The art and science behind Microsoft threat hunting: Part 1

microsoft.com/security/blog/2022/09/08/part-1-the-art-and-science-of-threat-hunting/

September 8, 2022

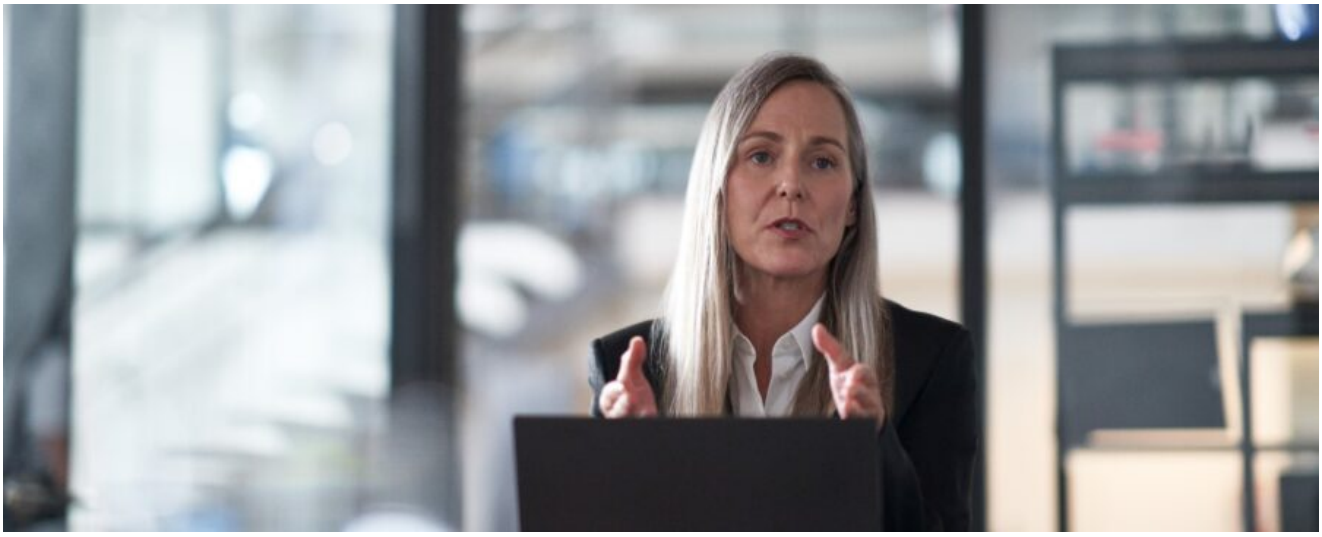

At Microsoft, we define threat hunting as the practice of actively looking for cyberthreats that have covertly (or not so covertly) penetrated an environment. This involves looking beyond the known alerts or malicious threats to discover new potential threats and vulnerabilities.

Why do incident responders hunt?

The Microsoft Detection and Response Team (DART) mission is to respond to security incidents and help our customers become cyber-resilient. This involves incorporating threat hunting as part of our proactive and reactive investigative service offerings to determine the following:

- Whether systems are under targeted exploitation through investigation for signs of advanced implants and anomalous behavior.
- Identifying groundwork for the recovery process of evicting the attacker from the environment.
- Strategic recommendations for protecting against sophisticated threat actors.

In reactive incident response investigations, threat hunting helps determine the full scope of the incident and informs an effective recovery and remediation strategy. In proactive investigations, a threat hunt can discover latent threats or existing compromises as well as demonstrate the effectiveness of current security controls and their security operations

processes. By uncovering novel attacker campaigns and previously undetected threats, DART provides valuable feedback to improve product detections, both for Microsoft security products and for the entire security ecosystem.

How do we approach threat hunting?

The canonical definition of threat hunting involves three interrelated things:

- **Targeted threat hunting**—We define targeted hunting as actively looking for and rooting out cyberthreats that have penetrated an environment, and looking beyond the known alerts or malicious threats to discover new potential threats and vulnerabilities. Targeted threat hunting has a scope where we are looking for specific classes of indicators. For example, given a recently revealed attack, an organization may want to assess its environment to see if it, too, has been affected.
- **Security monitoring**—Process of continuously monitoring the state of an environment to detect unusual or unauthorized activities. This involves a network operations center (NOC) and an SOC to ensure that networks are protected against disruptions and threats.
- **Incident response investigation—**An investigation to identify the root cause and develop a remediation plan to regain and retain positive control over the environment following the detection of unauthorized access or suspicious activity.

Each organization approaches threat hunting differently. Sometimes, the customer will have specific outcomes in mind that align with the known techniques. We center on a general approach based on anomaly detection and pivoting combined with a knowledge of the overall environment. This allows us to accomplish multiple goals, versus employing an approach solely focused on a targeted threat hunt where additional threats and risks may be overlooked.

We will go into more detail about hunting for anomalies later in this blog.

Threat hunting principles

Our forensic investigators at DART lean on the Alexiou Principle, which states four key questions for our investigators to answer:

1. What question are you trying to answer?

Threat hunting varies depending on the main objectives or questions that need to be answered. This involves trying to understand a threat actor's main objective, the cyber terrain in which they operate, and understanding how you can get closer to those objectives. Framing the question clearly helps us define the scope of every threat hunt.

2. What data do you need to answer that question?

To answer the previous question would involve a two-pronged approach with a focus on determining what data is required, and how to obtain that data. During DART investigations, we often get a variety of datasets while entering a customer investigation, such as live feeds and telemetry. We want to pick up everything that is currently in the environment, enumerate directories that we know bad actors like to live in, collect event logs that will potentially show us evidence of historical or current badness, registry keys that we see bad actors like to tamper with, and many more.

We use a tiered data collection model and start by collecting a snapshot of the densest, most indicator-rich data we can from every object and endpoint we can reach. This data is intended to provide information about any known threats, known attack patterns, and many (but not all) indicators of suspicious or anomalous activity. Where systems of interest are identified, we return and collect a larger, more complete dataset of logs and forensic artifacts.

3. How do you extract that data?

Now that you've identified the data, you'll need to capture it using various toolsets, such as a point-in-time snapshot tool or, if the customer doesn't have one deployed already, an endpoint detection and response tool, such as Microsoft Defender for Endpoint, to obtain the data. From the analytics captured, we can see things that are potentially good, bad, or interesting. Part of this phase also takes data ingestion into account. We consider how the collected data is consumed and how to efficiently separate threats from the background noise of a complex global enterprise.

4. What does the data tell you?

Looking at the collected data now becomes an exercise in data analytics. It's a question of evaluating prevalence and frequency by taking everything that occurs within an environment and trying to figure out what belongs and what doesn't belong. This train of thought can take a handful of different forms, that can be something as simple as "*How often does this secure hashing algorithm show up across the entire environment?"* to a more nuanced and precise way, such as asking "*How often does it show up only on domain controllers? On devices in this organizational unit? How about when it's seen with this other user account?"*

As it turns out, there are a lot of different ways for us to do this counting game. Our role as threat hunters is to figure out the most relevant, high-priority way to account for these interesting findings and see if patterns revealed themselves. We're looking for indicators of attack or compromise that maybe others haven't found. It all depends on what data is available to us and understanding it.

Understanding the data

We approach understanding the data by looking for anomalies, the current state, and the absence of data.

Where the rubber meets the road: forming the attack narrative

We believe there's a clear art and science to threat hunting, but at the end of the day, we seek to understand the anomalies in the acquired evidence. One way we do this is by using the knowledge of what is typical in an environment to identify what isn't. Understanding the typical scenario and marrying that with the knowledge of threat actor tools, techniques, and processes allows us to gain a deep understanding of the data and the systems we're looking at. Stringing these anomalies together can then create a pattern of anomalies, helping us form a story using analytical opinions based on facts, also known as the attack narrative.

The ability to identify anomalies makes for an important skill set for an analyst, but understanding the current state is just as crucial. Anomaly-based hunting will be discussed in more detail in the second part of this series when we go into general hunting strategies.

Looking at the current state

If an investigator is lucky enough, they might be dealing with forensic data for the anomaly hunt. But often, there will be times when our observations are limited to the current state of the environment. Even if we don't have the luxury of historical artifacts, looking at the current state can provide valuable information.

Our proactive Cybersecurity Operations Services prior to an incident allow organizations to gain better knowledge of their current security posture and risk exposure before an incident even occurs.

By understanding the current state and its configurations, you can determine where the potentially malicious or anomalous activity lies as an initial starting point.

Asking questions like *"How did it get into that state? Was that it in that state intentionally or was that the result of somebody doing something malicious?"* allows our investigators to build from something of interest, look a little bit closer, and then pivot from there until we find true signs of malicious activity.

Looking at the absence of data

The absence of data is just as important as understanding the presence of it. Often, we are provided with data that is lacking or missing, and so the questions gleaned from these observations become: *"Why don't I have that artifact(s)? What didn't happen? Was it because this data wasn't recorded? Was the data removed?"*

In the absence of data, we also try to determine what could have happened at a given stage of a compromise and what normally happens at that stage. With that information, we try and form our hypothesis about the stages of compromise, if it occurs in an environment. For instance, a customer during an incident response engagement might halt further investigation or response simply because they're not seeing data exfiltration activity on their sensors and logs.

The approach to understanding the data varies depending on the analyst, but the goal is to answer a series of questions and turn those questions into more questions, and then stop at some point so you can paint the most complete picture possible.

Knowing when to stop

Following an investigative trail results in some form of data aggregation. Knowing when to stop this trail can often be challenging. An indication of knowing when to stop is when the picture doesn't change even after pulling in more information, leaving you with a nexus of truth about that event or indicator. A comparison to this is the computer science algorithm of depth-first search versus breadth-first search, where investigators can potentially chase one single trail too far, investing too much time on one possible indicator of an attack, and running out of time to investigate other possible indicators. One approach we take to avoid the pitfalls of digging for data is to consult with fellow analysts to get a different perspective to ensure that you are looking at everything from every possible angle. Weighted risk analysis also helps us narrow down what leads to follow. We ask ourselves "what is the probability that a lead I'm investigating will turn out to be malicious?" Multiply that by the potential impact that malicious activity would have. Using that value to rank which leads are most important to follow first helps find higher-risk threats (ransomware, full-domain compromises) faster than low-risk threats (adware, coin miners).

We've just described DART's threat hunting principles and the art form that is understanding the data we're dealing with when it comes to our incident response work, combing through the data, and creating patterns of suspicious activity by applying critical thinking. In our follow-up post, we will talk about general strategies behind threat hunting and how we work with threat intelligence. Stay tuned.

Learn more

Go to our DART blog series to learn more about the Microsoft Detection and Response Team.

To learn more about Microsoft Security solutions, visit our website. Bookmark the Security blog to keep up with our expert coverage on security matters. Also, follow us at @MSFTSecurity for the latest news and updates on cybersecurity.