The Defender's Advantage: Using the Attack Lifecycle to Detect TTPs

Don't let Hollywood fool you: carrying out an attack doesn't resemble the plot of an action movie. Attackers don't automatically breach a network, immediately locate the information they want and then swiftly exit the organization. Attacks are complicated operations that unfold over multiple steps and take time, weeks and oftentimes months to achieve the desired goals.

The thought of adversaries carrying out a multiphase attack as they linger undetected in an IT environment for months may sound like a security analyst's worst nightmare. But defenders can use the lengthy period of time it takes to complete an operation to their advantage. Each phase of the attack lifecycle offers defenders an opportunity to identify and stop malicious activity before serious damage occurs, like the exfiltration of intellectual property.

Using behavior-based detection models to find the signs of an active cyber attack

Behavioral-based detection is ideal for discovering malicious activity, focusing in on what the attacker actually does, instead of relying on a set of signatures or known indicators of compromise (IOCs). Ultimately, if an analyst can't distinguish the good from the bad, then anomalies will remain noise, creating additional work for already strained analysts. By looking for the attackers' tools, techniques and procedures (TTPs), behavioral detection allows analysts to detect future unknown attacks as well as threats that are already in their environment. Unlike traditional IOCs, like malware signatures and hard-coded IP addresses, that are relatively simple to change, TTPs are challenging to modify and costly to develop. Like any software project, TTP development can take weeks or months and requires research, prototyping, testing, development and quality assurance.

This lengthy development cycle creates operational limitations for any campaign, even large sophisticated attacks. The bad guys have a finite supply of attack tools and the discovery of just one of them by the defender results in the loss of a strategic asset and jeopardizes the operation. Every TTP that's discovered depletes the attackers' arsenal. Given the effort it takes to create TTPs, they're meant to be used in multiple attacks. If enough TTPs are identified, the attackers may pause the operation and develop new tools before continuing the attack.

In one operation discovered by Cybereason, attackers suspended their operation after the PowerShell infrastructure they were using to carry out fileless malware attacks was detected and shut down. After a four-week hiatus, they returned with tools that allowed them to bypass the PowerShell execution restrictions that the company implemented.

All attackers, even the very skilled ones, leave behavioral fingerprints. Detecting just one provides defenders with the chance to piece an entire attack together and discover the full campaign. We're focusing on the penetration, lateral movement and command and control stages of the attack lifecycle since they are typically found in most attacks. But the principle of using behavioral analysis to detect malicious activity is applicable to all attack lifecycle stages.

Discovering code injection during the penetration phase

Every attack starts with penetration and establishing a foothold. The bad guys have to infiltrate a network somehow and establish a secure position to launch the attack's later stages.

Code injection is a great way for attackers to establish that foothold. Detecting malicious code injection provides an organization with a chance to detect the very early stages of an attack and reduce the damage. With code injection, attackers force an application to execute malicious code in addition to the code it normally runs. There are some techniques that allow one process to ask another process to run a specific code. This activity by itself isn't malicious, but there are vulnerabilities that allow attackers to map new code sections to an existing process.

Once attackers map whatever memory space they'd like, they can add malicious code, map it to the legitimate process and force this process to execute the code. This technique, called process injection, is very popular with attackers since it bypasses traditional security measures by using legitimate programs to launch malicious activity.

For example, take a protection mechanism that prevents any process from connecting to the Internet except those that are whitelisted. This whitelist will likely allow processes originating in Internet Explorer, Chrome and Safari to connect to the Internet (without this provision workers couldn't perform their jobs). If hackers can inject malicious code into one of those browsers, they'll have no problem connecting to an external, malicious IP address.

Looking for excessive port scans to discover lateral movement

When attackers move on to the lateral movement phase of the attack, they're scanning the network to learn what other machines are on it and what protocols they're using. They're looking for open ports that can be used later for exploits and lateral movement.

To obtain this information, the attackers will enumerate all the ports of another machine. For instance, let's say that there's a machine with the IP address 10.00.17. Attackers will try every port on that machine until they find one that's open. If the attacker discovers that port 443 is open, for example, they'll check to see if they have vulnerabilities or exploits that can be used on the open protocol.

Look for excessive port scans, which could indicate that attackers are conducting reconnaissance and attempting to map out your network. Intrusion detection tools can detect port scanners, but, admittedly, determining what's legitimate scanning and what's reconnaissance is tricky. Networks are filled with computers and applications that are constantly talking to each other so you have to filter out the noise, which can take awhile. But you can discover anomalies that indicate an attack if you know how many ports and destination the devices on your network typically access. Identifying port enumeration across machines can lead to detecting the early stages of internal reconnaissance and prevent the attack from spreading to other machines.

Detecting DGA use for command and control activity

Once attackers are inside a network, setting up a command and control infrastructure that allows them to remotely control the attack is crucial. Adversaries have stopped using hard-coded domain lists and IP addresses, which are useless once they're blocked, to communicate with the sophisticated malicious tools they've created. Instead, they're using domain generation algorithms (DGAs) to generate thousands of random domains. DGAs are easy to implement, difficult to block, impossible to predict in advance and can be quickly modified if the previously used algorithm becomes known. Attackers just have to register one. The malware will try all of them and eventually find the right domain. And if that domain is blocked, the attacker just has to pick another domain, register it and wait for the malware to connect to it.

But trying to find the right domain in a list of thousands generates lots of noise and failed connections. Check DNS logs for lots of failed DNS requests or requests that look like they were generated by an algorithm. To learn more about DGAs, check out this Cybereason Labs report, which dissects eight real-world DGAs.

Use an attack to your advantage

Behavioral detection challenges the notions that successful security means stopping every attack and it's game over for the good guys when a situation goes awry. Instead, attackers needs to successfully bypass every security measure the defender has in place once they're inside an enterprise. One wrong move or the slightest bit of noise and they risk exposing the entire operation. And carrying out a perfect operation is impossible given the complexities involved with infiltrating a target and maintaining persistence. Once you identify malicious activity in one portion of the attack lifecycle, you can use that information to detect the entire operation and stop it before more damage is carried out.

About Cybereason

Cybereason is the leader in endpoint protection, offering endpoint detection and response, next-generation antivirus, and managed monitoring services. Cybereason gives enterprises the upper hand over cyber adversaries. The Cybereason platform is powered by a custom-built in-memory graph, the only truly automated hunting engine anywhere. It detects behavioral patterns across every endpoint and surfaces malicious operations in an exceptionally user-friendly interface. Cybereason is privately held and headquartered in Boston with offices in London, Tel Aviv, and Tokyo.