451 Research® | PATHFINDER REPORT

Expanding Machine Learning Applications on the Endpoint

COMMISSIONED BY

CYLANCE

SEPTEMBER 2018

About this paper

A Pathfinder paper navigates decision-makers through the issues surrounding a specific technology or business case, explores the business value of adoption, and recommends the range of considerations and concrete next steps in the decision-making process.

ABOUT THE AUTHOR

FERNANDO MONTENEGRO

SENIOR ANALYST, INFORMATION SECURITY

Fernando is a Senior Analyst on the Information Security team, based in Toronto. He has broad experience in security architecture, particularly network security for enterprise environments. He currently focuses on covering vendors and industry events in the endpoint security and cloud security spaces.

Executive Summary

The constant evolution in technology has given rise to several industry trends that, taken together, point to increased importance for endpoint security considerations. These include dissolution of traditional perimeters, increased adoption of SaaS-based applications, and stronger network encryption technologies. Considering these and other changes, the endpoint becomes a key element of security strategy.

Endpoint security has so far been focused primarily on prevention/protection technologies. While a prevention/protection-first approach is highly recommended and can help alleviate security teams from dealing with preventable incidents, ignoring capabilities around detection and response to security incidents can have a detrimental impact on security. By contrast, there is significant potential upside in deploying technology that can integrate detection and response with existing protection/prevention capabilities.

Endpoint detection and response (EDR) has emerged as a key component of endpoint security strategy. Initial EDR functionality was targeted at larger organizations, but there is broader interest in deploying this tooling. Compared to existing prevention and protection tools, customers believe there is still opportunity for improvements in EDR.

One possible avenue for improving EDR is the broader adoption of machine learning techniques. While the term is often overused in security, machine learning methods have been effectively used in numerous areas, including spam detection, data loss prevention and malware detection within endpoint security. While there is a need for deep domain expertise and data science theory and methods, machine learning methods can be applied to different aspects of EDR.

As organizations look to deploy or improve their EDR practices, they should consider the possible applications of machine learning approaches. Machine learning for EDR can address different aspects of detection, investigation and response, and can nicely complement or augment efforts by existing security teams.

Introduction

In general terms, many security concerns boil down to an economics problem: how do we best allocate finite resources? While the list of new vulnerabilities is seemingly endless, and the adversaries don't seem to slow down on their innovations, the reality is that organizations have to make trade-offs in terms of security spending versus the benefits – usually reduced exposure – derived from it. To think otherwise is to ignore the reality of how organizations work, which can certainly push security practitioners even further away from the proverbial 'seat at the table' that is so desired as a way of supporting security initiatives.

Endpoint security has been a constant subsector of the security industry since the popularization of personal computers in the early 1990s, when it established a foothold. Since then, the industry evolved quickly through products such as antivirus, anti-malware, next-gen antivirus and, recently, endpoint protection. The market rushed to fill a need that organizations had to protect against increasingly sophisticated threats. Although initially seen as a nuisance, the onslaught of attacks became a critical concern for most organizations.

Endpoint protection is, thus, a key strategic requirement, one that organizations have been addressing with a variety of approaches. These include asking – and receiving – better security functionality from their default operating system vendors, as well as deploying increasingly sophisticated third-party products that leverage new techniques for protecting endpoints. Still, organizations should consider more.

Preventing and protecting against endpoint attacks remains a critical objective. Any strategy that ignores a prevention/protection-first approach to endpoint security risks overloading security teams with numerous incidents that could have been prevented. Resources are not infinite, and overall security posture suffers when teams are overwhelmed with preventable incidents; nevertheless, a prudent approach for most organizations is to think beyond reactive and legacy measures, anticipating the need to address incidents that bypass modern defenses. Organizations should consider how endpoints can be part of a broader incident response and digital forensics practice, working alongside the rest of the security architecture. To that end, detection and response capabilities have had to improve.

The industry has broadly adopted the term endpoint detection and response to refer to the capabilities that, when deployed on endpoints, allow for fine-grained detection of evidence of security incidents, investigation of said incidents and, should it be necessary, some form of response. Initial adoption of EDR has come from larger organizations that, besides having lower risk appetites or more stringent compliance mandates, have been able to bear the higher costs in terms of skill set and operations that required more manual efforts. There is now stronger demand for EDR capabilities by a broader set of organizations, not all of which can maintain the necessary in-house expertise or use external managed detection and response services.

Vendors have been quick to respond to this demand, proposing offerings that include a variety of features and innovations. One of the more interesting areas for innovating EDR is the increased application of machine learning.

Understanding Current Needs and Trends

Prior to discussing technology options, it is important to understand some of the pressures on security teams. The increased importance of cybersecurity has led to the widely held view in the industry that there is a shortage of security skills. Organizations seem to struggle with having enough staff with the right level of expertise to properly address infrastructure security concerns.

The cybersecurity skills shortage manifests itself in multiple ways. There's increased difficulty in hiring professionals with the necessary skill sets to address security needs, just as there's difficulty in updating the skill sets of existing staff. Furthermore, once employees attain a certain level of expertise, it becomes more difficult for organizations to retain them. Data from 451 Research's Voice of the Enterprise confirms this. Figure 1 shows that there is some difficulty in hiring across organizational boundaries.

Figure 1: Difficulty in hiring

Source: 451 Research's Voice of the Enterprise: Information Security, Organizational Dynamics 2017
Q. Does your organization currently face a skills shortage in information security?

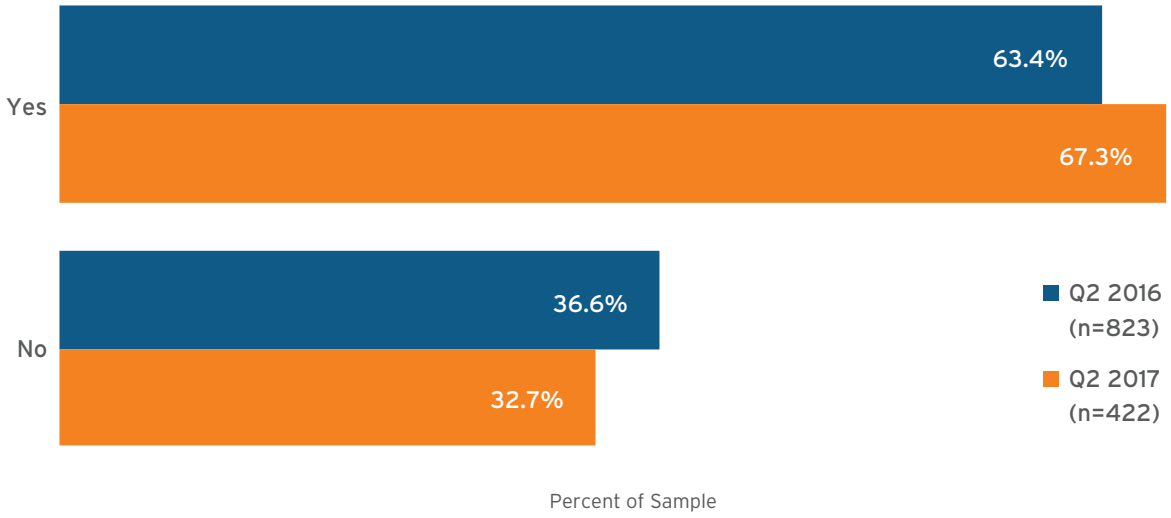

A key step when considering where machine learning can contribute to endpoint detection and response is to understand where customer demands are. A recent 451 Research Voice of the Enterprise survey asked respondents about their level of satisfaction with their endpoint security tooling as it relates to various use cases. Results are below:

Figure 2: Endpoint sentiment

Source: 451 Research's Voice of the Enterprise: Information Security, Workloads and Key Projects 2018

Q. On a scale of 1 to 5, where 1 is 'very ineffective' and 5 is 'very effective', how would you rate your current endpoint security solution against the following use cases?

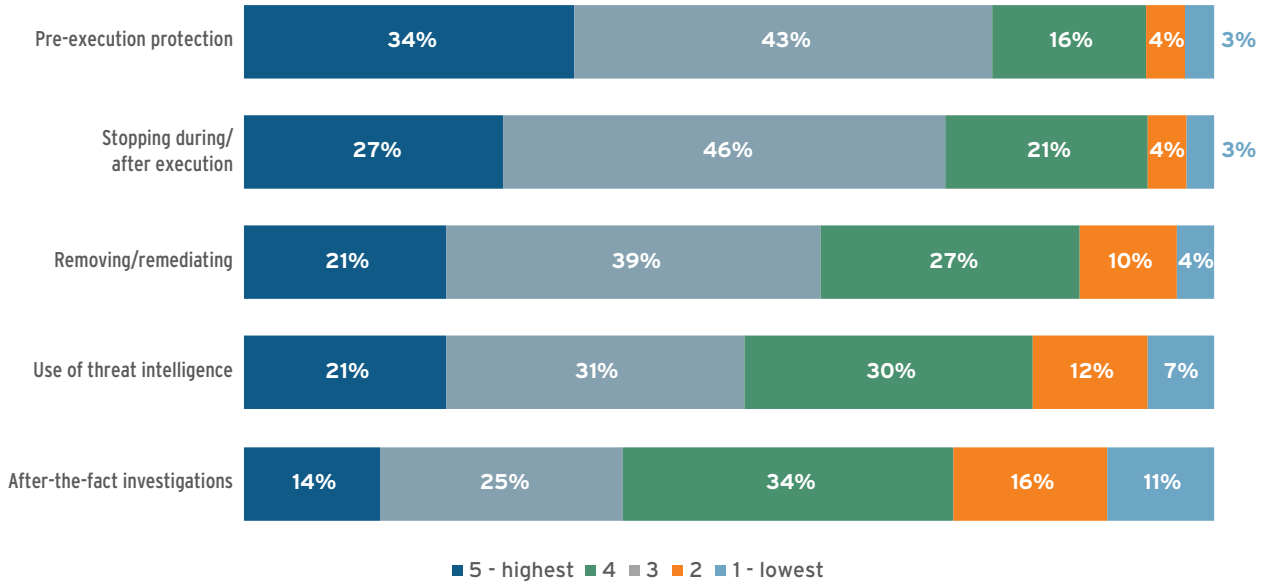

Figure 2 shows that the aggregate response from participants – spanning small and large organizations across several industries – varies significantly by use case. The level of satisfaction with pre-execution protection is notably higher than for additional phases of the incident lifecycle. Notably, there is much more dissatisfaction with support for investigations.

Figure 3: Response to ransomware from respondents with no exposure to ransomware

Source: 451 Research's Voice of the Enterprise: Information Security, Workloads and Key Projects 2018
Q14. If your organization became the victim of a ransomware attack, how would it most likely respond first?

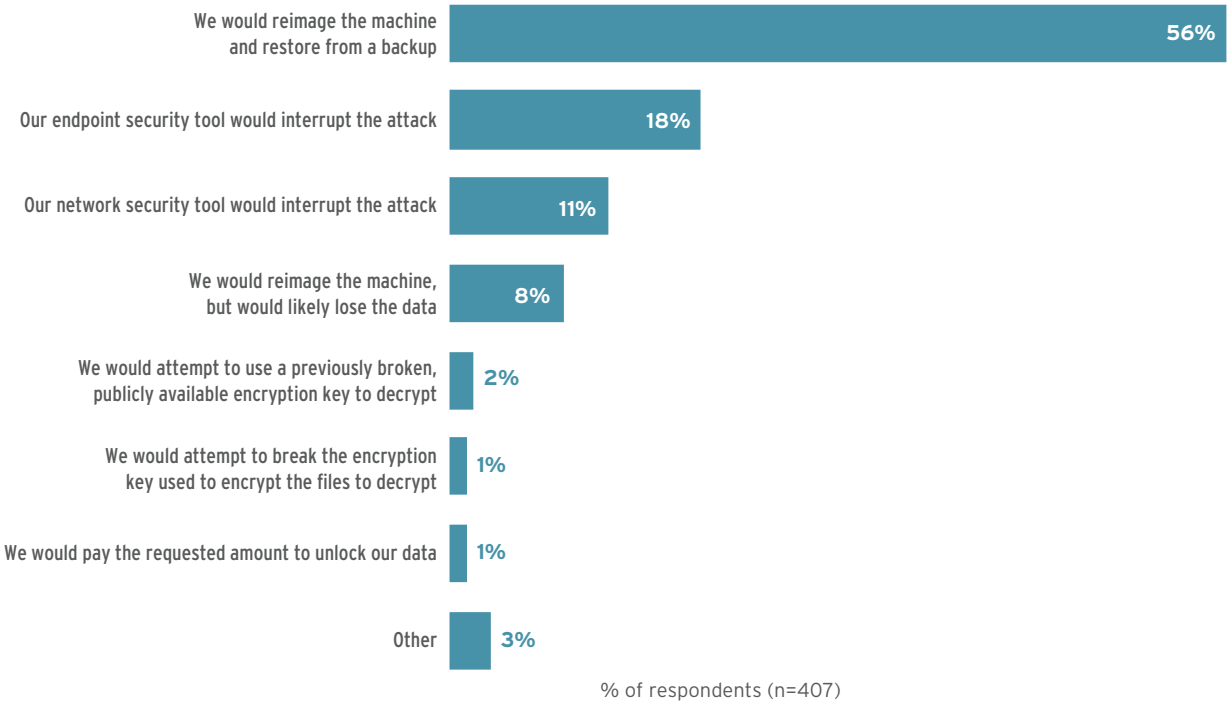

Respondents' expectations regarding ransomware incidents were also interesting. About half of respondents that have not had exposure to a ransomware incident indicated they would need to reimage and restore from a backup, and just shy of one in five indicated they would expect their endpoint solution to interrupt the attack. These expectations, however, don't necessarily align with the responses from those that have experienced ransomware attacks.

Figure 4: Response to ransomware from respondents with exposure to ransomware

Source: 451 Research's Voice of the Enterprise: Information Security, Workloads and Key Projects 2018
Q13. How did you handle ransomware in the case, or the most prominent case if multiple, you experienced?

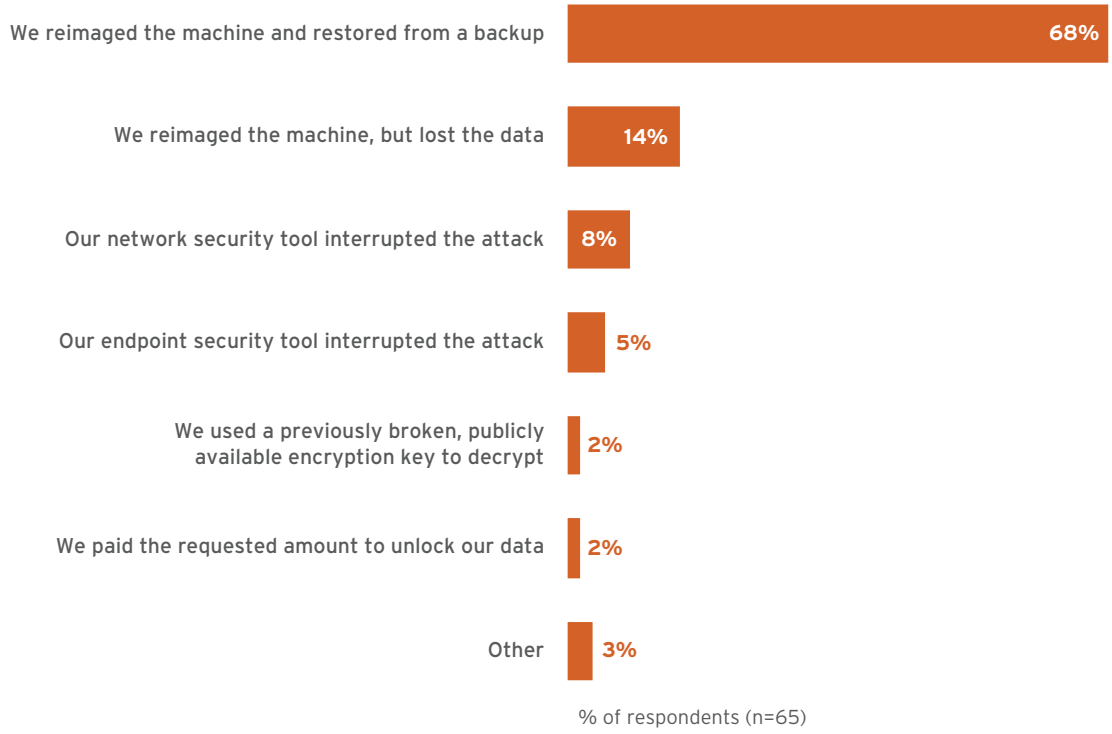

In this case, the percentage of respondents who said they reimaged the machine and restored from backup was significantly higher. Noticeably, however, the share of respondents who indicated that the endpoint tooling was able to stop the attack dropped to one in 20.

Taken together, these data points indicate there is a clear opportunity for improving incident response. Since successful attacks are not just an atomic event but a sequence of steps, the ability to block the completion of the sequence can result in positive impacts to the organization.

Current Approaches to EDR

At a high level, the endpoint detection and response toolset covers capabilities across the following set of tasks:

- Data/telemetry collection. This means collecting relevant information from the endpoints. This information typically includes details about files, filesystems, processes, memory, network connections, user activity and system configuration. Depending on the product architecture, this collection can be done locally on the endpoint and periodically sent to a centralized location, or it can be done centrally. There are arguments for and against both approaches to data collection.
- Exploratory data analysis. As the system collects information from the endpoint fleet, an analyst should be able to interact with the data that has been gathered to answer questions related to the various use cases for EDR. This interaction between agent and the system will vary, but can include elements of ad hoc queries, data visualization, prerecorded sets of queries and scheduled reporting, etc.
- Analytics and enrichment. In addition to interaction with the analyst, the EDR tooling is expected to perform some level of independent analysis with the goal of triggering alerting based on security use cases. This analysis often includes the application of statistical methods including various machine learning techniques to assist with anomaly detection, clustering or predictive modeling. This analysis can be enhanced further by integrating the EDR system with third-party data-enrichment sources such as threat intelligence feeds and malware sandboxes for analyzing runtime behavior.
- Response capabilities. Whether triggered by an automatic alert from its analytics engine, an external alert by another system, or invoked manually by an analyst/operator, the EDR system should offer a range of responses that support investigation or containment use cases. These responses normally range from restricting endpoints – either via the EDR agent or in collaboration with the rest of the environment, and may be done at the process, network, file, system or user levels, with optional rollback to previous states – to triggering additional data collection – often for further analysis or, in some cases, interaction with law enforcement. Importantly, some systems support the use of playbooks (predetermined response plans) to be used, either manually or in some automated manner.

Different products will have different strengths and weaknesses. Also, particularly when it comes to analytics (detecting behavior that is considered malicious), the distinction between endpoint prevention or protection and endpoint detection and response may be less well defined: is the automatic blocking of malicious activity considered a prevention/protection or detection/response?

The more common form of analytics in EDR currently consists of rules-based engines tied to configuration options. In this mode of operation, dozens to hundreds of rules are created, covering the wide range of options for potentially malicious behavior – should a script execute from a temporary directory? Should a word processor process spawn code that modifies memory of another process?

This approach is valid but ends up requiring more effort to keep current in the face of newer attacks, changes to attack patterns, or changes to how endpoints are used. As organizations look to automate their operations, the increased use of scripting, for example, may have significant impact on more brittle detection rules.

Machine Learning

Few technologies have had as great an impact on security recently as machine learning techniques. Faced with increasing volumes of data, increasingly sophisticated attackers that are quick to respond to updates in defenses, and a shortage of resources, organizations are looking to machine learning as a key component in updating their defenses. Before applying machine learning, though, organizations should clearly understand where machine learning can fit within their security architecture, and that requires that they understand key points about how it works.

In truth, machine learning methods have been applied to security for several years. Protection against unwanted email ('spam') has benefited greatly from techniques such as Naïve Bayesian processing. Fraud analytics approaches have long used anomaly detection via similar Bayesian methods. Other techniques such as natural language processing assist in areas such as e-discovery and data leakage prevention. There has been tremendous evolution in machine learning, driven by increasingly more capable resources such as GPUs and cloud-based resources.

Hyperbole aside, applying machine learning to a specific problem requires a clear two-phase process: first, there is significant effort in creating the prediction model to be used, then there is ongoing application of that model via ongoing predictions. Hidden within these two phases are significant amounts of effort, research and nuance.

Model creation involves having a significant understanding of the problem being addressed. This is often referred to as 'domain knowledge.' Within security specifically, this means understanding attack patterns, operating system mechanisms, vulnerability management, application security considerations and numerous other topics.

Creating successful models also requires deep expertise in data science itself, both in terms of data engineering – collecting massive amounts of information and massaging it into useful formats – and the mathematical and statistical foundation theories and techniques of machine learning – variations on supervised and unsupervised learning, among others, and the different trade-offs that must be made between multiple models. This often requires an experienced data science team working in conjunction with domain experts.

These elements are brought together in a process that involves analyzing the data collected for characteristics (or 'features') that can be used to 'learn' what a particular object represents, then iteratively analyze these features to derive mathematical models that can later be used for prediction. This analysis process is usually time- and resource-intensive and is done before the models are used.

When it comes time to use the machine learning models in real life, the endpoints must collect the new data, extract the necessary features, and process them through the model that was created. There are significant variations in this step as well because that processing can be done locally within the endpoint or remotely within cloud/centralized infrastructure. As with other topics involving machine learning, there are fundamental trade-offs.

On one hand, local determination and response can be particularly effective against fast-moving attacks. Quick identification and interdiction of malicious activities can prevent further damage to systems and greatly simplify cleanup efforts. The potential trade-off, in this case, is that the models that make predictions must do so with fewer resources.

Cloud-based processing, on the other hand, may bring to bear more sophisticated analytics and easily incorporate additional datasets or updated information, but at the expense of response time. This response delay could, in some cases, result in a larger impact on the target system. Understanding these trade-offs and choosing the right balance requires the deep domain expertise mentioned above to understand how different threats impact endpoints.

Typical Applications in Endpoint Security

There are several applications of machine learning for security in general, and for endpoint security specifically. In endpoint security, the more prominent application has been the use of machine learning methods as a replacement for traditional antivirus signatures: because malware writers create techniques that mutate the payloads, applying machine learning techniques to derive common features from the payloads and detect and block malicious binary payloads was a significant enhancement.

Other applications so far have included aiding security researchers by performing automatic clustering of elements (files being analyzed, function call patterns, etc.) and performing anomaly detection via a variety of statistical methods. Machine learning techniques can also be applied in supporting areas as well. The use of natural language processing has aided with providing more human-friendly interfaces in security products, which can lead to higher productivity with lower costs.

Not All Machine Learning is Created Equal

While machine learning offers significant benefits, there are also limitations and drawbacks that organizations should consider. Much research is being conducted across several aspects of machine learning to mitigate those drawbacks, meaning that what is 'state of the art' continues to evolve, but even so, organizations should be aware of some limitations.

First and foremost, proper design and application of machine learning requires significant domain expertise and access to meaningful datasets. Understanding the nuances of a specific problem – whether malware detection, natural language processing or other topics – requires a broad swath of expertise covering data engineering, mathematical methods and machine learning techniques.

Hidden within the application of machine learning methods to perform analysis is an inherent trade-off between 'false positive' and 'false negative' results. Organizations need to strike a balance when tuning their systems: they don't want to overwhelm their teams with false alerts, nor do they want to fail to send legitimate alerts.

Traditional machine learning methods assume a certain stability in the environment: if the features that were chosen for the model change, or their relative importance to arriving at the final classification change, then that model may no longer be useful. This is particularly important when dealing with cybersecurity use cases because this knowledge of machine learning limitations is also available to the adversaries. Significant industry effort is being put toward studying 'adversarial machine learning' to develop methods that are resistant to attacks.

Applying Machine Learning to EDR

One of the expected benefits of applying machine learning to EDR is the ability to better analyze complex system behavior to identify and potentially interdict newer attacks.

Machine Learning has been adopted with some degree of success in areas such as endpoint prevention/protection, as well as network anomaly detection, anti-fraud and email security. The application of machine learning to EDR can touch both functional and non-functional aspects. One of the expected benefits of applying machine learning to EDR is the ability to better analyze complex system behavior to identify and potentially interdict newer attacks. In contrast to traditional approaches based on static rules, machine-learning-based EDR may be able to detect attacks sooner, potentially interrupting more complex action chains that would result in further compromise.

The key functional aspect is the application of machine learning methods to the analysis phase of EDR, while non-functional use can cover human interaction aspects. The analysis phase of EDR consists of sorting out malicious patterns across a wide dimension of elements: process information, network activity, command execution, and file and directory activity. This analysis – done with somewhat more rigid rules – can benefit from machine learning because models can be created that capture the scope of malicious activity.

This means that EDR tooling with machine learning may be able to leverage machine learning detection on its own, but also enhance rules-based policies with machine-learning-derived insights. In this scenario, machine learning acts as a higher-fidelity signal that an analyst can code into a specific policy.

Other areas of interest include aiding incident response analysts by performing more sophisticated anomaly-detection analysis: if a machine-learning-enabled investigative tool can automatically outline clusters of behavior, potentially also applying data enrichment, it will allow analysts to perform analysis much faster.

Analyst interaction can also benefit from natural language processing: if the system can parse user inputs in a more natural manner, it will reduce the learning curve during product adoption, as well as potentially reduce errors during an investigation.

Regardless of application – functional or non-functional – one key point to consider is that proper use of machine learning methods generally requires having access to significant amounts of suitable training data. Specifically, one point of caution is that compared to prevention/protection use cases, EDR use cases are quite different and may require significantly more effort in both collecting data and the resulting models.

Conclusions and Recommendations

In our opinion, organizations are rightfully placing increased importance on the selection of their endpoint security tooling. This reflects that endpoints are truly becoming a critical component in modern architecture as workforces become more mobile, and application-level security places increased importance on the endpoint as a control point.

Organizations can benefit from considering prevention/protection components for their endpoint strategy, and survey results indicate there is a higher level of satisfaction with this tooling than with EDR. This translates into an opportunity to consider improvements to EDR. Better application of machine learning techniques to EDR use cases represents one such avenue for improvement. Machine learning has been widely deployed in the security industry, with good results across different use cases.

From both functional and non-functional aspects, EDR presents opportunities for applying machine learning. Machine-learning-based EDR may be able to achieve better outcomes than traditional tooling during the analysis phase by quickly, efficiently and oftentimes autonomously mapping out and stopping malicious behavior. There is also the potential for increased analyst performance by using investigative tools that are assisted by machine learning methods.

Deploying machine-learning-enabled EDR during analysis, investigations and responses can provide significant value to organizations. The potential benefits include not only severely limiting or even entirely preventing newer advanced attacks, but also improving the productivity and efficiency of existing security resources.

It is well understood that endpoint security is growing in importance and must be maintained constantly. Organizations that can improve security outcomes across their endpoint detection and response practices in an efficient manner will be able to support the agility that their business and the modern threat environment demand.

About 451 Research

451 Research is a preeminent information technology research and advisory company. With a core focus on technology innovation and market disruption, we provide essential insight for leaders of the digital economy. More than 100 analysts and consultants deliver that insight via syndicated research, advisory services and live events to over 1,000 client organizations in North America, Europe and around the world. Founded in 2000 and headquartered in New York, 451 Research is a division of The 451 Group.

NEW YORK
1411 Broadway
New York, NY 10018
+1 212 505 3030

SAN FRANCISCO
140 Geary Street
San Francisco, CA 94108
+1 415 989 1555

LONDON
Paxton House
30, Artillery Lane
London, E1 7LS, UK
+44 (0) 207 426 1050

BOSTON
75-101 Federal Street
Boston, MA 02110
+1 617 598 7200

451 Research®