dark web sites and other on-line sources of cyber security conversations can be collected and analyzed to predict which vulnerabilities are the most likely to be exploited, which security experts are most concerned about, and how those trend over time.

Collecting and evaluating millions of such posts over time is impractical with human resources, but can be accomplished continuously with Neural Networks and Natural Language Processing (NLP) techniques like modelling and document embedding, AI technologies that can discern meaning, positivity/negativity, and even more importantly, can extract precise technical information from text.

05 Assessing the Reliability of Detections

An element of vulnerability management that is often unappreciated by those outside the field is the challenge of vulnerability detection with confidence, or detecting vulnerabilities correctly. Indeed, it is generally accepted by VM professionals that it is relatively easy to find vulnerabilities, maybe a bit too easy. AI can be employed in

this part of the vulnerability management process to help reduce the number of false positives, essentially "detecting the mis-detections." Factors such as services running on the asset and the detection mechanism that flagged the vulnerability can be used to assess the probability that the identified vulnerability is, in fact, a legitimate one. And, as the collected data analyzed by the AI engine increases over time, its ability to accurately predict false positives versus legitimate vulnerabilities will improve.

To improve the reliability of vulnerability detection, Bayesian networks can be used when there is uncertainty in some result; in this case, whether an identified vulnerability is legitimate. The technique allows other observations to be included as evidence in the assessment, for example, how frequently does the detection mechanism being used generate false positives, or whether this vulnerability has often been identified manually as a false positive in the past. Effectively, employing Bayesian networks allows a more intelligent analysis that balances imperfect scanning techniques without expert human knowledge.

Applying AI to actual vulnerability remediation data across multiple organizations can yield insights based on the collective judgement of many hundreds or thousands of IT and security peers, and as discussed previously, the larger that peer group grows, the higher the probability the decisions are sound.

06 Leveraging Industry Vulnerability Remediation Priority

All modern vulnerability management products today are either cloud-based or have a cloud-based component. Although there are myriad benefits to a cloud-based vulnerability management platform, one of the most valuable, yet typically under-appreciated, is the user data that can be anonymized and culled from the application. Every organization is often remediating vulnerabilities on multiple assets daily. Multiply several daily remediation activities across dozens, hundreds, or thousands of customers, and a cloudbased vulnerability management product has a rich data source on which to apply an AI engine.

Using this ever-changing and growing data source can reinforce or contradict conventional vulnerability remediation prioritization.

Which assets are enterprises patching the most frequently?

Which vulnerabilities appear to be the most concerning to peer organizations?

Which are lower priorities?

We all learned in high school that copying one classmate's answer on a test question is not only unethical, but a risky proposition given there's no assurance that you picked the right classmate to copy. However, if you could determine that 90% of the class chose a specific answer, you'd have significantly more confidence that the answer was the right one.

Using a machine-learning technique known as Gradient Boosted Tree Regression, user behaviors and preferences can be blended with their history of remediation to predict what is important. Using this ever-expanding database of cloud-based users and their

remediation activity, the contribution to the vulnerability risk score becomes a dynamic element that reflects the constantly changing nature of the threat landscape.

07 Developing Remediation Plan Recommendations

Once a context-driven priority list of vulnerabilities is established using some of the AI methodologies detailed here, optimizing remediation work is the final step in the vulnerability management process. Here, AI has a role to play as well.

Most medium to large enterprises can identify more vulnerabilities on their networks than could practically be remediated in any reasonable timeframe, so developing remediation plans that maximize risk reduction while minimizing remediation activity is essential to any modern vulnerability management program.

AI can be leveraged to address this challenge as well. Specifically, a Risk-Aware Recommender System can be used to generate multiple remediation scenarios. Similar to the algorithm used to make Amazon recommendations to consumers - only its objective is to reduce risk and optimize time, not maximize sales - a vulnerability management Recommender System would also take into account the risk reduction afforded by each remediation scenario using individual vulnerability risk scores generated using some of the AI techniques discussed in this paper.

The AI Reality

Whether the fault of Hollywood movie makers or overzealous vendor Marketing teams, Artificial Intelligence technology has been simultaneously over-hyped and misunderstood. Despite that, AI researchers and data scientists have found practical ways to exploit the technology for the benefit of businesses and consumers alike. Cyber security, and in this case, vulnerability management in particular, is another discipline where AI can, and has, substantially changed the way previously labor-intensive activities are completed.

Secureworks®

Secureworks® (NASDAQ: SCWX) a global cybersecurity leader, enables our customers and partners to outpace and outmaneuver adversaries with more precision, so they can rapidly adapt and respond to market forces to meet their business needs.

With a unique combination of cloud-native, SaaS security platform and intelligence-driven security solutions, informed by 20+ years of threat intelligence and research, no other security platform is grounded and informed with this much real-world experience.
www.secureworks.com

Corporate Headquarters

United States
1 Concourse Pkwy NE #500
Atlanta, GA 30328
+1 877 838 7947
www.secureworks.com

Europe & Middle East

France
8 avenue du Stade de France 93218
Saint Denis Cedex
+33 1 80 60 20 00

Germany
Main Airport Center,
Unterschweinstiege 10 60549
Frankfurt am Main Germany
069/9792-0

United Kingdom
One Creechurch Place,
1 Creechurch Ln
London EC3A 5AY
United Kingdom
+44(0)207 892 1000

1 Tanfield
Edinburgh EH3 5DA
United Kingdom
+44(0)131 260 3040

United Arab Emirates
Building 15, Dubai Internet City Dubai,
UAE PO Box 500111 00971 4 420
7000

Asia Pacific

Australia
Building 3, 14 Aquatic Drive Frenchs
Forest, Sydney NSW Australia 2086
1800 737 817

Japan
Solid Square East Tower 20F
580 Horikawa-cho, Saiwai-ku
Kawasaki, 212-8589
Japan
81-(44)556-4300
www.secureworks.jp