INTEL471 CYBORG SECURITY

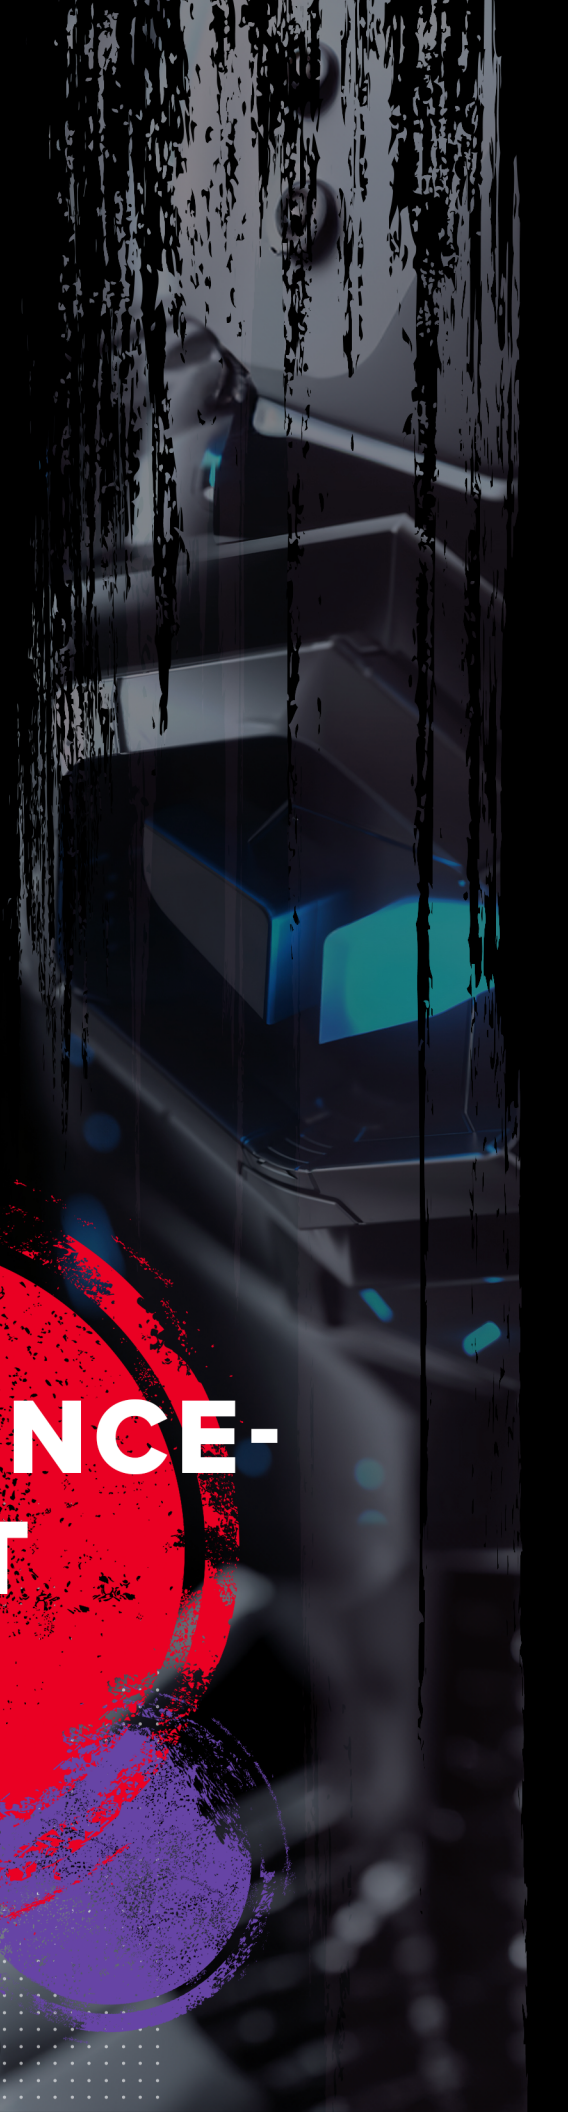

WHITEPAPER

SUPERCHARGE YOUR SECURITY WITH **INTELLIGENCE-DRIVEN THREAT HUNTING**

Security teams are faced with a reality: sometimes, adversaries are going to compromise an environment. A user may click on a link in a phishing email that leads to the download of malware that's not caught by antivirus software. A threat actor may exploit an unpatched vulnerability in an internet-facing appliance that was not on an organization's asset register. Compromised credentials could lead to an attacker taking over a highly privileged account, lending access to a domain controller. However, not all is immediately lost. Once an adversary gains access, there may be time before attackers begin to move laterally and the threat can be detected and removed. This dwell time can vary in length, but this gap offers an opportunity for threat hunting.

Intel 471's HUNTER is a powerful threat-hunting platform that is designed to let security teams proactively detect these kinds of cyber threats before they manifest into a serious event. The platform contains a suite of tools for managing and executing threat hunts. At the heart of the platform are hunt packages. These are pre-written threat hunt queries that allow security operations teams to query their SIEMs, XDR, EDR, Splunk or other logging platforms or security tools to hunt for suspicious activity that could indicate a compromise. The queries empower security teams to jump immediately into threat hunting and spend less time crafting queries and more on investigations and remediation. This provides scale and efficiency, allowing for more hunts. The queries are derived from Intel 471's leading intelligence, which collects real-time malware indicators and the tactics, techniques and procedures (TTPs) of threat actors. This whitepaper will discuss some of the concepts around threat hunting, how threat hunts can be conducted and how an intelligence-driven behavioral threat hunting program can lead to positive security outcomes.

Why Threat Hunt?

The premise of threat hunting is to look for clues of a possible compromise. Security tools have long employed indicators of compromise (IoCs) as a means of detecting attacks. IoCs include hashes of malware samples, malicious domains associated with malware downloads and IP addresses. These types of indicators, however, are usually short-lived. Adversaries know these types of files and infrastructure will be quickly flagged and blocked, so they change them frequently. It's trivial for malware developers to change a few bytes of the code to create a piece of malware with a different hash that has not yet been flagged as malicious. Malware distributors routinely change domains and hosting URLs to new ones that have not been linked to malicious behavior. It's much more difficult, however, for threat actors to change behaviors linked to intrusions. These TTPs are at the top of the Pyramid of Pain, which is a model of indicators illustrating the relative difficulty threat actors have in changing them.

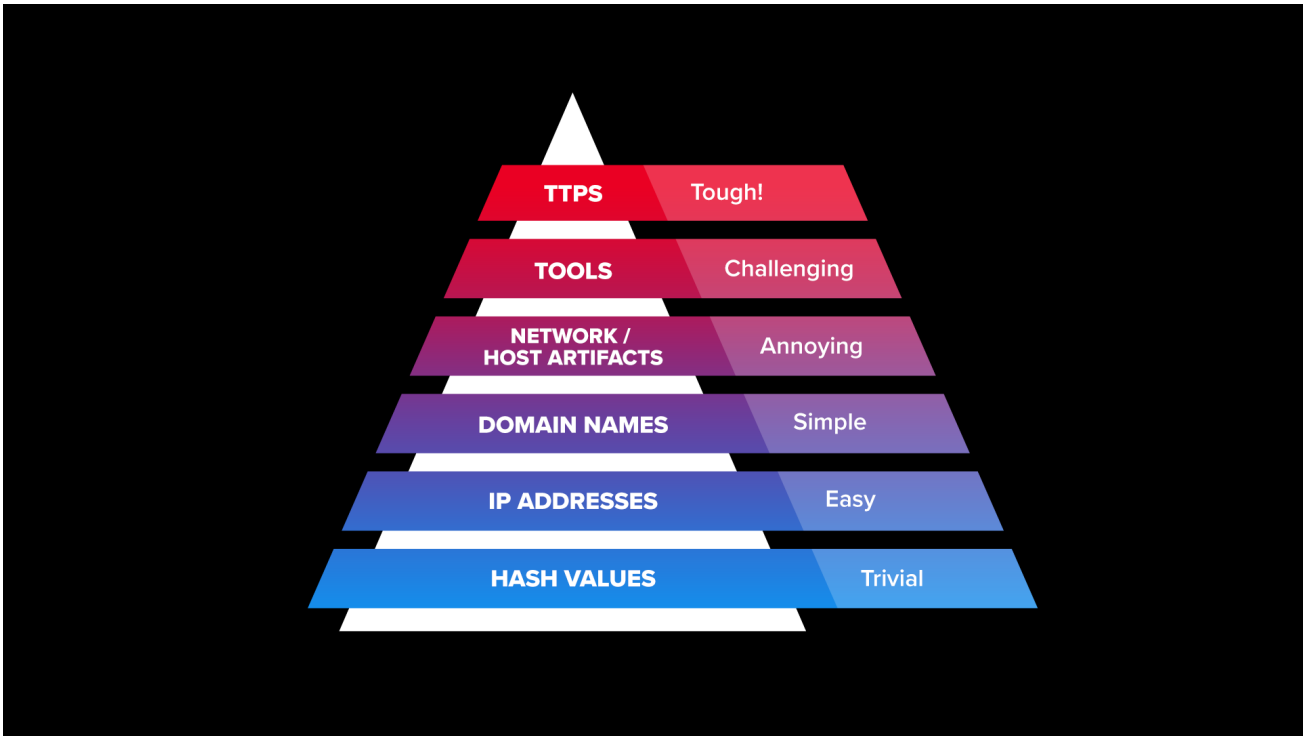

The Pyramid of Pain, which is a model of indicators of compromise ranked by the difficulty threat actors have in changing them.

Malicious Behaviors

Focusing on the peak of the Pyramid of Pain offers a more reliable way of detecting intrusions. Threat actors and malware often perform consistent actions on newly compromised machines. If these behaviors are found on a system, it's a starting point for further investigation. TTPs of threat actor groups are discovered through incident response investigations and the analysis of malware and infrastructure. TTPs are distributed through a variety of ways, from national cybersecurity agencies, published vendor reports, closed trust groups, blog posts, social media and information sharing and analysis centers and organizations (ISACs and ISAOs).

As an example, one behavior-based TTP used by many threat actor groups is disabling security software. On Windows systems, ransomware actors will routinely try to disable Microsoft Defender prior to deploying file-encrypting malware. They do this by manipulating the registry keys, such as the DisableAntiSpyware key, which is done by using Powershell commandlets to change its value from a 0 to a 1. If discovered using a behavior-based threat hunt, this type of action would merit a deeper look as it's somewhat unlikely to have been undertaken by an administrator. The Hunter platform contains a query for this

behavior. The query is written for Microsoft's Defender and Sentinel as well as other SIEMs such as Elastic, QRadar Query and Splunk that ingest Windows Event Logs.

In another example, ransomware actors also often try to delete backups. The Volume Shadow Copy Service is a framework provided in Microsoft Windows to perform volume backups or for creating consistent, point-in-time copies of data (known as shadow copies). These snapshots are targeted by ransomware actors prior to launching encrypting malware, as it could make it more difficult for a target to recover and thus make it more likely to consider paying a ransom. This behavior has been used by many ransomware groups including REvil, Ryuk, Maze, CL0P, Netwalker, CryptoJoker, Snatch and Phobos. This tactic is typically carried out with tools such as Powershell, WMI command-line (WMIC) and other management tools for shadow copies, such as vssadmin and vssvc. However, behavioral threat hunting can uncover malicious command line arguments that have caused backups to be deleted or changed, which could be a sign of a pending ransomware attack. HUNTER contains custom hunt queries to find this behavior in CarbonBlack Cloud — Investigate, CarbonBlack Response, CrowdStrikes, CrowdStrike LogScale, Elastic, Elastic Signal, Microsoft Defender and Sentinel, Palo Alto Cortex XDR, QRadar Query, SentinelOne, Splunk and Trend Micro Vision One.

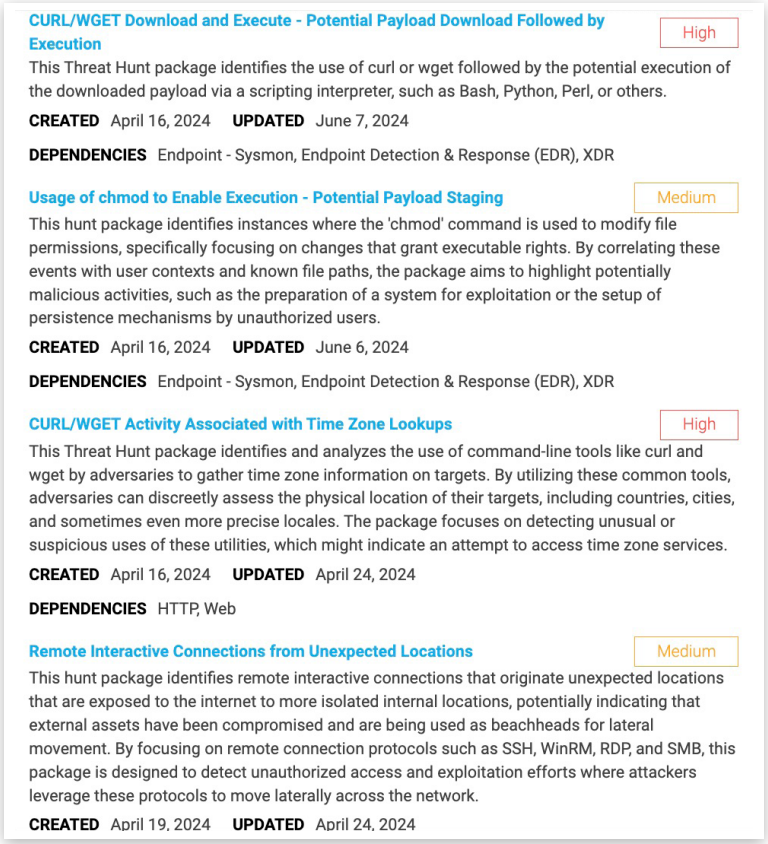

Four threat hunts that are intended to find possible exploitation of CVE-2024-3400, a command injection vulnerability in PAN-OS.

It's also possible to hunt for signs of vulnerability exploitation. If the exploitation behavior is recorded by a logging platform, that can be hunted. Take CVE-2024-3400, which is a command injection vulnerability in certain versions of Palo Alto Networks' PAN-OS, which runs in GlobalProtect, a secure remote access and firewall product. This critical vulnerability ranked a 10 on the Common Vulnerability Scoring System 3.1 and was added to the Cybersecurity and Infrastructure Security Agency's (CISA) Known Exploited Vulnerabilities catalog on April 12, 2024. After this vulnerability became public, we created four hunt packages with queries to help assess if users of that product had been exploited.

What to Hunt: Behavior-Based Threat Hunting

The cyber threat landscape is enormous. Threat hunters have limited time and want to ensure that hunts are turned to maximize the chances of discovering potential threats. How should threat hunters decide what potentially threatening behavior to hunt? It's one of the hardest questions a security team can ask, but should start with cyber threat intelligence (CTI).

Organizations ingesting CTI should undertake an intelligence planning exercise. This exercise is intended to gain a greater understanding of what key stakeholders — ranging from senior management, security operations, incident response, forensics, legal and risk management professionals — regard as the most important or valuable types of cyber threat intelligence. Intel 471 has developed an open-source framework, the Cyber Underground General Intelligence Requirements Handbook (CU-GIRH) as a baseline tool to assist in organizing, prioritizing, measuring and producing cyber underground intelligence (for more information, see our blog post "Open Source Release of Intel 471 Intelligence Requirements Framework)." Central to the CU-GIRH framework are General Intelligence Requirements (GIRs). GIRs describe threats and activities that pose risks to organizations — such as malware, vulnerabilities, access brokering, etc. — and the relevant questions around those activities that practitioners should focus on to create actionable intelligence products. These GIRs can be selected to narrow down Priority Intelligence Requirements (PIRs), which are specific to organizations based on their own threat modeling and assessments. PIRs are the most important intelligence requirements for an organization.

PIRs can be a great starting point for threat hunting. Say, an organization's top three PIRs are:

- 1 Malware > 1.1 Malware variants >1.1.6 Loader malware
- 1 Malware > 1.1 Malware variants >1.1.5 Information stealing malware
- 5 Adversary Tactics and Activities > 5.2 Post-attack tactics > 5.2.8 Lateral movement tactic

This can be used to guide threat hunts. For example, threat actors use so-called living-off-the-land binaries (LOLbins) and tools once they have compromised a system. LOLbins, such as Microsoft's Powershell scripting language, are often used by attackers because these tools are already on a system and may not raise security alarms. Malicious actors may use PowerShell to gather system information, perform network scanning, or facilitate lateral movement within compromised networks. They also may use specific PowerShell modules, such as Invoke-Command, to execute commands on remote systems. If threat intelligence shows that a ransomware group has been recently invoking PowerShell in a unique way, that would be a prime threat hunt that aligns with an organization's PIRs, as it would fall under the lateral movement GIR. HUNTER contains many queries related to malicious use of Powershell. For example, the hunt package below identifies PowerShell processes started with parameters meant to modify the execution policy of the run, run in a hidden window and connect to the Internet. This combination of command-line options is suspicious because it's overriding the default PowerShell execution policy, attempts to hide its activity from the user, and connects to the Internet. This behavior has been used by the BlackByte, Nokoyawa and Rapture ransomware groups and has been seen related to XMRig, a cryptominer.

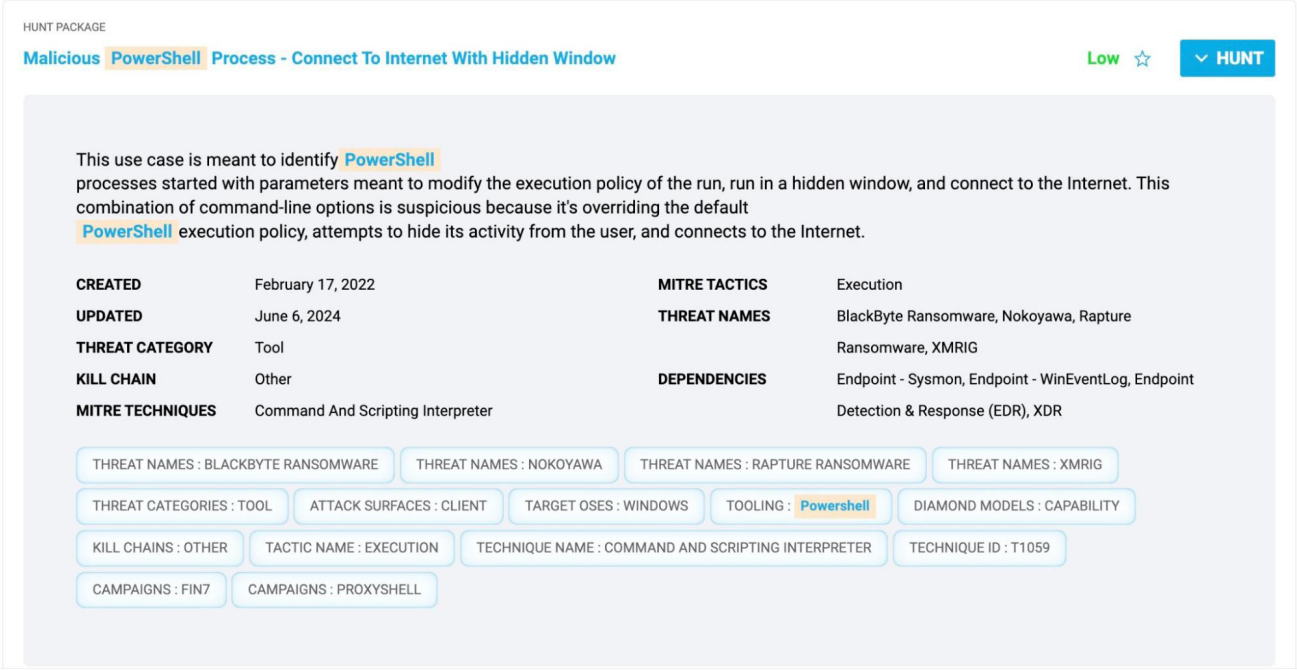

A Hunter package that looks for malicious use of Powershell that attempts to hide its use and connect to the internet.

Intelligence-Driven Threat Hunting

Every organization faces threats from coordinated malware distribution campaigns. These campaigns are pervasive, persistent and come from a variety of vectors. Spear-phishing emails, which seek to intentionally target specific people in an organization, may contain malicious attachments or links. Malicious actors manipulate search rankings using search-engine optimization (SEO) techniques, tricking people into visiting fraudulent websites that purport to contain legitimate software but actually are malware. Malvertising is another method in which threat actors buy search-engine or display advertisements, with those advertisements leading to deceptive sites hosting malware.

All of these methods consistently result in infecting machines with malware, from infostealers to "loaders," which are code used to load other harmful software. Botnets run by cybercriminal groups seek to infect computers and then sell that access to other cybercriminals, an arrangement known as initial access brokering. Security software often misses malware because indicators such as hash values and command-and-control services used to communicate with the malware frequently change.

SECURITY SOFTWARE OFTEN MISSES MALWARE BECAUSE INDICATORS FREQUENTLY CHANGE

Malware can also be detected based on its behavior. We periodically publish in-depth reports on malware campaigns that pose a wide threat. These reports can help organizations understand real-time threats and focus threat hunts. In May 2024, we published a malware campaign report about Gootloader, which is a type of loader or initial stage malware that loads other malware. GootLoader remains a significant threat within the cybersecurity domain and has maintained its core operational tactics without significant changes for years. Gootloader has been used to deploy tools for reconnaissance and lateral movement as well as deliver malware such as the REvil ransomware, the Gootkit banking trojan, the Kronos banking trojan and Cobalt Strike, a post-exploitation attack framework. Gootloader usually consists of a loader and core component. The core component allows for deployment of reconnaissance and lateral movement tools, including well-known utilities such as Rubeus, SharpHound and SystemBC.

Once it infects a machine, Gootloader performs many behaviors. One of those is creating suspicious scheduled tasks, a tactic used by malware for persistence. While scheduled tasks often have legitimate purposes, adversaries have been known to create and use scheduled tasks to execute code, escalate privileges, move laterally or establish persistence. In the Gootloader sample we observed, it created scheduled tasks that invoked JavaScript files. As a result, we created a threat hunt package for our Hunter platform to look for three types of suspicious scheduled tasks:

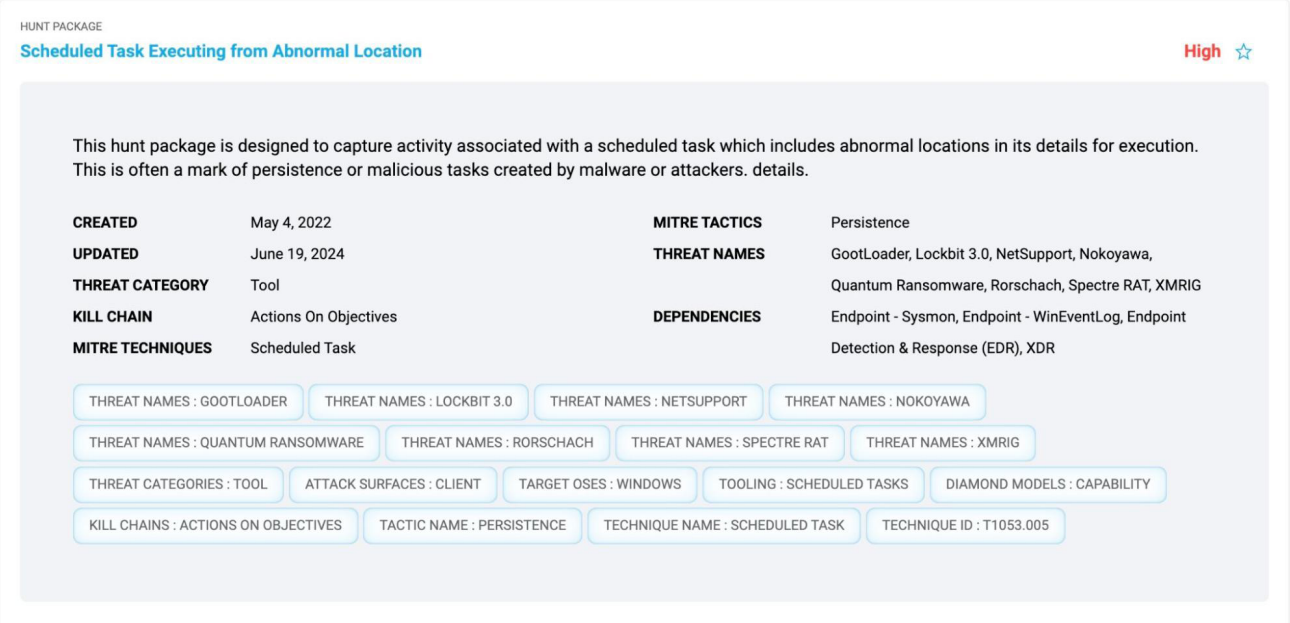

Screenshot of a hunt package that is crafted to detect suspicious scheduled tasks.

Next Steps

Organizations of all sizes can undertake intelligence-driven threat hunting. While large organizations have more resources and staffing for their security teams, that should not discourage smaller IT or IT security teams from trying threat hunting. With a focused scope, threat hunting can result in a significant return on investment, particularly when compared with the costs of a data breach or ransomware event. For more information on threat hunting and how to set up an intelligence-driven threat hunting practice, please contact Intel 471.