LIGHTCYBER

CLOSING THE BREACH DETECTION GAP

FLIP THE ODDS

USING ACTIVE BREACH DETECTION AGAINST ADVANCED ATTACKERS

WHITE PAPER / FEBRUARY 2015

Introduction

Advanced attacks such as the well publicized breaches against Target, Home Depot, JP Morgan Chase, and Sony Pictures have proven that today's attackers have the odds stacked in their favor. Attackers can launch unlimited intrusion attempts, consequence-free, until they find a successful technique to circumvent a target company's prevention systems. IT security operators meanwhile have to deploy and manage IT security infrastructure to stop every single intrusion attempt, or face the consequence that an attacker gains unfettered access to the corporate network. In short, the "bad guys" have unlimited opportunities with no risk, and the IT security "good guys" have to be right every single time. Not a very enviable position.

The rising success of targeted attacks in the past two years challenges us to determine how to flip the odds on the attacker. Twenty years of experience with threat prevention systems tells us that we can't reliably and comprehensively stop all intrusion attempts with systems like NGFW's, IPS, AV, Sandboxing and Endpoint Protection technologies. Presuming that targeted attackers can and will continue to successfully penetrate the network, and given that such intrusion is only an early step towards the goal of theft or damage, the obvious question should be can we build an internal security system to reliably detect active attackers once they have landed on the network? Can we stop attackers in their tracks?

We believe the answer is a resounding yes.

Rigging the Game to Lose

First, how did we get to the state we are in? The answer is simple: by focusing all of our effort on combating cyber attacks where the attacker has the advantage: at the network perimeter and endpoint. Currently the vast majority of security investment is focused on prevention, keeping the attacker out of our network and off our endpoints. What this translates to is increasing effort in constructing a perimeter with firewalls, next generation firewalls, sandboxes, etc. or endpoint protection in the form of anti-virus, host IPS and personal firewalls.

Problems with Technical Artifacts

Further, these tools rely primarily on the detection of technical artifacts, which muddies the waters. Technical artifacts include things such as virus signatures or engine detections, indicators of compromise, IPS signatures, blacklisted domains, etc. The problem with such detection is twofold. First, since these artifacts are static, definitions must be developed for each newly discovered attack method. As we've seen, no one can keep pace with the ever-changing attack arsenal and thus attackers can and will always be able to fool such systems. Second, detection of such artifacts doesn't necessarily correspond to a real and active (let alone advanced) attack, and therefore results in a terrible signal-to-noise ratio that paralyzes most security operations teams with hundreds and even thousands of alerts. While the detection of malware passing the perimeter may be accurate, it may not be linked to an actual detonation worthy of action. Indeed most alerts are just an example of a prevention tool bragging at a block (with an attention-stealing alert) or a detection (sandboxing) tool triggering an alarm on malware that may not have actually landed on a vulnerable system. With the millions of malware files flying around, this translates to a problem of both high false positives (a lot of noise), and high false negatives (misses). Not to mention that the volume of alerts creates camouflage within which advanced attackers can hide.

Giving the Attacker Unlimited Shots

We've rigged the game to lose because in this arena – connected to the Internet – the sophisticated attacker can launch an unlimited number of intrusion attempts with no consequence. Phishing email after email can be sent, drive-by-download sites setup, etc., and the cost and risk to the attacker is zero. These attacks are trivial to keep anonymous, and thus an advanced attacker can easily conceal their actions within the flood of "normal" bad traffic that pervades the

[1] http://blog.trendmicro.com/trendlabs-security-intelligence/joomla-and-wordpress-sites-under-constant-attack-from-botnets/
[2] http://www.darkreading.com/attacks-breaches/dropbox-wordpress-used-as-cloud-cover-in-new-apt-attacks/d/d-id/1140098?

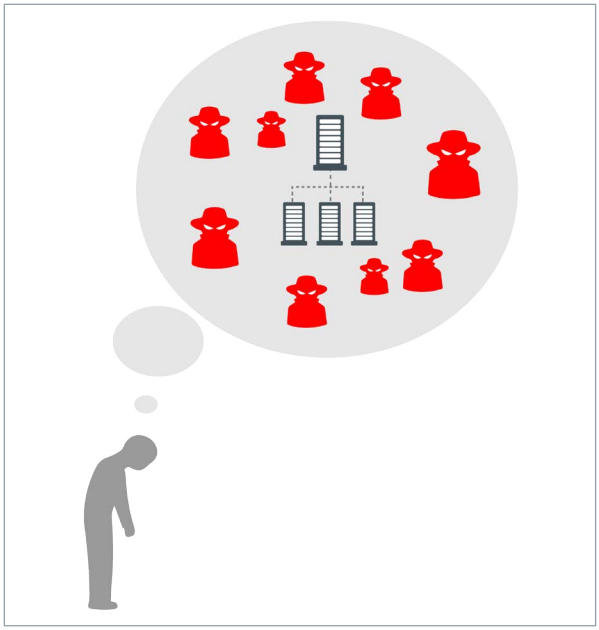

Before: IT security managers are overwhelmed and insufficiently armed to respond to attackers

Internet (e.g., opportunistic malware, bulk phishing scams, etc.).

And the attacker can keep trying until something gets in. Of course, focused attackers can also bypass many of these defenses through numerous mechanisms that make no use of malware at all: recruiting insiders, social engineering, even physical break-ins (though most networks are porous enough not to require such extreme measures).

Unfortunately, this means both that some opportunistic attacks land, but also that a persistent attacker will, with time, always be able to find some way to circumvent prevention tools.

The consequence is that we've set up the contest such that the odds are stacked in favor of the attacker. They can launch thousands of attacks, millions if needed, and if the defender misses even a single one, all is lost. And of course "missing" is a team sport: it isn't just about IT staff or security operations – if any end user clicks on the wrong email, follows the wrong link, plugs in the wrong USB, etc., the attacker is in.

Once the Attacker Gets In

So what happens once the attacker is inside? In most cases, the answer is also quite simple, if disastrous in scope: the attacker has free reign to expand their footprint in the network, identify assets of value, and exfiltrate them (or worse, exfiltrate and destroy as happened to Sony). The industry data back this up. The vast majority of attacks are identified by external agents (banks, law enforcement, etc.)[3], not by the victim organization. Further, the median time an attacker operates within the network before detection is 229 days[4]!

Given that the average attacker knows very little about the network and resources they are attacking, it is surprising that they can perpetrate their attack and remain undetected for so long. The attacker must be incredibly active in this period: they need to ensure they can sustain repeated access (often termed command & control, but which involves many steps), expand beyond their initial beachhead, map the subnet they are on, and find and explore others. In other words, they must perform active steps for reconnaissance and lateral movement. They need to identify assets of value, and determine mechanisms to gain access to same. Generally these steps require extensive network activity: port scans, administrative operations, outbound communications, SMB access, compromising additional credentials (even administrative ones) in a network-wide version of escalation of privilege, and finally affecting the movement of data. And none of this is subtle, even if the attacker attempts to go "low and slow."

The Defenders Unutilized Advantage

Meanwhile within the network, the defender should have the advantage. The defender already knows the topology of the network, where critical assets are located, who has administrative credentials and who does not, what limited set of servers typically produce scans, where critical data is and is not authorized to be stored, etc. The problem is, while this knowledge exists, it is not centralized, and it is not modeled in a way that enables the defender to quickly and automatically detect the variations (anomalies) that indicate attacker activity. If the defender could quickly isolate a specific case, it is often easy to realize attacker activity is ongoing. But, absent the proper tools and processes, it is impossible to gain such visibility through the vast noise of typical network activity. With current tools, it is actually quite difficult to gain this visibility even after the fact (i.e., after the 229 day period when an external entity has

[3] Verizon 2014 Data Breach Investigation Report
[4] Mandiant 2014 Threat Report

informed the victim of the attack), and usually requires a significant investment in forensics tools and/or services.

Further, the industry's historical reliance (nay, hope!) on comprehensive intrusion prevention via the detection of technical artifacts leads us down the wrong road since an active attacker's movements within the network heavily utilize standard IT tools and protocols, not necessarily malware5. Even in cases where malware is used, the attacker at that point has the opportunity to discover what malware detection tools are in use on the target network and select (and test) their malware tools to ensure they are ones that will not be detected.

If Knowing is Half the Battle, Automating Knowledge is the Missing Half

With the right technology, defenders can leverage the advantage of their control of their own networks to flip the odds on the attacker. Automated attack detection is possible, especially with systems that focus on attacker behaviors, rather than the technical artifacts of ever-changing tools and malware. An attacker must take thousands of disparate actions (i.e., behaviors) after successful intrusion in order to actually perpetrate an attack. If the defender can gather and analyze the signals given off by such universal attacker activity, the defender can get ahead of the attacker – e.g., figure out the credentials that are compromised and reset them, determine the server that is being used to cache data for later exfiltration and wipe it, and clean RAT's and other malware off compromised machines. In other words, IT can knock the attacker back out of the network, and keep knocking them out, until they

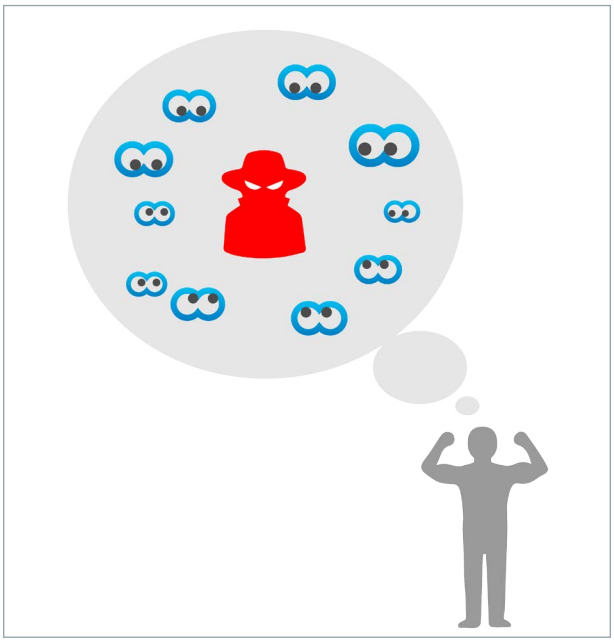

Flip The Odds: IT security managers armed with Active Breach Detection solutions can hunt down attackers

Conclusion: Flip The Odds

We call this flipping the odds, and we believe a new class of tools called Active Breach Detection is the way to do it. By automatically profiling typical user and device behavior, an Active Breach Detection system can continuously monitor network and endpoint behavior to detect anomalous attack behavior that deviates from the baseline. Rather than having to be perfect and detect every move an attacker makes, the defender only needs to detect and respond to one of the many behaviors the attacker takes en route to completing the breach. This puts all the pressure on the attacker. Unlike in legacy network and endpoint prevention, where the defender has to achieve the impossibility of perfection, now the attacker has to. And when the attackers are inevitably detected, they lose!

We are not saying that the current investment in prevention has been wasted; it is needed to stop the roughly 95% of known malware that it does stop. The pain of opportunistic attacks would be unbearable without it. But, further focus on either prevention or technical artifacts won't help against the targeted attack, because it is focused in an area where the attacker will always have the advantage. Instead, it is time to try something different.

It is time to flip the odds and make life tough for the attacker. Once we do, they'll move on to easier targets – ones that haven't deployed Active Breach Detection.

LIGHTCYBER
5050 El Camino, Suite 226
Los Altos, CA 94022

Ph: (844) 560-7976
www.lightcyber.com

[5] High profile cases of autonomous worms like Stuxnet are very rare, more the exception than the rule. And even such forms of malware still engage in many of the same ***behaviors*** as human attackers to spread and identify their targets, and are thus susceptible to the same behavior-based detection techniques.