10 Machine Learning Secrets to Stop Modern Attacks

How to Block Endpoint Attacks and Detect MITRE ATT&CK® Techniques with Machine Learning

Machine Learning Is Everywhere

Today, we are surrounded by artificial intelligence or AI. It's built into the devices we use every day, from smartphones to computers, cars, and even robotic vacuums. It powers the customer service chatbots that answer our questions, the navigation systems that tell us the fastest route to our destination, and the movie recommendations displayed on our favorite streaming sites. Just like chocolate, it seems to make everything a little better; it speeds up manual processes, reduces human error, and makes the seemingly impossible possible.

Machine Learning and Cybersecurity

Nowhere is the impact of machine learning more evident than in cybersecurity. Fraud detection, web and email filtering, and network security, among others, depend on machine learning. The most impressive advancements today are occurring in endpoint security and detection and response. Signature-based antivirus is rapidly evolving into AI-driven next-generation antivirus. Static, rule-based threat detection is giving way to more advanced analytics and machine learning.

Machine learning, despite its promise, is only as good as its application. Applied ineffectively, it will deliver the same results as manually defined rules and signatures: excessive false positives and the continued inability to detect critical attacks.

Machine learning requires AI expertise, continual investment, and, above all, reams of good data from which to train machine learning models. This paper will reveal the 10 machine learning secrets you need to know to stop cyberattacks, including five foundational requirements and five ways machine learning can be applied to detect and prevent threats.

Foundational Requirements for Machine Learning

1. Quantity, Quality, and Diversity of Data

The number one requirement for effective machine learning (ML) is high-quality data. ML requires lots of data from multiple sources correctly labeled and stored over time. Since machine learning algorithms analyze training data to build ML models, it should be no surprise that better data generates better ML models. As a case in point, consider how machine learning can be used to detect and block malicious files. To build an ML model for local analysis, data scientists first define a comprehensive set of file attributes that could indicate whether a file is malicious. They then use machine learning to parse the training data and determine whether a file is malicious or benign.

For supervised machine learning—which uses labeled data to build ML models—data must be labeled correctly. If data is labeled inaccurately—for example, a benign file is identified as malware—then the resulting ML models will be flawed. Therefore, data scientists must rigorously evaluate every labeled data asset to ensure it is labeled accurately. Adding multiple types of labels—such as a grayware or a ransomware category for malware files—enables the model to clearly learn border cases, which will reduce false positives.

Building robust processes to label data with internal feedback loops, such as incorporating policy exceptions, allow lists or incident resolutions from real-world customer deployments, can help improve data labeling correctly. High-quality data is essential for machine learning. Each new malware sample or benign file added to the database increases the accuracy of the ML model because it allows the model to detect more malware signals. As a result, the system gets better and better at identifying malware and legitimate files. The more extensive the database of information, and the more often it is updated with rich data, the better the results. Gathering a varied set of data from multiple sources that represent many different types of attacks strengthens machine learning functions.

2. Expert Feature Engineering, Including Normalization and Stitching of Data

Besides the amount and diversity of data, ML benefits from comprehensive data with deep, contextual information. Details such as application, user, content, and device information can improve the precision of machine learning and analytics used by security tools. However, if each data source labels this contextual information differently, or if the ML functions do not understand the information, then they won't analyze it correctly.

To build accurate ML models, data must first be parsed and normalized into a common message schema. For example, different log sources might identify a single user as "john.doe@company.com," "jdoe," "host123\jdoe," "company\jdoe," "8c6bfd4a-4cb2-11ea-b67e-88e9fe502c1f," and other labels. A security system that applies ML and analytics to data must first normalize the data into a common format, such as account "host123\jdoe" and domain "company.com" before processing it.

If security systems collect log data from multiple sources, they should automatically stitch related log messages together into a single, augmented log record or "story." By stitching together data from multiple sources, ML engines can apply cross-data analytics to these stories for more accurate threat detection. Data stitching provides more context for analytics and deduplicates event data if the same event was observed by multiple sensors.

3. Optimal, Tailored ML Algorithms

The third secret to stopping cyberattacks is to select the best ML algorithm based on the type and amount of data collected and the overall objective of the ML analysis. There are numerous approaches to machine learning, including supervised, unsupervised, semi-supervised, active, reinforcement, and deep learning. Within these approaches, data scientists can select from various ML algorithms, from tree-based to linear regression and neural networks to clustering.

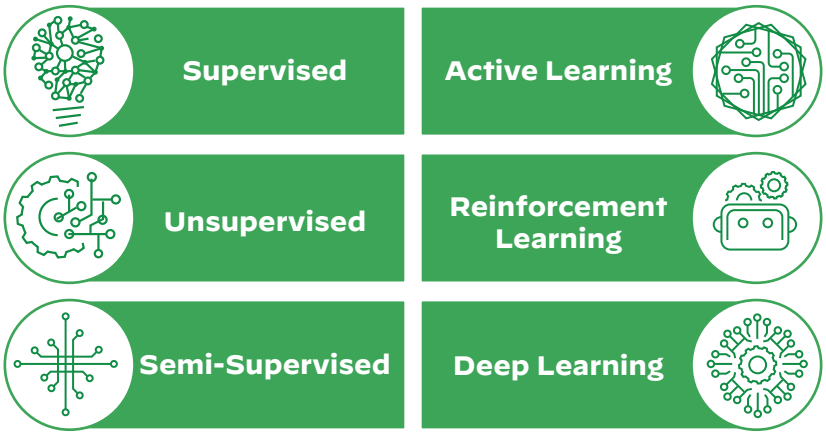

Figure 1: Common approaches to machine learning

Supervised and semi-supervised learning are the most important approaches in cybersecurity because they can deliver definitive and actionable results. Unsupervised learning, which examines unlabeled data, can help data scientists identify patterns and anomalies in data. However, unsupervised learning often lacks the context and precision required to block attacks.

Each ML algorithm has its strengths and weaknesses in terms of accuracy and performance. To detect every step of an attack, a security solution might implement multiple ML algorithms. While no algorithm is necessarily perfect, the key to maximizing security outcomes is to carefully evaluate and select the best ML algorithm for each task.

4. Fast Iteration of ML Models

The security community often disparages signature-based antivirus for not keeping up with rapidly evolving threats. But the dirty, dark secret of the cybersecurity industry is that some "next-gen," AI-powered antivirus solutions lag behind so-called "legacy" antivirus products because they only update their machine learning models once a quarter or once every six months. While a machine learning approach is, in general, more impervious to new malware variants, some novel strains can require updated ML models.

Cyberattacks evolve quickly, and ML models must evolve with them also. Old ML models lead to poor detection of new samples and techniques. Look for security solutions that update ML models frequently and that complement AI-powered local analysis with other layers of protection, such as threat intelligence, behavioral threat protection, analytics, and incident and user scoring.

5. Cloud Computing

Security vendors sometimes herald the cloud, like machine learning itself, as the solution for all the industry's ills. The reality is more nuanced, yet the cloud undoubtedly provides the scale and the speed needed for modern machine learning. Security solutions can leverage cloud native deployment models to collect the massive amounts of data needed for machine learning—all the malicious and benign files, events, and other essential data that power machine learning. Cloud computing can leverage cloud scale to run ML frameworks and quickly build new models, helping to ensure ML models are up to date. Lastly, cloud native security services can unlock security innovations such as community-based analytics.

Machine Learning Applications to Stop Threats

1. AI-Powered Local Analysis to Block Known and Unknown Malware

Antivirus software that relies solely on static signatures fights a losing battle against the tidal wave of new malware introduced daily. Because attackers can make multiple small modifications to malware to evade signature detection, antivirus makers must create an exponentially increasing number of signatures to block attacks. The most promising weapon in the endpoint security arsenal is machine learning, with its ability to quickly learn, make instant decisions, and proactively block threats. One of the main advantages of machine learning is that it can handle minor deviations in a way that signature-based approaches cannot.

Adding AI-powered local analysis to endpoint security allows decisions to be made instantly based on what has been learned in the past while continuing to learn and improve as new malware variants are seen. AI-powered local analysis can identify real threats more accurately than legacy approaches without bogging down endpoint performance with continual signature scans. Local analysis can quickly decide whether a file is malware and stop it before it can execute and cause damage.

Manual, signature-based detection engines are only as good as their own analysts. An AI-powered local analysis engine built using diverse data, high-quality threat intelligence, and product feedback loops can leverage the power of analysts in the world. An ideal local analysis engine should incorporate all the foundational requirements for machine learning. However, a diverse and correctly labeled dataset and a frequently updated ML model are essential for local analysis.

2. Behavioral Analytics to Detect Stealthy Threats

Advanced adversaries and malicious insiders can often operate under the radar. Instead of relying on traditional malware, they can take advantage of seemingly benign applications and services to further their attacks and achieve their objectives. To pinpoint hard-to-detect attack activity, such as a login by a dormant account, multiple discovery commands run on a cloud compute instance, or "the impossible traveler." Without generating false positives, a security solution detects the changes in behavior that attackers cannot hide. This starts by profiling activity with behavioral analytics. By building a baseline of expected behavior for users, devices, and endpoint processes, a security solution can identify the changes in behavior indicative of an attack. Behavioral analytics can reveal malware activity, command-and-control beaconing, internal network discovery, credential theft, lateral movement, and data exfiltration by detecting unusual activity compared to past behavior or peer behavior.

The key building blocks for behavioral analytics are comprehensive data and the automated stitching of data. Because attacks can come from any direction, organizations should integrate and inspect endpoint, network, cloud, and identity data for potential threats. The stitching of malicious data and benign data enables analytics to extract relevant artifacts from across data sources for more precise analysis.

By collecting all events—not just data after an alert is triggered—and storing data for a minimum of four weeks, behavioral analytics engines can profile behavior and zero in on suspicious activity indicative of an attack.

3. Cloud-Based Malware Analysis

To complement local malware analysis performed on the endpoint, security solutions can send unknown files to cloud-based malware analysis services for deeper analysis to detect unknown malware. Cloud-based services offer high-fidelity detection and evasion-resistant discovery. They can perform AI-powered file analysis in the cloud and detect known threats by analyzing the characteristics of samples prior to execution.

Profiling user behavior requires more than just analyzing one, two, or seven days of data. For accurate baselines, be sure to collect at least four weeks of data.

Much like local analysis, cloud-based file analysis uses ML models to unearth malware. However, the sheer compute power of the cloud offers much deeper and more comprehensive analysis than local analysis due to compute and storage limitations inherent with endpoints. Cloud-based malware prevention services can also detonate unknown files to examine real-world effects and behavior.

4. ML-Based Analysis to Identify High-Risk Users and Entities

The ninth machine learning secret is to identify suspicious users and entities to speed incident response. Security solutions can leverage machine learning and analytics to analyze user behavior and assign a user risk score. An ML-based risk score lets analysts prioritize high-risk users and incidents for investigations and track score trends over time.

In addition, machine learning can reveal risky entities, such as bulletproof hosting providers. By analyzing multiple inputs with machine learning, including domain reputation, percentage of domains that are malicious for a related group of IP addresses, and search engine results associating the IP addresses with "bulletproof," a security solution can locate bulletproof hosting providers. Because attackers often use bulletproof hosting providers for command and control and malware distribution, behavioral analytics can use this information as a factor when detecting command-and-control traffic. Machine learning not only detects attacks directly but can be used to enhance the accuracy of other detection algorithms by exposing high-risk users and entities.

5. Community-Driven, Global Analytics

Whether the challenge is fast-moving threats or the risk of accidentally blocking legitimate scripts, the power of the security community can help. By gathering and analyzing events from deployments around the world, security providers can uncover new threat vectors. If the same threat is observed across multiple organizations, community-driven analytics can reveal emerging threats such as new supply chain attacks and zero-day vulnerabilities being exploited in the wild.

They can examine how incidents are resolved, uncover false positives, and adjust out-of-the-box detection and protection rules accordingly. Crowdsourced threat intelligence can bolster the speed and accuracy of threat detection as well as endpoint protection, ensuring that all enforcement points are instantly protected against the latest threats. Community-driven analytics harness huge volumes of data, machine learning, and analytics to fortify the entire community against dangerous threats.

Machine Learning Capabilities in Cortex XDR

Using these 10 machine learning secrets as a benchmark, it is easy to see how Cortex XDR®, the industry's first extended detection and response platform, leverages machine learning to stop modern cyberattacks.

AI-Powered Local Analysis on the Endpoint

The Cortex XDR agent includes, as one of its multiple layers of endpoint protection, an industry-best, AI-powered local analysis engine that runs on the endpoint without requiring any internet connectivity. Built from a comprehensive, curated data set and a state-of-the-art machine learning framework, the local analysis ML model accurately identifies and blocks malware with virtually no false positives. Leveraging an expansive set of training data from WildFire, the local analysis engine includes a unique agile framework for rapid ML-model updates to stay ahead of attackers' evolving techniques.

MITRE | ATT&CK®

Cortex XDR delivers the best combined protection and detection scores with 100% threat prevention and 100% detection of all 19 detection steps in the MITRE ATT&CK Round 4 test.

The local analysis ML model is built by examining over 20,000 attributes across more than 16 million malicious and benign files to identify known and never-before-seen malware with laser precision. Careful file analysis, through examination with Yara rules, WildFire analysis, and multiple layers of reputational analysis, ensure that files are correctly labeled before running the machine learning framework. The Cortex XDR agent offers a complete prevention stack, including exploit protection, network packet inspection, threat intelligence, behavioral threat protection, an anti-ransomware module, and more, but the centerpiece of its protection capabilities is its adaptive AI-driven local analysis engine that's always learning to counter new attack techniques.

Behavioral Analytics

Cortex XDR unearths stealthy attacks using analytics and machine learning, allowing security teams to cut dwell times and swiftly contain threats. Cortex XDR starts by analyzing comprehensive data from across the organization, providing complete visibility and eliminating blind spots. It stitches together network, endpoint, cloud, and identity data to accurately detect attacks and simplify investigations.

Cortex XDR tracks more than 1,000 dimensions of behavior, including attributes that are nearly impossible to ascertain from traditional threat logs. It then profiles user, device, and process behavior using:

- **Time profiles and peer profiles:** Cortex XDR baselines user and device behavior, performs peer group analysis, and clusters devices into relevant groups of behavior. Based on these profiles, Cortex XDR detects anomalies compared to past behavior and peer behavior to detect malicious activity, such as malware behavior, command and control, lateral movement, and exfiltration.
- **Entity profiles:** Cortex XDR inspects data to classify different types of devices, such as a Windows® computer, an Apple iPhone®, or a vulnerability scanner. Cortex XDR also learns which users are IT administrators or normal users. By classifying entities, Cortex XDR recognizes deviations from expected behavior based on the type of user or device, reducing false positives.

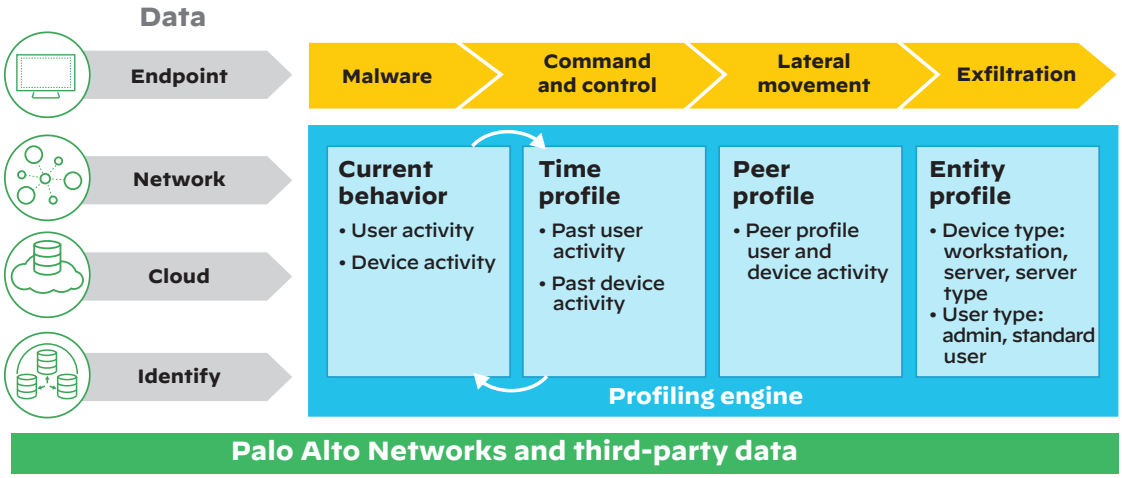

Figure 2: Behavioral analytics architecture for Cortex XDR

View the list of all Cortex XDR behavioral analytics alerts.

Cloud-Based Malware Analysis

Cortex XDR integrates with WildFire® malware prevention service to detect and stop unknown malware and advanced persistent threats. WildFire offers the preeminent solution for cloud-based malware analysis, combining dynamic and static analysis, innovative machine learning techniques, recursive analysis, and a groundbreaking custom-built analysis environment to identify and prevent even the most sophisticated and evasive threats. Its machine learning engine extracts thousands of unique features from each file, training a predictive machine learning model to identify new malware. WildFire also extracts valuable intelligence to provide context for security analysts, generates training updates for ML models, and shares intelligence with Cortex XDR to prevent other attack vectors. WildFire has identified 16 billion unique malware samples since its inception, and it blocks over 260,000 threats every day.

Cortex XDR integrates with WildFire in multiple ways. First, it sends suspicious files to WildFire for an additional layer of analysis and verification as needed. Second, it leverages the massive number of files and malware observed by WildFire to build a robust and diverse training set for its own local analysis engine. Lastly, Cortex XDR applies threat intelligence gathered from WildFire to block known attacks.

ML-Based Analysis of Users and Entities

Cortex XDR identifies high-risk users and entities which leverage machine learning. The Identity Analytics capability in Cortex XDR lets security teams detect risky and malicious user behavior that traditional tools can't see. A dedicated "360-degree user view" shows a given user's risk score, calculated using ML-based analysis of each user's risk. Besides the risk score, the user view displays associated alerts, incidents, and activities. The user view can be toggled when investigating a specific incident, alert, or event.

Cortex XDR also uses machine learning to understand the reputation of entities, such as websites and servers. Patented ML technology, for example, identifies bulletproof hosting providers. By analyzing the risk and reputation of users and entities, Cortex XDR can hone the accuracy of other detection techniques, such as behavioral analytics, and speed investigations by prioritizing the threats that matter most.

Global Analytics

As a cloud-delivered service, Cortex XDR not only provides unmatched scale and agility but also leverages the collective data and intelligence of Cortex XDR deployments around the world to uncover new threats and improve the fidelity of threat detection. Because Cortex XDR automatically parses, normalizes, and stitches together rich data spanning multiple data sources, the Cortex XDR cloud offers an unprecedented foundation for analytics.

Global Analytics reveals emerging threat vectors and novel attack techniques that would be difficult to uncover without a global perspective. Global Analytics can expose new, large-scale threats such as supply chain attacks or zero-day vulnerabilities being exploited.

How does Global Analytics work? It starts with the endpoint. The Cortex XDR agent continuously monitors endpoint behavior, collects granular process data, and sends this data to the cloud-based Cortex XDR service. The Cortex XDR service automatically analyzes this data to generate behavioral profiles of signed processes for each customer. Profiles include which domains and IP addresses a process accessed, which protocols and ports it used, which modules it dynamically loaded, and much more.

If Cortex XDR detects aberrant process behavior for a subset of customers, it will automatically generate an alert. For example, if an accounting software process suddenly starts dialing out to a new IP address using an unusual port, and this behavior is only observed in a small fraction of Cortex XDR tenants that have deployed the accounting software, Cortex XDR would automatically trigger an alert of a behavior with a low global prevalence. This detection method is able to detect supply chain attacks, as well as additional techniques used by attackers, such as DLL side-loading, rootkit-based thread injection, zero-day exploits and more. Global Analytics allows Cortex XDR to detect extremely sophisticated attacks automatically by leveraging cross-customer intelligence and insights.

Cortex XDR can also harness global threat data and anonymized metrics to improve security accuracy. A prime example of the power of community-driven insights is the speed at which Cortex XDR has introduced new Behavioral Threat Protection rules to Cortex XDR agents. The Behavioral Threat Protection (BTP) engine examines the behavior of multiple related processes to uncover attacks as they occur on the endpoint.

By deploying silent BTP rules to Cortex XDR agents, evaluating when they triggered, and if they would have blocked actual attacks, the Cortex XDR research team can release new, active BTP rules to all Cortex XDR agents with unprecedented speed and precision.

Battle-Tested Against Targeted Attacks

Cortex XDR successfully blocked the SolarWinds supply chain attack before the attack was publicly disclosed with its Behavioral Threat Protection capability. Cortex XDR also detected post-exploit activity associated with four critical, zero-day Microsoft Exchange Server vulnerabilities in 2021, before the vulnerabilities were announced, stops Log4Shell attacks, and has blocked countless high-profile ransomware campaigns.

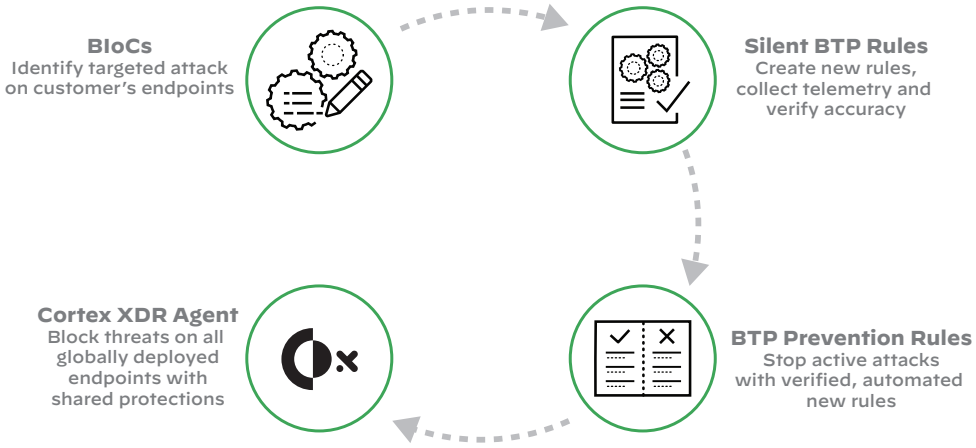

Figure 3: By applying insights from the global community, Cortex XDR can convert detection into prevention rules in less than 24 hours

Global Analytics optimizes detection and protection, which translates into more accurate alerts, less noise, and lower risk to the business.

All data analyzed by Global Analytics is fully anonymized for data privacy and confidentiality; customer-specific identifiers such as names, IP addresses, and more are removed, except for indicators of compromise (IoCs). Global Analytics leverages huge volumes of data, machine learning, and analytics to fortify the entire community against dangerous threats.

Safeguard Your Business with Cortex XDR

The probability of falling victim to an attack is growing far faster than traditional tools or human analysis can handle, and machine learning is clearly the answer. It is a key component for effective endpoint security and detection and response. But not all machine learning solutions are equal. When evaluating a security solution based on machine learning, get answers to key questions to ensure the solution you choose will be truly effective. Once you understand the 10 machine learning secrets to stop modern cyberattacks, it will be clear that Cortex XDR from Palo Alto Networks is the best choice to protect your organization.

paloalto NETWORKS

3000 Tannery Way
Santa Clara, CA 95054

Main: +1.408.753.4000
Sales: +1.866.320.4788
Support: +1.866.898.9087

www.paloaltonetworks.com