LMNTRIX

BE THE HUNTER I NOT THE PREY

Thinking of starting a adversary hunting program?

Then you must do these to get it right.

Enterprise IT security strategies, processes and technology stacks are fundamentally defense-based, requiring known information about the adversary. They depend on prior knowledge of the adversaries' tools to drive detection and prevention. Adversaries continue to be on offense, targeting specific organizations, identifying attack paths to valuable assets, and deploying customized malware variants, intrusion techniques, and tailored towards organization's infrastructure.

In fact, 70-90% of the malware used in data breaches are unique to the victim organization, often leading to complete circumvention of signature-based enterprise defenses. Adversaries compromise at will, penetrating defenses in ways that leave companies ignorant of a breach for 3-9 months depending on which research you read. A dollar of offense always wins against a dollar of defense. Traditional security programs are bureaucratic and compliance-minded, while adversaries are committed, creative, and nimble.

Security teams must be successful 100% of the time, while attackers only need to succeed once to enter enterprise networks and cause damage and loss. A different approach is needed. Enterprises must assume that their networks are compromised and implement an offense-based strategy. This requires a shift in mindset, wherein enterprises think like the adversary and deploy the same creative and nimble tactics, techniques, and procedures that the adversary uses against them. Enterprises must hunt for adversaries within their networks.

Security Teams must think like adversaries, actively identifying adversaries without known indicators of compromise, and evicting them before data is exfiltrated or systems are disrupted. But even after following online tutorials, attending webinars and workshops – you might be struggling to achieve any justifiable success with your adversary hunting program, or simply lost with the large amount of data generated from the first hunt.

In this paper we have laid down the top 10 most important tasks to perform to make your adversary hunting program a success.

But first a quick definition of what is adversary hunting?

Adversary hunting is the stealthy and surgical detection and eviction of adversaries within your network without prior adversary knowledge or known indicators of compromise. The goal of hunting is to detect and evict adversaries that have bypassed defenses before damage and loss can occur. To do so, a hunter must be able to enter the network undetected, identify the adversary at any stage of the kill chain, and evict them without disrupting running systems. There are three key components of adversary hunting: stealth, early detection, and surgical response.

Stealth

Adversaries are looking for you as much you are looking for them. They hide and adapt their behavior upon detection of any traditional security tools. Enterprises must be stealthy and hide their presence from these advanced and adaptive adversaries.

Early Detection

Enterprises are often informed by a third party about a compromise on their networks, about 53% of breaches are detected by third parties, by then the damage has already occurred. Adversaries need to be rapidly detected at all phases of the kill chain to stop them from gaining unauthorized access to critical systems and reduce the damage they can to inflict on the enterprise.

Surgical Response

Most adversaries target mission critical systems within enterprises, which are crucial to daily business operations. Once the source of compromise has been identified, these adversaries must be stopped. It is key to remove them surgically without any business disruption.

Top 10 Adversary Hunting Best Practices

1. Define a Standard Methodology for Hunt

Much like any other business process, the adversary hunting process needs to be defined using a methodology that is repeatable and can be verified. SOC analysts should be able to perform steps outlined by the process and generate a consistent finding report that allows all the team members to understand, collaborate, or, verify the results. A standard process helps the team to independently review the findings and provide their input about threat hunts that indicate if an incident has taken place or if it was a false positive, which in turn increases the level of the hunting maturity for the organization.

The following is the LMNTRIX hunt cycle that you can use. We have pioneered the LMNTRIX Hunt Cycle to implement this offense-based hunt approach. Our hunting methodology enables security analysts to stay ahead of attackers by detecting them at all stages of the kill chain. By automating the hunt cycle, security analysts can stop adversaries before damage and loss occurs. To detect and evict adversaries in the network, LMNTRIX's unique methodology consists of four phases: Survey, Secure, Detect and Respond.

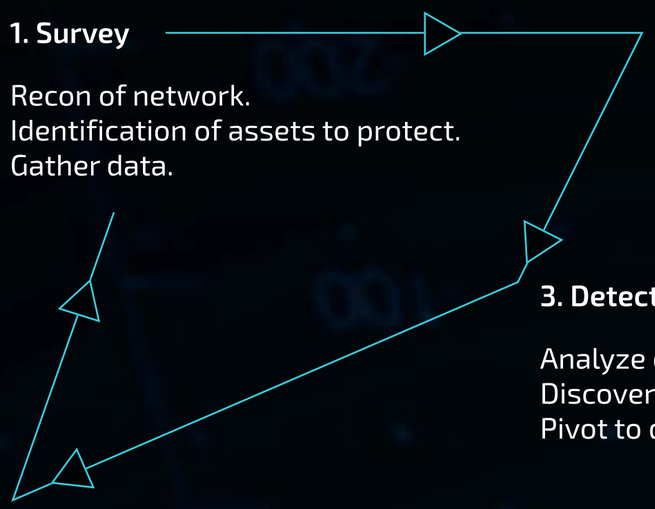

2. Baseline the environment

It is of utmost importance to understand what "normal" looks like within your organization. Identifying the baseline of your environment can be a tedious task, but once you know how noisy or quiet things are in your environment, finding an anomaly becomes significantly easier. This step might be done differently for an in-house SOC and an MDR, as the in-house SOC should already know about the IT administrative tools and processes, while it will take an MDR a longer period to understand the same.

Then there would be those tools or activities that a certain group within the organization (non-IT) may be performing that needs to be overlooked or allowed, such as, developers would compile programs most of the time which might trigger some generic Yara rules, or Finance team might be working on an Excel sheet which has macro-code to pull live forex currency rates, to do some calculation and save an output.

3. Use High-Quality IOCs

Using large number of Indicators of Compromise as they are consumed from threat intelligence suppliers will only generate unwanted noise and alert fatigue by your SIEM. This practice is not suitable for adversary hunting. The adversary hunting team needs to carefully curate IoCs. Instead of sweeping the environment with a large quantity of IoCs, hunters must use a smaller number of high-quality IoCs that fit their environment's threat profile. Curation not only requires selecting a subset from a vast pool of IoCs but also putting context behind every single IoC to give them meaning.

Finding common IoCs for a known threat actor, say Turla malware campaign, might point your adversary hunting team towards digging deep and performing more hunts to confirm the presence of the Turla malware or it's dropper and other variants, or at least attempts of exploitation of 0-days known to be used in Turla malware campaign. Notice that IOCs, just does not mean using IP Address, Domain or Hash, but also Tools, Techniques and Procedures of an attack. STIX is what we use to convey a threat profile by describing relationships with traditional IOCs and TTPs.

Top 10 Adversary Hunting Best Practices

4. Anomaly Detection

Anomaly detection generally creates a baseline and detects outliers by using machine learning algorithms or implementing proprietary static algorithms. In the context of threat hunting, anomaly detection is a more iterative and open-ended process. Ideally you would perform your anomaly detection processes at regular intervals, such as, 7 days, 15 days, 30 days, 45 days, 90 days. Each time interval will help you dig out different anomalies and based on the baseline, it will help your team determine if they were confirmed anomalies or a false positive.

One example we use for example is data stacking, where an analyst acquires a set of data, such as a list of all running processes within the environment. The analyst then counts the occurrence of every unique process throughout the environment to create a baseline. Because targeted malware is the exception rather than the norm, it will show a low frequency of occurrence. The power of stacking lies in the combination of different stacks that skilled hunters build dynamically based on what they find during the hunt.

5. Hypothesis Driven Hunting

Most online articles and blogs simply state, "Ask a question and go perform adversary hunting to find answers". Simply stating, 'Go and threat hunt' does not help as a structured process is required that defines steps to conduct threat hunts. Developing a hypothesis is a crucial step in the adversary hunting process. A hypothesis is formulated on how an attack could happen. That hypothesis relies heavily on threat intelligence about the organization's specific risk profile. Building a proper hypothesis can help identify ongoing attacks and even rule out which hypothesis is not relevant. This however requires high visibility in the environment.

Making hypothesis starts from looking at the organizations' threat model. You can make “what-if” hypothesis, such as:

- Open source/Commercial tools, like PowerShell Empire or Cobalt Strike are used to perform attacks
- Spear phishing Link is delivered via legitimate email contacts
- Public facing web application is being exploited to gain remote access
- A disgruntled employee is trying to take out confidential files from office workstation
- Chinese/Russian/North Korean hackers are already inside the enterprise network

Top 10 Adversary Hunting Best Practices

6. Reduce Visibility Gaps

Speaking of having high visibility to create hypothesis-based hunts brings us to our next point, that is, identifying and reducing visibility gaps. The desired process of identifying visibility gaps is to have a hypothesis first and then figure out where and how to get the data needed to accept or reject the hypothesis. If that data is not available, the hunters have identified a crucial visibility gap. Only visibility gaps that come with a sound hypothesis are a sign of good adversary hunting practice.

Using examples from the previous point, we can make a rough list of data sources needed to perform hunt and validate our hypothesis:

- Process Audit log, PowerShell logs, WMI logs, Network logs Etc.
- O365 logs, Exchange logs, Network logs. Etc
- Network logs, Application audit logs, Process audit logs. Etc.
- Network logs, Removable device audit logs, File audit logs, O365 logs. Etc.
- Process audit logs, PowerShell logs, WMI logs, Network logs, Full packet capture, O365 logs, File audit logs. Etc.

7. SSL Decryption for Deep Packet Inspection

Even with high level of visibility, it is impossible to understand the flow of encrypted data. Most solutions and tools used by a SOC will record host communications, ports, protocols, and traffic volume, while network flow data gives a high-level overview of the patterns of network communication that are ongoing within the network. However, analyzing live decrypted packet captures provides a much better granular level of visibility to analysts.

Full packet capture comes at a cost of disk space, but it is worth the visibility in critical networks even if retained for maximum of 30 days. Start by using open source tools such as Squid MITM proxy with tcpdump. Run Zeek scripts on your packet captures to find suspicious activities. In our experience we find meta-data instead of full packet capture is sufficient for investigating over 90% of threats.

8. Adopting MITRE ATT&CK

MITRE ATT&CK is a breakdown of all the Tactics and Techniques used by the attackers and, understanding and using those terms and concepts in adversary hunting operations will provide context to the findings. Each MITRE ATT&CK Tactics contains a set of Techniques that not only describes the technical description but also about the threat actors' intentions when they employ a certain technique.

Once this comes into practice, it becomes easier for threat hunters and management to focus on specific threats targeting their industry, attempts of attacks, or, the presence of attackers within the organization. Perhaps the most useful outcome of MITRE's ATT&CK Matrix is the ability to differentiate between the tradecraft of threat groups.

Top 10 Adversary Hunting Best Practices

9. Automation

Automation is widely used in many different IT processes, and adversary hunting is also a candidate for automation. Analysts can focus on manual data analysis and work on new hunting scenarios while the automation takes care of the repetitive tasks of existing hunts.

However, we can never fully automate threat hunting. We can automate data retrieval and transformation steps most of the time, but the analysis and interpretation of data require human analysts. Machine learning solutions can help, but decision making based on the analysis still needs human judgment.

The following is a small sample of the daily automated hunt that we run at LMNTRIX across our client networks.

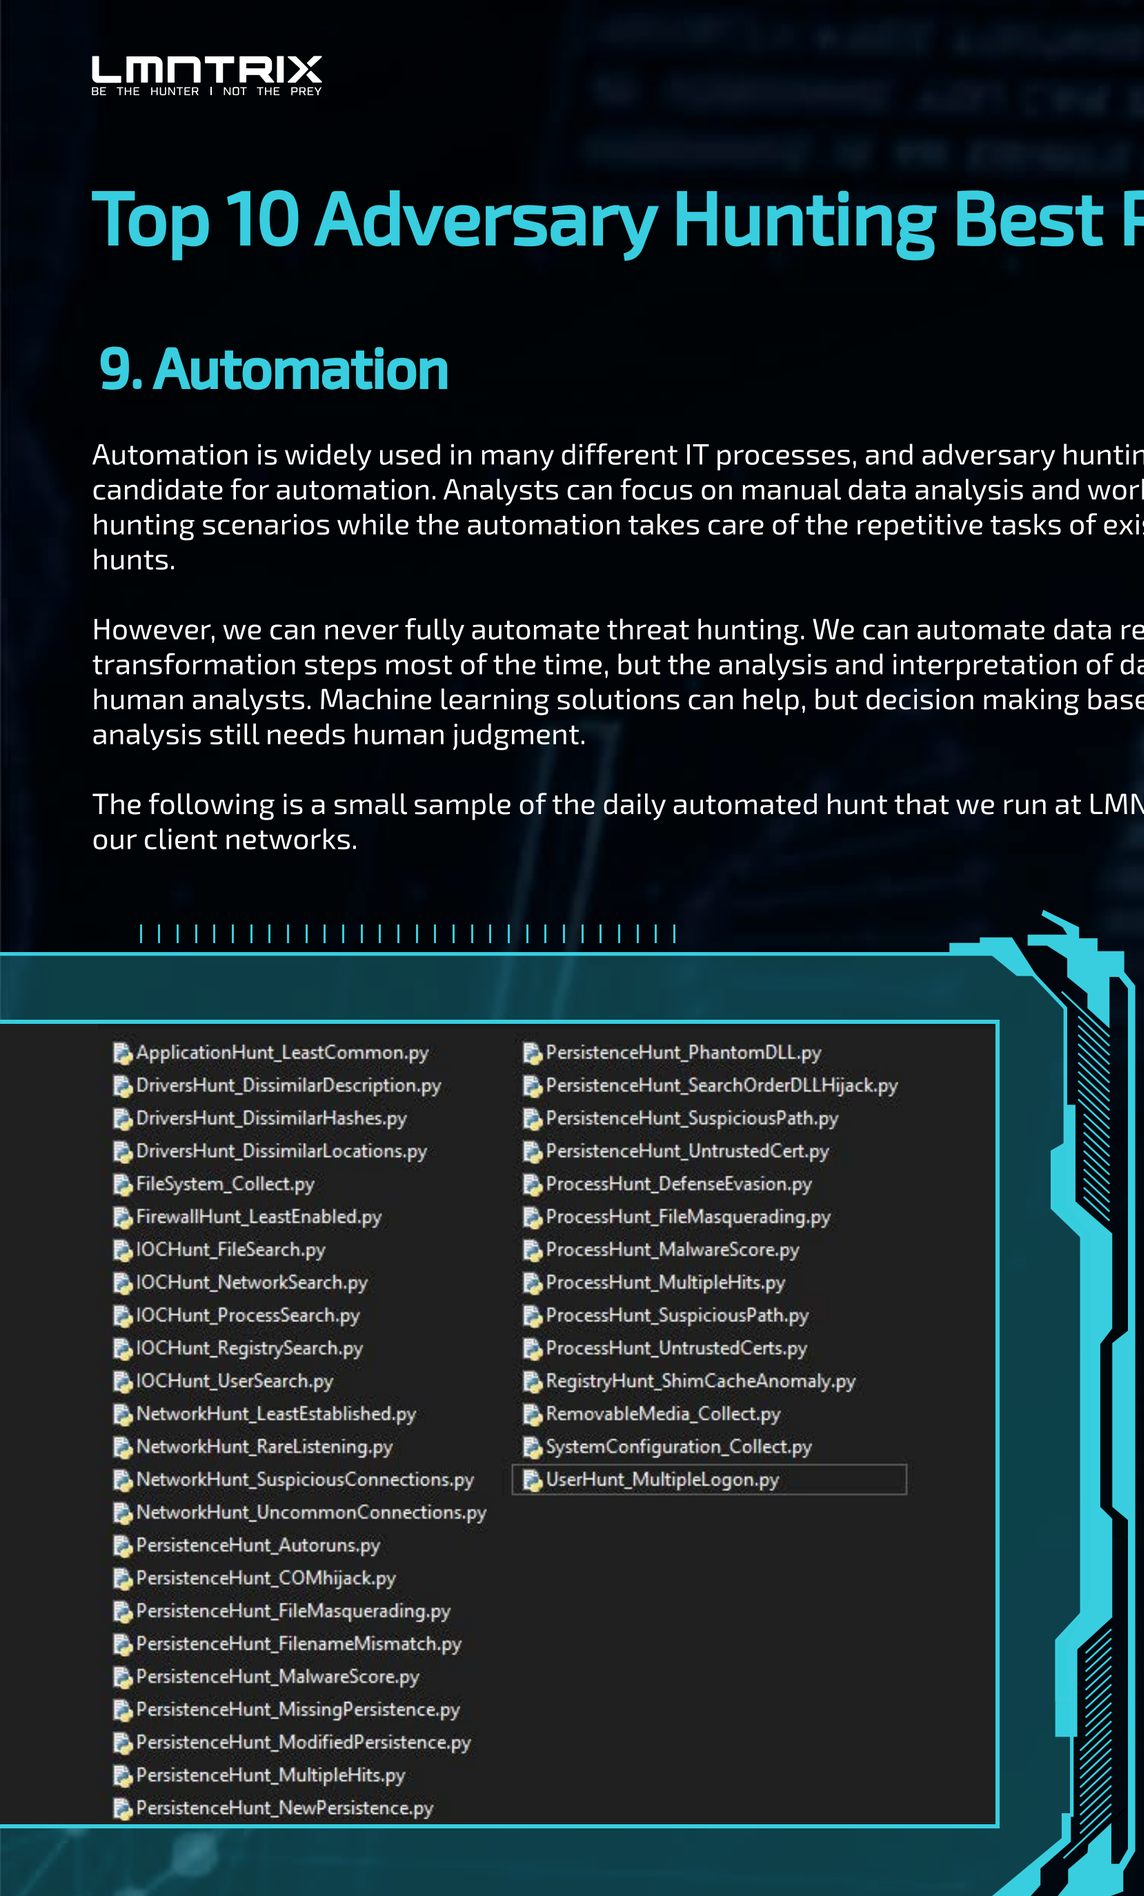

10. Measure Success

Move away from the “amount of attack detected” assessment approach to a “risk-based damage assessment” approach, that is, measure success by the amount of damage averted from a hunt. That means besides financial harm, you would also need to factor in the loss of reputation, legal implications, theft of intellectual property, etc.

When adversary hunting activities catch the attacker early on, it is usually impossible to calculate the impact of what would have happened. Organizations that sport a mature risk management process might be better off because the various risks that support adversary hunting metrics may already have price tags attached. A common approach is by labelling critical servers, workstations, and users as high value targets. While helpdesk and front desk users and workstations would be medium value and low value targets, respectively. This however is subjective to each organization.

In Conclusion

Enterprises need to think offense, adopting an adversary mindset to proactively and dynamically tackle the challenges of the modern threat landscape. Hunting operationalizes this mindset, and is an increasingly modern necessity for enterprises to protect their critical assets. The key to hunting effectively is to enter the network undetected, detect the adversary early and at every stage of the kill chain, and evict them without disrupting running systems. Hunting allows enterprises to proactively protect their most vital assets before any damage and loss occurs.