



A new threat detection model that closes the cybersecurity gap

Blind to the cybersecurity gap

The cybersecurity gap exists between the time an attacker successfully evades prevention security systems at the perimeter and the clean-up phase when an organization discovers that key assets have been stolen or destroyed.

Inside this gap, attackers have a huge advantage over traditional prevention-based security products. Although prevention tools and techniques are widely used today, cybercriminals routinely outsmart them by using complex and intelligently constructed attack methods.

Cyber attacks are no longer simple smash-and-grab jobs driven by preprogrammed malware. They are controlled by highly skilled, creative and intelligent humans. Ongoing coordination allows a human attacker to progressively learn more about the target network, adapt to any defensive measures, and advance the attack over time.

While attacks have made an evolutionary leap in complexity, security defenses have not. Defenses are overwhelmed trying to find threats using fast pattern-matching signatures of known threats and malware.

As threats became more intelligent and evolved in sophistication over time, traditional security remains dependent upon making snap judgments based on incomplete information.

Traditional security remains dependent upon making snap judgments based on incomplete information.

Today, this imbalance gives attackers a significant advantage. To keep pace, organizations need a more intelligent approach to security – a new class of security that can learn, evolve and think.

This paper lays out the requirements for a new methodology that identifies threats based on what has been learned from the past as well as local context, and then connects events over time to reveal the progression of an attack.

The signature challenge

Security defenses have tried to keep pace by using more and more signatures and delivering them faster and faster. Signatures are the bedrock of traditional security technology and are written to identify exploits, malicious URLs and known malware.

Signatures can quickly identify and block known threats at scale. However, their weakness is that they are inherently reductive – they reduce a known threat to its simplest fingerprint in order to give a single yes or no answer within microseconds to avoid slowing the flow of application traffic.

This reductive focus on immediate and simple answers has created an advantage for attackers who are willing to adapt. Signatures only work by fingerprinting a known threat, and attackers have learned to avoid signatures by using new, unknown threats.

The 2015 Verizon Data Breach Investigation Report illustrates this trend in stark detail, indicating that 70-90% of malware used in data breaches were unique to the organization that was infected.

This means that every organization would need a unique set of signatures to protect itself – a requirement that doesn't scale. But if an attacker uses an unknown, or zero-day, threat, no signature could possibly exist to detect it. The importance and ease of avoiding signatures has not been lost on attackers.

While attackers are able to stay ahead of signatures, it is the persistence of the ongoing attack that has truly turned the tables. Once an organization's outer defenses are compromised, attackers can blend in with the network, progressively spy, and spread deeper until they find high-value assets to steal or destroy.

This process typically involves multiple compromised hosts, a variety of tools and malware, and the theft and misuse of valid user credentials. The important point is that the threat itself is ongoing while attackers evolve their operations and adapt over time.

The threat itself is ongoing while attackers evolve their operations and adapt over time.

The reductive nature of signatures that identify threats at the atomic level is particularly ill-equipped for recognizing the more complex chemistry going on around them. This intelligence gap is precisely why a new security model for threat detection is so vital.

The new threat detection model

The newest, most advanced threat detection model does more than simply plug the gaps in traditional security technologies. It squelches the strategic advantage that attackers have enjoyed for too long.

Coarse-grained detections with a long shelf-life

Detections that use traditional signatures become obsolete when attackers adapt by moving to a new domain or adding a few bits to known malware so signatures no longer match it. This gives them a first-mover advantage where even the most trivial changes keep attackers several steps ahead of defenders.

One of the core goals of the new threat detection model is to deliver detections that remain valid for long periods of time. This requires a shift from fingerprinting every individual instance of a threat to recognizing the fundamental attack characteristics that every threat has in common.

When applied to packet-level traffic, data science and machine learning become extremely powerful tools to identify the fundamental characteristics that distinguish threats from normal traffic.

Focus on attacker actions and behaviors

Traditional detection models attempt to find snippets of exploit code, a known sample of malware or a malicious domain. This leads to an intractable job of constantly finding and fingerprinting an infinite number of malicious occurrences. The task is never-ending and attackers always remain steps ahead by using a new exploit.

To break this cycle, the new threat detection model shifts the focus from trying to name all possible bad things to identifying the unique indicators of attack behaviors and actions.

In other words, the goal shifts from identifying what a thing is, to identifying what the thing does. Although attackers can hide their threats by making slight changes to malware or buying a new domain, the actions and objectives of an attack are always the same.

For example, virtually every attack must establish some form of hidden communications in order for the bad actor to coordinate and manage the attack. The attack also needs to spread internally, compromise more internal devices and credentials, and ultimately destroy assets or exfiltrate them from the network.

Virtually every attack must establish some form of hidden communications in order to coordinate and manage the attack.

By focusing on attack behaviors, defenders can fight and win the asymmetric cybersecurity war by shifting the math of security back in their favor. Instead of using thousands of signatures to find every variant of a threat, they can focus on a few dozen key behaviors that attackers must perform in order to succeed.

Recognize threats over time

One of the most recognizable traits of modern network data breaches is that they evolve over time. This low-and-slow approach has become standard operating procedure for sophisticated attacks, and for good reason. Traditional security suffers from short-term memory and a post-breach form of perfect amnesia.

Traditional security suffers from short-term memory and a post-breach form of perfect amnesia.

The new threat detection model recognizes threats in real-time and identifies the signs of attack that evolve over time. One does not preclude the other. For example, small timing anomalies and cadences within a network session can reveal hidden tunnels and remote access tools used by attackers.

Conversely, recognizing when an employee's credentials have been compromised may require learning the user's normal behaviors over a period of days, weeks and months. While the time scale can be very small or very long, both cases require a keen understanding of threats in relation to time.

Recognize attacks, not just techniques

In order to provide value, security must identify real business risks to an organization and not simply deliver a list of alerts. This requires security solutions to understand how individual events are interconnected and the impact those threats have on an organization's assets.

This necessitates a combination of threat context and organizational context. The ability to connect the dots between phases of an attack is precisely what distinguishes a targeted attack from the stream of commodity threats that inundate networks on a daily basis.

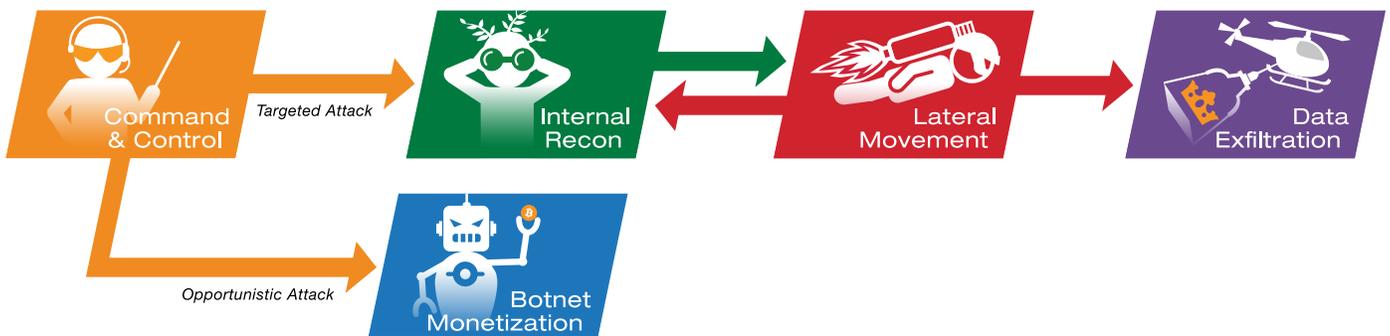


Figure 1: The lifecycle of a modern cyber attack.

Detecting threats with data science

In order to meet these requirements, data science and machine learning techniques can be applied directly to network traffic. The newest threat detection model uses both to proactively reveal hidden attacks inside a network.

Why use data science?

Data science and machine learning have become buzzwords in the industry, with a seemingly endless array of claims and applications. It is important to understand that these are simply tools and not a cure-all for every security problem.

To avoid marketing hype, it is essential to understand exactly what these approaches bring to the table, how they are different from other approaches, and the strengths and weaknesses they bring.

To learn about the data science that drives Vectra threat detections, [download the white paper](#).

Data science represents a fundamental shift in security. Unlike a signature-based approach that delivers a 1-for-1 mapping of threats to countermeasures, data science uses the collective learning of all threats observed in the past to proactively identify new ones that haven't been seen before.

Think of it as a student learning a new subject in school. Memorizing the answers to a test might result in a passing grade but this approach misses the mark when it comes to learning how to solve a problem.

Long term, it is essential to understand what, when, why and how. Actual knowledge and intelligence is far more advantageous when evaluating and solving new problems that have not been encountered before.

This is a critically important distinction when using data science to detect threats. For the traditional model to work, all of the answers must be known ahead of time. For example, the domain ACME.com has been seen behaving badly in the past, therefore it is bad.

Data science expects to be asked real questions and applies collective learning to evaluate an unknown.

In a different scenario, ACME123.com has never been seen behaving badly in the past, but traffic to and from this domain is showing four different behaviors, the combination of which is consistent with command-and-control attack behavior.

From the collective knowledge of threats gathered from the real-world, it's possible to identify the domain as acting bad based on its behavior.

Signatures



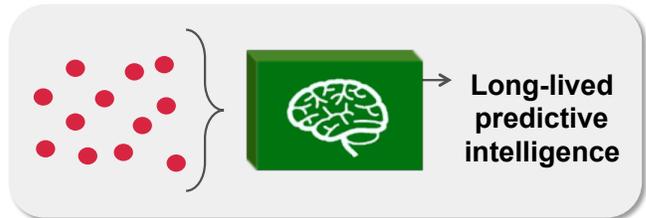
How the threat looks

Find threats that you've seen before

Snapshot in time

No local context

Data Science



What the threat does

Find what all threats have in common

Learning over time

Local learning and context

The importance of direct, first-hand data

Data science models are naturally dependent on the quality of data they analyze, and this is especially true in the cybersecurity field. It is essential for cybersecurity solutions to find stealthy threats that sneak past traditional controls, and this job requires direct, first-hand access to all network traffic.

The vast majority of approaches that use data science to detect threats perform data mining against large databases of event logs. While this approach may find correlations between logs that were previously missed, it has some troubling limitations.

Logs are a secondary source of data that briefly summarize an event. Information that is not in a log is lost and unavailable for analysis. Additionally, logs are only as good as the system that generates them. If an upstream firewall or security device fails to detect a threat, there will be no log to analyze.

The limitations of log data are disconcerting. The role of a cybersecurity solution is to detect threats that evade the standard layers of defense. Re-analyzing summaries from devices that already failed to detect a threat is hardly logical.

Re-analyzing logs from devices that failed to detect a threat is hardly logical.

NetFlow and other flow summaries suffer from similar limitations. Flow data monitors and tracks performance in the network plumbing. These summaries are limited to tracking the direction and volume of network traffic, and lack the direct, first-hand visibility needed to find a hidden, highly evasive threat.

Data science models perform much better when applied to higher quality data. Instead of simply mining data from flawed sources, the new threat detection model applies data science and machine learning to network traffic at the packet level.

This direct, first-hand access to traffic represents a completely new style of threat detection. Instead of correlating events or learning simple baselines, advanced threat detection builds real-time models that recognize the behaviors of malicious traffic.

Cyber attacks always evolve. But by retaining first-hand visibility into traffic, the new threat detection model can always adapt to identify new attack techniques and strategies. This is a stark contrast to log and flow-based systems, which are always dependent upon upstream data sources.

The role of machine learning in data science

The popularity and interest in data science and machine learning has turned both terms into slick catch-phrases, making it difficult to distinguish one from the other.

Data science is widely concerned with the many ways that knowledge can be extracted from data. Its sweeping perspective spans a broad set of *disciplines* that include mathematics, statistics, machine learning, and a variety of analytics just to name a few.

It is important to note that machine learning is a subset of data science. The new threat detection model also leverages a wide range of data science *techniques* that include supervised and unsupervised machine learning, mathematical heuristic models of detection, statistical modeling and behavioral analysis.

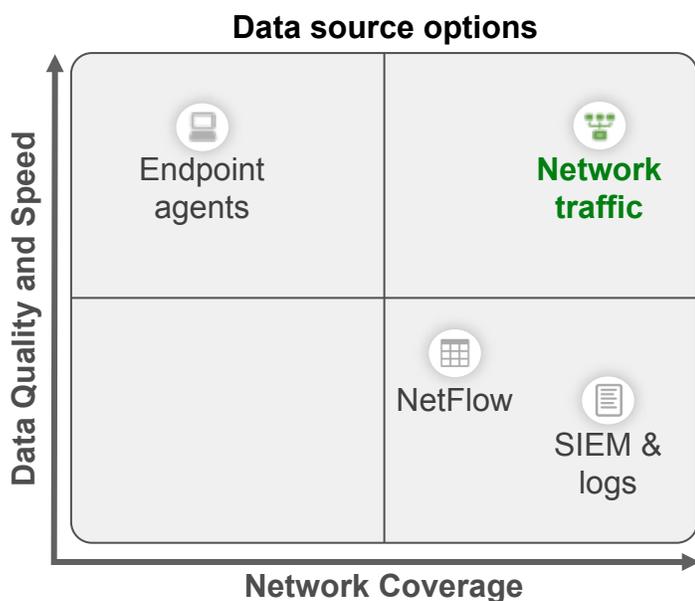
Machine learning enables software to iteratively learn from data and adapt without being explicitly programmed. In the context of detecting threats, machine learning identifies and learns patterns of behavior that reveal an attack.

In the context of detecting threats, machine learning identifies and learns patterns of behavior that reveal an attack.

Supervised and unsupervised machine learning

Detecting threats requires two types of high-level data sets. The first is a global set of experiences that indicates how threats differ from normal or benign traffic. The second is a local set of experiences that reveal unusual or anomalous behaviors in a particular environment.

The first approach reveals behaviors that are always bad, regardless of the network in which it occurs, and the second reveals threats based on local context. Both are crucial to detecting threats and they must work cooperatively.



Supervised machine learning addresses the first approach by analyzing known malware, threats and attack techniques. It is guided by security researchers and data scientists who identify fundamental post-exploit behaviors that are consistent across all variants. This analysis then feeds algorithms that detect underlying malicious behaviors in network traffic.

While global intelligence is extremely useful, some attacks are only revealed based on understanding the local context of the target network. *Unsupervised machine learning* refers to models that proactively recognize what is normal for a particular network and when behaviors deviate from that norm.

Both styles of machine learning are essential and work together to detect hidden threats. Likewise, both styles support detection algorithms based on information that is observed over extended periods of time.

Instead of detecting in a few milliseconds based on a single packet or flow of data, the new threat detection model learns and identifies attack behavior patterns over periods ranging from seconds to weeks.

Conclusion

The newest, most advanced threat detection model combines a wide array of industry-leading intelligence and detection techniques to see threats from all angles in real time. It represents a new, more effective and highly proficient detection methodology that leverages data science to detect threats that are missed by traditional security models.

To learn about the data science that drives Vectra threat detections, [download the white paper](#).