InfoWorld®

EDITOR'S CHOICE: THREAT HUNTING

GET TECHNOLOGY RIGHT®

The threat hunter's guide to securing the enterprise

You're already breached. Here's how to track down attackers on your network before they wreak havoc.

It's time to face facts: Attackers are stealthy enough to evade your monitoring systems. If you're sitting back waiting for alarms to go off, there's a good chance you're already hosed.

Despite spending more than $75 billion on security products and services, enterprises are frequently compromised, highly sensitive data is stolen, and the fallout can be devastating. Worse, enterprises don't discover they've been breached for weeks to months after initial compromise, taking between 120 to 200 days on average to even detect an attack. That's a six-month head start on reconnaissance and exploitation -- more time on your network than most of your recent hires.

Needless to say, existing approaches to threat detection aren't working. It's time to strap on your threat hunting gear and proactively look for malicious activity in your environment. Here's a plan to track down threats.

Hunt in your own backyard

Threat hunting is a set of technologies and techniques that can help you find bad actors before they cause too much damage to your environment. Although threat hunting can involve both manual and machine-assisted techniques, the emphasis is on investigators looking at all the pieces in context and uncovering relationships, says David Bianco, a security technologist at Sqrrl.

Security automation can help collect data from network and endpoint segments, and machine learning can speed up analysis, but in the end, it's up to you to assemble a series of diverse threat hunting activities into a comprehensive process for sleuthing out your adversaries.

While a successful hunt requires you to think like a hacker, that doesn't mean you should be tracing attacks back to the originating machine, immersing yourself in Dark Web forums, or engaging in questionable practices to uncover potential issues. That may be the case for investigators and hunters from the U.S. Department of Defense or the Federal Bureau of Investigation, but threat hunting is purely defensive in the enterprise. You hunt by forming hypotheses about how an attacker can get into your network, then you look for evidence within your environment to prove or disprove those hypotheses.

Build a baseline of knowledge

Assessing security risk is a central facet of threat hunting, and the process can be split into three phases. First, you must understand the threats most likely to target your organization, whether they be persistent adversaries, particular sets of malware, or a certain type of attack. Second, you must identify your vulnerabilities, such as unpatched software or processes susceptible to human error. Third, you must assess the impact a successful threat may have in targeting your vulnerabilities. Once you can calculate these risks, you can then prioritize your threat hunting activities to target them.

Before you can start hunting, you need to understand the environment you are hunting in. This goes back to basic IT administration, such as having a clear picture of the number of systems, what software and which version is running, and who has access to each one. The network architecture, patch management process, and kind of defenses you have in place are all critical pieces of information in understanding your threat landscape.

IT teams need to know the weaknesses to identify potential points of entry.

Here, adopting an adversary mindset is key in determining your attackers' moves. Your attackers' motivations may vary wildly, but they often have similar goals and frequently share similar techniques. An adversary intent on cybercrime will typically behave differently from one focused on economic espionage or sabotage, for example.

Threat intelligence is one way to receive information about the kind of attacks hitting similar-sized organizations in the same industry

If a number of competitors has been under attack by a gang using a Flash exploit, it makes sense to prioritize investigating potential Flash-based attacks over other types. Knowing exploit kits and other types of malware are all pushing the same dropper payload is helpful.

It's also essential to ascertain what might interest an attacker most about your organization right now. This could be a new product your organization is working on or rumors about a potential acquisition. When you know what might trigger interest from potential attackers, you can better predict what techniques they will use and how they will traverse your network to get what they want.

Map the kill chain

A few years back, Lockheed Martin put forth the "cyber kill chain," which divides targeted attacks into seven distinct phases: reconnaissance, weaponization, delivery, exploit, installation, command and control, and action. Attackers typical move through each step, from initial compromise to theft,

getting a lay of your environment well before exfiltrating any data. A targeted attack takes time to develop; detecting the breach and blocking the attack as soon as possible will minimize damage..

During reconnaissance, criminals collect information about potential targets and avenues of attack. In the case of an acquisition, an attacker will collect information about executives and assistants who could potentially be working on the deal. Based on the information gathered, the criminals develop a course of action, such as creating a phishing campaign.

A successful hunt involves examining each phase of the kill chain and assessing specific tactics and techniques attackers may employ. That may involve mining social media postings to determine whether anyone working on a possible acquisition may have identified themselves as working on the deal and creating a list of employees who may be potentially targeted by a phishing email. If you believe phishing is the likely entry point of a targeted attack, then you can make assumptions about what the attack scenario will look like along each phase of the kill chain.

Actively hunt for threats

Your assumptions and hypotheses about potential attacks provide places to start your hunt. Successful hunting involves examining a specific segment of your network without trying to see everything that may go wrong. It's about closely scrutinizing an endpoint for specific indicators of attack rather than getting a bird's-eye view of system security.

Most threat intelligence efforts focus on indicators of compromise that don't help with threat hunting. The factors tend to be cheap, fragile, and inexpensive for adversaries to change. Consider domain names or the name of the weaponized Word document carrying the payload. It is trivial for attackers to generate new domain names and to change the messaging in an email accompanying an attack file to bypass security filters. Instead, hunters should focus on patterns of attack.

For example, you should look out for attempts to open a remote desktop session to create new admin accounts within Active Directory. It doesn't matter what the new accounts are called -- you should be searching for unexplained accounts.

It's trivial for an attacker to change the domain of a command-and-control server, but far more expensive to give up using a Flash exploit delivered via a malicious advertisement to remotely execute code and open a backdoor on the compromised machine. Look for attackers using legitimate tools such as PowerShell and WMI. See where account credentials are being used. Patterns of attack reveal more about attackers than indicators of compromise because they are relevant for a longer period of time.

Next-generation firewalls, anomaly detection platforms, and logs all provide a wealth of information, as do threat intelligence platforms and network threat detection systems. In many cases, there is a silo effect, with information locked within each system, making it difficult for defenders to see all the related pieces. Threat hunting forces defenders to break out of the tendency to consider systems in isolation. When a process touches different segments and systems, hunters must pay attention to how they relate to each other.

Build up security response

Once you find signs of a breach, threat hunters should step aside to let traditional incident response teams take over. The hunter's job is to make guesses as to where the attackers may be within the network, but they aren't necessarily those with the expertise to block attackers. Incident response will be in charge of mitigating the attack and remediating issues.

It may be tempting to create specialized hunt teams because they pinpoint problem areas and find the attacks, but that shouldn't be at the expense of basic IT administration, network monitoring, and defense-in-depth strategy. Threat hunting starts with the assumption "I have been breached" and looks for evidence to support that assumption, and dedicated incident response and forensics kick in when that evidence has been found and the damage has to be contained. They are very distinct skill sets, and both are necessary. Defenders need all of these elements to work together.

Stop the cancer

Threat hunting isn't a new concept, and many organizations have already adopted some form of the practice as part of their overall security plan. In a recent SANS Institute survey, 86 percent of IT professionals said they had implemented threat hunting processes in their organizations and 75 percent claimed threat hunting had reduced their attack surface.

As with every other aspect of information security, there's a time and place for threat hunting. Enterprises should look at the Hunting Maturity Model developed by Sqrrl's Bianco to judge if they are ready to begin hunting. The model defines maturity based on three factors: the quality of data collected, the tools available for accessing and analyzing that data, and the skills of those performing the analysis. A skilled enough analyst with high-quality data can compensate for deficiencies in the toolset, but for the most part, organizations should focus on all three factors.

"In order to get anywhere, you must first know where you are and where you want to be," Bianco wrote in a blog post outlining the model.

Enterprises need to reduce the breach detection gap -- more than half a year to discover a breach is unacceptable. Start with the assumption that attackers are already present and keep looking until either the compromise has been found, or there's conclusive proof that your environment hasn't been compromised.

Think of the enterprise as a biological system that has been infected, and threat hunting as a way to discover how far the infection has spread and what kind of damage it is causing.

"Enterprises should look at the Hunting Maturity Model developed by Sqrrl to judge if they are ready to begin hunting."

InfoWorld