There's a Hole in Enterprise Cyber Defense: Detection-based Protection is Not Enough

By Eirik Iverson

Industry analysts say that enterprises rely solely or mostly on detection-based cyber defense technologies. Their reports also demonstrate why detection-based tools alone are not enough. Breach volume increased 33% from 2019 to 2020 to 5,258 following a 96% increase from 2018 to 2019, per the Verizon DBIR. Mandiant's 2020 cyber incident investigations found that 65% of the organizations did NOT discover the attack within the first week. Machine learning did not magically fill these costly detection gaps. The remainder of this blog focuses conceptually on why detection is missing the mark. It concludes with a brief introduction of what the enterprise needs and the steps to get it.

Two Lessons from the SolarWinds Orion Supply Chain Compromise

Some of the targeted organizations had log data with many of the indicators for this attack. But, their SoCs were unable to interpret the data until after industry guidance was disseminated. It's yet another example of the skills gap that affects detection-based cyber defenses. Other organizations were found to be retaining log data too briefly with some affected firms deleting data after seven days.

Living-off-the-Land Attacks: Mimic Legit Workflows, Fool Detection

Industry analysts say these attacks are among the most difficult to detect with high confidence. Note the last qualifier. Pentesters and actual threat actors openly state they intentionally make their attack trees similar to legitimate workflows to avoid detection.

Threat Actors Fool Detection by Altering Known Malicious Files

File manipulation methods are numerous and varied. One can readily use a "packer" to alter an executable without breaking its functionality and then upload the file to VirusTotal, which then returns false negatives from dozens of detection vendors. This short video shows someone altering a 2017 WannaCry executable so AV/EDR tools do not recognize the file uploaded to VirusTotal. Adversaries still successfully recompile source code via a different programming language to slip files past detection. Uploading files to clouds for analysis has not been enough either because malware sometimes goes to sleep for a while or behaves differently in virtual environments to evade detection.

Machine Learning (ML) is not Artificial Intelligence; it's just Statistics

ML-based detection tools excel at detecting what has been seen previously and what is very similar to what has been seen. Evading them just requires a little extra effort. Fundamentally, ML can be fooled because it cannot leverage abstractions. It statistically correlates things to other things with no notion of what each thing is. So irrelevant things added to a basic attack tree disrupt the correlations, which confounds recognition. For example, adversaries alter previously recognized attacks by adding in pauses between stages, inserting unnecessary stages, and appending innocuous binary snippets from legitimate files to malicious files to fool ML. There are many anti-ML techniques. Some literally use ML tools to figure out how to fool ML-based detection tools. Behind malware attack headlines are ML-based detection tools that were fooled.

Disabled Endpoint Agents Hear and Report Nothing

Mandiant found that 9.8% of its 2020 cyber incident/breach investigations involved the T1562 technique from Mitre ATT@CK, which terminates or incapacitates endpoint security tools. Ryuk, REvil, SolarWinds, and many other attacks used this technique. An agent offline for about an hour is routinely not investigated, if even noticed. Attackers can be quite noisy within this window. This underscores the importance of blocking attacks in real time.

When Attacks are not Blocked Immediately, Credentials Go First, other Endpoints Go Next

The 2021 Verizon DBIR found that credentials are the most commonly stolen type of data. Over 40% of all enterprise breaches investigated involved stolen credentials. The first endpoint is followed by many others. There are tools in the hands of adversaries that automate finding, capturing, and using privileged Windows credentials to compromise Domain Admin credentials. These can run fast and noisy or slow and stealthy. The 2021 Verizon DBIR also observed that credential theft is not just used against the big enterprise. Attack tools make this too easy not to do. Some are explicitly designed to evade EDR detection.

Analysts say Human-Controlled Attacks Almost Always Beat Detection Tools

Human-controlled attacks require adversaries to get at least one remotely controlled process running on a target

host. The attack tools themselves, see below, are modular, consisting of numerous different attacks they can run à la carte. Once they discover what host protection is running, they can tailor their actions. Different attack payloads are quiet or noisy, better against some defenses than others, and feature multiple ways to either disable or evade detection.

Some industry analysts say such attacks always succeed in penetrating an enterprise. However, such characterizations are anecdotal and not statistically significant inferences. These broad assertions are founded in the understanding that attackers quickly get feedback as to what works and does not. Trends indicate threat actors are making more use of attack tools, such as those in the chart below, because they are successfully defeating detection-based defenses.

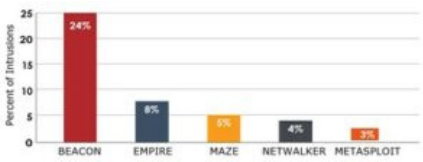

Cobalt Strike, originally a pentest tool, is the preferred weapon of choice by threat actors conducting human-controlled attacks (chart is from Mandiant Mtrends 2021).

The Missing Piece: Non-Detection based Endpoint Protection

At this point, cyber industry articles typically prescribe better cyber hygiene and two-factor authentication. Who hasn't heard that a thousand times already? The missing piece is a capability that neutralizes malware without having to recognize it. That addresses the fundamental shortcoming with all detection-based tools. They only succeed if and when they recognize malice.

All non-detection protection methods strive to prevent the adversary from successfully completing those actions necessary to achieve its goals. There are many options to consider. Meanwhile, it's important to realize that one needs to better understand what non-detection protection tools really need to accomplish so that you don't seek a unicorn.

The adverse impacts from detection tools missing the mark are crucial to formulating your expectations for a non-detection protection tool. It requires adding something to your cyber stack to subtract workload volumes from the rest of it. Once the IT/Sec-Ops symptoms are understood and qualified, then one can assess what a non-detection protection tool must do and not do. Any such tool requires some tuning. Focusing its mission minimizes the tuning. Test driving candidates separates the difficult from the easy. For now, the most essential point to take away from this is that detection-based protection alone is not enough; enterprises must add non-detection protection.