BlueVoyant

Threat Hunting

In Today's Threat Landscape

MICHAEL SCUTT

ABSTRACT

There are many variables involved in threat hunting today. Threats are evolving at a rapid pace, attackers are becoming increasingly persistent, and attacks have become progressively more difficult to remediate.

It's vital to block attacks before they do harm to your customers and reputation. Detecting emerging threats isn't enough. Predicting what your adversary might do next is necessary in today's ever-changing threat landscape.

Attacks occur across every industry, in every vertical, and every size business. Your organization must have the skills and resources to keep up in this constantly changing landscape.

Attacks vary depending on the adversary and their motivations. An opportunistic "smash and grab" attack may last from a few minutes to an hour. Other adversaries have a specific goal in mind that might require reconnaissance, which can take weeks, months, or in some instances, years to ultimately complete the attack lifecycle. You need to understand your adversary in order to understand where you are in the attack lifecycle.

This eBook examines the management of the Cyber Attack Lifecycle when threat hunting, the behavioral aspects of a threat, and how BlueVoyant manages this process.

AUTHOR BIO

Mike Scutt

Director of Threat Operations

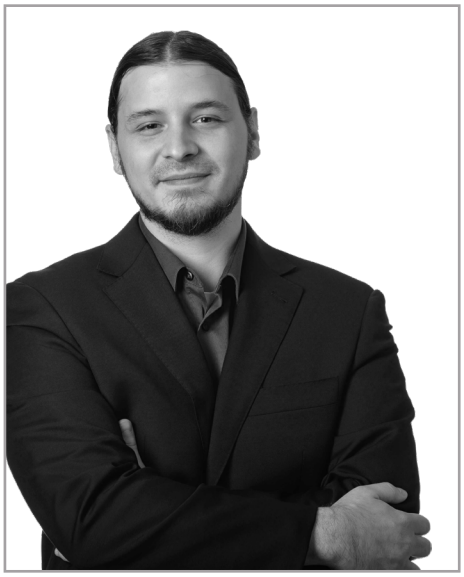

Mike Scutt leads Threat Hunting services at BlueVoyant, assisting clients in uncovering advanced adversaries, cutting edge malware, and attacker infrastructure. With a lengthy background in responding to breaches by nation-state threat actors and hundreds of incident response engagements, Mike applies threat intelligence, forensics, and malware analysis techniques in the search for attackers.

In his role at BlueVoyant, Mike oversees the creation of analytics to identify malicious activity commonly unseen by traditional security tooling and the implementation of threat intelligence for Managed Services clients.

Mike comes to BlueVoyant from CrowdStrike, where he served as the Director of Security Research, overseeing detection strategy for the Falcon EDR platform. Prior to CrowdStrike, Mike was a founding member of Rapid7's Incident Response and Managed Services businesses and was the Endpoint Lead Incident Handler for Mandiant. With over a decade of experience working in the Incident Response, Digital Forensics, and Malware Analysis fields, Mike has spoken at multiple industry conferences and holds patents in malware detection.

ABOUT BLUEVOYANT

BlueVoyant is an analytic-driven cybersecurity company whose mission is to protect organizations of all sizes against agile and well-financed cyber attackers. Founded and led by experts in the cyber security and government security sectors, BlueVoyant's offerings are built with real world insight and applicability, plus an eye on the threat horizon.

Through our Advanced Threat Intelligence, Managed Security Services, and Incident Response Services, we excel in intelligence gathering, cyber security defense, detection of attacks and response coupled with remediation.

TABLE OF CONTENTS

INTRODUCTION

There are many security tools on the market designed to stop low-sophistication threats from getting into your system. However, as threat actors become more agile and sophisticated, their attacks aren't easily detected by standard automated security tools. Every industry and vertical is at risk of attack.

Threat actors are steadily devising meticulous, well-planned attacks that evade your traditional security measures. It's not uncommon for an attacker to be inside of your network long before you notice their presence. Frequently, by the time an attack on your network is detected attackers have had an opportunity to set the stage for a persistent attack, and they've evaded your automated security measures. Threats are becoming increasingly difficult to detect and prevent.

This eBook offers you insights into the cyber attack lifecycle and examines the human elements of an attack. Key takeaways include an understanding of the pre-planned, sophisticated, and malicious nature of attacks in which the attackers have breached your system for days, months, or even years and BlueVoyant's unique approach to solving these challenges.

CHAPTER 1: CYBER ATTACK LIFECYCLE

The cyber attack lifecycle and the manner and methodology in which adversaries work is particularly important when trying to detect them through all stages of an attack. The cyber attack lifecycle consists of 8 stages:

- Reconnaissance
- Initial Exploitation
- Establishing Footholds
- Privilege Escalation
- Internal Reconnaissance
- Lateral Movement
- Persistence
- Data Theft / Destruction

RECONNAISSANCE

The first stage is reconnaissance. Initially, we tend to see adversaries do some level of reconnaissance or probing, or if we're looking at organized crime, talking to others about what they intend to do or who might be an interesting target for credit card theft or other financial damage. You must monitor deep and dark web forums, and also the internet in general, for chatter related to potential attacks.

Indications that an adversary is probing you include registering lookalike domains or actively scanning your infrastructure or third parties with which you work. Evidence of these activities indicate that you should be prepared for an attack and should set off internal alarms.

INITIAL EXPLOITATION

Once threat actors identify you as a target and understand how they might be able to manipulate or exploit you, they move into the initial exploitation stage of the cyber attack lifecycle. Often, the initial exploitation begins with phishing emails or similar schemes. It's common for threat actors to use social media, email, and instant messaging to entice you and your unsuspecting employees to download their malware or visit one of their compromised sites.

It's also becoming common for threat actors to compromise the third parties you do business with to masquerade as that third party. Once the threat actor appears to be a third party you usually do business with, they can easily trick you into changing bank numbers, downloading and running malware, or any other seemingly usual activities. At this point, you likely have no idea that you are compromised.

ESTABLISHING FOOTHOLDS

Once a threat actor has a foothold in your environment, through malware, a vulnerable system interface, remote desktop, or TeamViewer, their next step is to establish a means for command and control, either through malware, exposed RDP, or VPN access. Attackers may enter your system as an unprivileged user or assume the identity of the initial victim, perhaps your unsuspecting employee.

PRIVILEGE ESCALATION

Entering your system as an unprivileged user may not give attackers the access they need to accomplish their mission or allow them to identify where the targeted information resides. To gain the necessary access, your adversary may start dropping tools like credential dumpers or key loggers into your system. They might start manipulating data and gathering other credentials in order to gain the privileges required to move laterally or perform internal reconnaissance. Many times your adversary disguises themselves as a domain owner or the owner of a system of interest, for example, a database server, to gain trust and get one step closer to accomplishing their mission.

INTERNAL RECONNAISSANCE

Attackers need to get their bearings after compromising a system within a target organization, even after performing extensive reconnaissance of your company from the outside. Typically the “patient zero” in a compromised network is still an opportunistic attack; the attacker may only send a phishing email to three victims in an attack, but they can’t anticipate who will open the email and which system they will gain access to. Following a successful compromise, attackers need to enumerate accounts, systems, and business processes to find the best path to mission success.

LATERAL MOVEMENT

Once an attacker has established which systems contain the data required to accomplish their mission they will need to move laterally from compromised system to others in the network. Attackers will frequently use built-in capabilities of the operating system to accomplish lateral movement: RDP, SSH, SMB, and internal phishing. Identification of lateral movement is challenging due to most attackers’ use of valid accounts and pathways that are open to enable day-to-day business operations.

PERSISTENCE

Attackers recognize that workforces are largely mobile and that they may trip an alarm during their efforts. To decrease the likelihood of losing access, attackers will deploy additional means of access, including creating new accounts, deploying additional backdoors, and creating new command-and-control infrastructure. Additionally, attackers want to ensure that their means of access are resilient - surviving reboots, network changes, and defensive measures. Persistence takes many forms, but most commonly attackers use built-in Auto-Start Execution Points (ASEPs) native to common operating systems to maintain access to compromised systems.

DATA THEFT / DESTRUCTION

The final stage of the attack lifecycle is typically data theft (intellectual property, competitive intelligence, financial records, PII) or data destruction (ransomware, denial of service). The attacker's end goal is not always clear, once data has been stolen it is difficult to understand how an adversary, nation-state, or organized crime group will make use of it.

At this stage, your adversary has identified you as having data of interest, made the effort of compromising your system, and established a firm foothold within your system. Your role, as the unsuspecting victim, is to continue to generate data, build new products, and add users. Your adversary stays in your system environment and collects this new information as it's generated.

At BlueVoyant, we base threat hunting on the cyber attack lifecycle and encompass each phase in our process. Our threat intelligence program not only addresses the cyber attack lifecycle, but also the human aspect of breach functions.

CHAPTER 2: KNOWING YOUR ADVERSARY

On the other side of every attack is a human, whether the developer of a malware family or a hands-on-keyboard attacker. The most important thing to remember throughout threat hunting and detection is that even when you're dealing with mass malware, there's a human on the other side of the keyboard, and humans have predictable patterns.

Our deep understanding of the human aspect of breach functions enables us to anticipate which variant of malware the adversary may use to get into your environment, where they host their infrastructure, and the methods they will use to carry out their attack, which are known as “Tools, Tactics, and Procedures” (TTPs). Our goal is to identify their presence, track them, and evict them from your environment.

Human adversaries tend to repeat behaviors: They have certain ways of typing commands. They have specific techniques they use to maintain their presence. So, even if they swap the back door they’re using, or bring in a new piece of malware, we can identify them through the footprints they leave.

Humans make mistakes and threat actors are human. As defenders and threat hunters, you need to recognize and take advantage of subtle mistakes or repetitive behaviors and actions that your adversary takes.

Your adversaries have goals. For example, Nation-State adversaries tasked with identifying and acquiring intellectual property around a mechanical design always have similar goals. That means we can anticipate the types of data they’re going to go after. We can anticipate the industry verticals and even the specific clients that they are attracted to. This means that we can better prepare ourselves, and build out our hunting methodologies, well before an adversary has set their focus.

We can use behavioral data to take an in-depth look at our adversaries and determine their infrastructure. We watch for their most recent infrastructure changes - within the last 24-48 hours. Then we use the information that we collect to better protect our clients from potential threats.

CHAPTER 3: DEVELOPING A THREAT HUNTING PROGRAM

We developed the threat hunting program at BlueVoyant because we understand that even the best technologies on the market, while very good, often fall a little behind the curve. Next-generation antivirus, next-generation firewalls, and many other detection technologies rely on data that comes in from users. When the data, or events, reach a critical mass, it becomes something that they incorporate into the security solutions. Primarily, developers look into new campaigns by volume and then use that to develop new detection or preventative tools.

BlueVoyant understands many of the technologies on the market don't fit our client profile very well. Waiting for an antivirus vendor to develop detection for an emerging threat doesn't serve our clients, especially if we're looking at something targeted.

USING A BEHAVIORAL APPROACH

We hunt proactively using our deep insight into how adversaries work, how they move, and the tools they use, to limit the exposure for our clients through the rapid detection of compromises. From a practical standpoint, we continually and proactively hunt for threats on the horizon before they can enter our clients' environments.

Our research and queries are run minute by minute. We're able to detect most types of adversary techniques the moment that they land on an impacted system, or the moment that they start probing one of our clients, which brings us very close to the moment of initial compromise. Often we are able to thwart attacks because we have advanced knowledge of reconnaissance or have observed chatter on forums where the adversaries are chatting and making plans.

By getting in front of adversaries or, in some cases, just a step or two behind them, we're able to limit the overall impact that they have. This process is hugely beneficial for our clients because they can close backdoors and vulnerabilities before they become public knowledge.

Approaching attacks from a behavioral standpoint rather than looking at a discrete piece of malware or command and control traffic is essential. When using a behavioral approach, if the adversary modifies their techniques or their tools in some way, we are often still able to detect them early in the attack lifecycle based on their patterns and habits.

A WELL-COORDINATED TEAM

Threat hunting requires a well-coordinated team. Threat Hunting at BlueVoyant involves several teams of highly focused people working around the clock to gather predictive data from various forums. Each team has a particular role throughout that cycle to ensure data is continuously refreshed and up-to-date.

BlueVoyant's Threat Fusion Cell is a team of threat analysts, embedded within Managed Services, who continuously perform security research, gather new intelligence, and expand our intelligence data. They document emerging threats in regularly produced threat assessment reports outlining new campaigns, malware, and changes in tracked adversary groups. These include the most recent tools and infrastructures that we see. They allow us to generate new threat hunting methodologies and new soft rules. The reports give our clients context to in order to understand threats, events, and incidents that might occur.

The Threat Hunting Team examines behavioral aspects of activity in a client environment in order to identify any irregularity. Unusual activity sets off a chain of action on our end.

A typical chain of action begins with the Threat Hunting Team informing SOC Analysts of the situation. The Threat Fusion Cell is alerted to the new malware, or presence of command and control traffic, or domain. A new rule is implemented and information on triage and how to identify the threat is made available to the SOC. The Threat Fusion Cell incorporates the information into their intelligence gathering and security research. They identify any variants in the wild and search further to determine if there are other examples of work by this author. This proactive detection and prevention mechanism is shared to ensure that if any of these adversaries hit, we can not only see them but in most cases, we can block them.

CHAPTER 4: LIVING IN DATA

Successful threat hunting requires living in data; a threat hunting team spends the majority of their time sifting through data. At BlueVoyant, we take the technologies that clients have deployed and we feed it into a centralized location. There we can normalize it, make sense of it, and then start building automation and various queries against it to identify threat actors. We receive your data unfiltered, straight from your security devices. This process provides us an advantage as the data is raw, and no one else has influenced it by reviewing it before it reaches us.

In addition to your data, we receive threat intelligence feeds curated and generated by the BlueVoyant Threat Intelligence team. We're continually getting bleeding edge data about newly registered domains from adversary groups, or a new sample, or the type of methodology or technique we've seen an adversary use. We analyze the data searching for unusual activity or behavioral anomalies. Once detected, we predict what will happen next. For example, we might know that there are three more steps in this particular attack and we need to go look for all of them.

BUILDING HUNTING METHODOLOGIES

Once we have all of that data coming in and it's enriched with threat intelligence as well as the human context behind it, we start building out hunting methodologies. Our hunting methodologies are behavioral approaches to finding unusual or malicious activities that are at a higher level than you would see in most antivirus solutions or most SIEM queries. Instead of looking for a particular piece of data that's coming in from the threat fusion cells, we'll look at the chain of an attack, determine the adversary, and identify the adversary's usual behaviors.

We may identify an adversary that drops mimikatz. Through our methodology, we know that before dropping mimikatz, this adversary tends to move laterally using a remote desktop protocol, and their initial ingress into environments is often exploiting Apache struts or WordPress followed by a web shell. Rather than look for something specific like the China Chopper web shell, we would abstract that one layer and design a behavioral approach to identifying web shell activity. That way, if they modify the underlying component, we're still able to find them, and we're still able to track them. This methodology also accounts for a great deal of new malware and new techniques that come into play.

A TWO-PRONGED APPROACH

As we find various activities or various pieces of malware in your environment, we take two approaches depending on whether the threat is known or unknown.

If we can identify unusual activity and we have a significant amount of intelligence around it (known-bad); we report it to the SOC. The SOC informs the client while we begin remediation. We use this process for mass malware like Emotet or TrickBot, which are well documented.

If the suspicious activity is something we haven't seen before, and don't know about, we deeply analyze the affected endpoint system and web traffic. We identify any other unusual activity, look for the root cause, the role of the system on which it occurred, and which user executed it. The goal is to get a better understanding of anything else that may have happened prior to or after that suspicious event to ensure we don't miss any attacker activity. We also want to ensure we aren't reporting something that is legitimate in your environment. If we have determined that it's malicious, our next step is to update our detection capability.

Updating our detection capabilities ensures we can detect it immediately the next time. It also ensures that we have processes in place to remediate, notify the client, and mitigate the threat before it happens again.

We also update our threat intelligence teams. We let them know what we found and any context that we can provide to help them expand their scope of research and better understand what that adversary or developer is doing. Lastly, we present the client with a remediation and mitigation plan, which includes:

- How the adversary got in
- The root cause
- Suggestions for future prevention

EXAMPLE: MANAGING A TRICKBOT INFECTION

TrickBot is a malware dropper that's really everywhere. It's all over the internet. We will illustrate what a TrickBot infection looks like when we are in a client environment that:

- Has traditional antivirus
- Doesn't have Next Gen AV
- Doesn't have active, hands-on incident responders
- Doesn't have another detection mechanism that might catch it earlier

We are focusing on the initial delivery mechanism of TrickBots.

Quite frequently, TrickBot comes into your system through a phishing email with a malicious document containing a macro. It is highly successful in environments that don't have sophisticated detection technologies and don't employ SOC analysts or other preventative measures.

TrickBot moves through an environment using common “live off the land” utilities. It embeds applications in the operating system to allow the adversary to execute code or modify the system at elevated privileges. As soon as that macro drops, Wscript executes and runs JavaScript to continue pulling down additional samples, which then go on into a full-blown infection resulting in second-stage malware, such as Ryuk ransomware deployment.

Because TrickBot iterates so quickly and the developers of that malware are paying attention to the defensive measures and antivirus detection rates, they’re constantly developing security bypasses. They’re continually building it better or trying new mechanisms to bypass antivirus, which means more often than not, TrickBot itself isn’t detected, but whatever it drops might be detected.

If you have SOC analysts monitoring and next-generation antivirus deployed, you should immediately catch Trickbot and its associated persistence mechanisms. NextGen antivirus would ultimately block the command and control and all of the other flow, and you would have been alerted by a SOC analyst. However, it still leaves a document on disc that we know has malicious macros; it has an embedded dropper and download in it. These leave you open to a potential compromise in the future.

When we approach this same scenario from a threat hunting perspective, we do things a little differently, and we detect the TrickBot behaviorally almost as soon as it happens.

When the user downloads a phishing document and executes the malicious macro with VB script, though we can’t necessarily behaviorally detect the opening or downloading of a document, we can immediately identify that Microsoft word executed wscript. Although the simple pairing between downloading the document and executing wscript may be benign and may have perfectly normal business use cases, it’s of immediate interest to threat hunters.

When JavaScript reaches out to a network connection on a public IP address, it generates another behavior of interest that we capture. Additionally, we have creation of a new Windows service, we have svchost.exe invoking different binaries, and we have an additional application dropped from what is already at this point in the threat hunting cycle, something that's very suspicious.

We approach these instances with the assumption that all of your other defensive mechanisms may not work. We make this assumption because there are circumstances where antivirus is either disabled, not installed, or misconfigured. There may also be other instances where we don't have full coverage or full visibility into your network.

Our approach is to write different triggers for every behavior that a sample like this may take even though antivirus could catch aspects of this, and Next Generation antivirus could catch even more.

Attackers know that we are looking at various behaviors that an executable may take and that we are proactively blocking them. Borrowing techniques from historically more advanced adversaries, attackers are performing a great deal of reconnaissance. The attackers compromise your environment, and instead of using automated droppers and drive-by downloads, they are embedding applications by hand. They dump credentials, move laterally, identify backup servers, identify valuable data, and instead of stealing it, they take the time to disable the detection and prevention measures you have put in place and then they compromise your backup servers first.

This allows your attacker to look like a normal user in the environment. They take their time; they stage their ransomware everywhere; they deploy the ransomware, and they wait for the money.

Our assumption of lack of detection and our assumption that hands-on keyboard activity is going to keep happening, is the value and the safety net that comes with threat hunting services. We're looking for users doing things that they've never done before, like accessing a brand-new system or installing a new service. As we gather all of that data and review it, we profile that user and identify if there's something unusual.

A DEFENDER'S DUTY

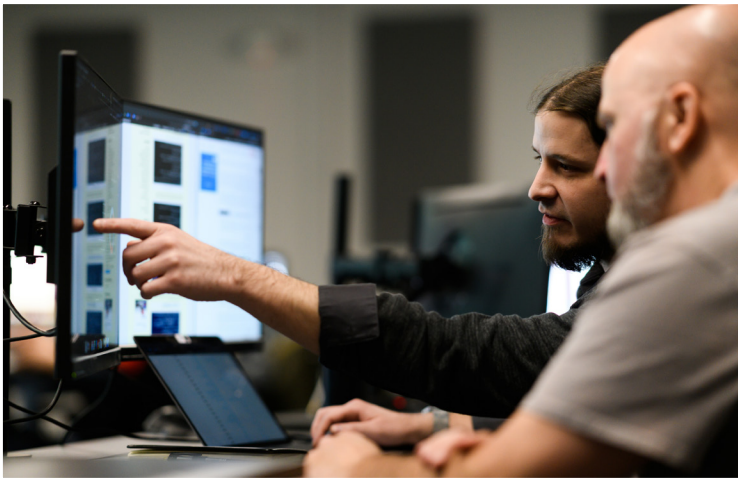

As defenders, we must remain diligent. Keeping up with the quickly changing landscape is a challenge, but not impossible. Defenders must constantly research information from the larger security industry, review new samples and new tools, and evaluate new security tools. Defenders must understand how an adversary can misuse and abuse the new tools, and then proactively look for this activity.

You need to match the sophistication of your adversary, not necessarily in your ability to deploy code, but in understanding what they do as a business and what they look like as they walk through an attack. You must weed out benign activity from something potentially malicious. You also need to continually look for tools that adversaries write themselves, new malware variants, and also tools used by system administrators, for example, PsExec or TeamViewer, that may be abused by your adversaries.

You must review updates to offensive security tools. The tools that are built and designed for testing security frameworks, penetration testing, or red teaming are often co-opted by adversaries and then modified to some degree and ultimately used to compromise your environment. Profiling your adversary's infrastructure and command and control servers, understanding which forums they frequent, knowing what their domain generation algorithm looks like, and using that information to query raw data gets you ahead of your attacker. This data allows you to identify new command and controls because you know your adversaries' protocols.

CREATING THE PLAYBOOK

Through an extremely high-level threat hunting process at BlueVoyant, we have gathered and analyzed a great deal of data. We not only present this information to our clients; we also create repeatable processes. We use new information and new threats to create policies and playbooks.

Playbooks are the built-in automations that allow us to take a small event and expand it through various API and intelligence lookups, and various contexts from other sources, to provide our SOC analysts with a narrative indicating exactly what happened. The playbooks gather all of that data, score it, put it in front of our SOC analysts along with the triage steps necessary to determine whether this was a malicious activity, something benign, or perhaps a new variant.

Our policies and automated playbooks are maintained on our Security Orchestration Automation and Response (SOAR) platform. They take us closer to bridging the gap from detection, remediation, and notification - from potentially hours to an average of three minutes. Those three minutes include the entire playbook, triage, notification, and the start of remediation.

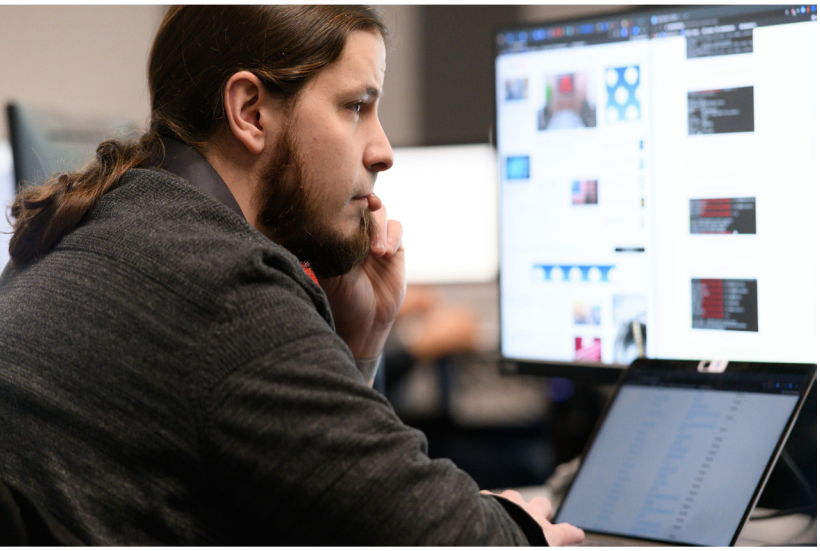

A SOAR platform allows us to take new and emerging threats, make them digestible, and essentially old overnight. The SOAR platform assists us in democratizing detection for you.

There are thousands of people performing threat hunting and SOC analysis. The people have the experience and expertise necessary to perform the analysis of threats, but people don't scale. By using the SOAR platform, we empower our experience across all of our clients without having false negatives or dropping data. The SOAR platform is custom-designed to integrate with the data provided.

CONCLUSION

Many organizations take newly reported information, put it in their threat intelligence system, and start searching to capture any instance of that one particular thing happening. This is a reactive approach to threat hunting. While it's a valid way to try to get ahead of your attackers, it does not account for the variability of the human element of adversaries. It does not account for the possible actions of adversaries on a larger scale. Your adversary may exhibit behavior that closely resembles a typical user in your environment, leaving you at risk.

The value of threat hunting is to detect emerging threats. The goal is to get ahead of the adversary as much as possible. Predict what the adversary is going to do next by understanding the attack lifecycle. Be proactive.

The ultimate value of threat hunting is enhanced protection. It allows you to track your adversary before your vendors, Next-Gen AV, or Next-Gen firewall. Your adversaries are sophisticated and constantly evolving; you must stay ahead of them.

Learn more at www.bluevoyant.com and follow us on LinkedIn for news and industry insights. Have questions? Email us at contact@bluevoyant.com.