The art and science behind Microsoft threat hunting: Part 2

microsoft.com/security/blog/2022/09/21/the-art-and-science-behind-microsoft-threat-hunting-part-2/

September 21, 2022

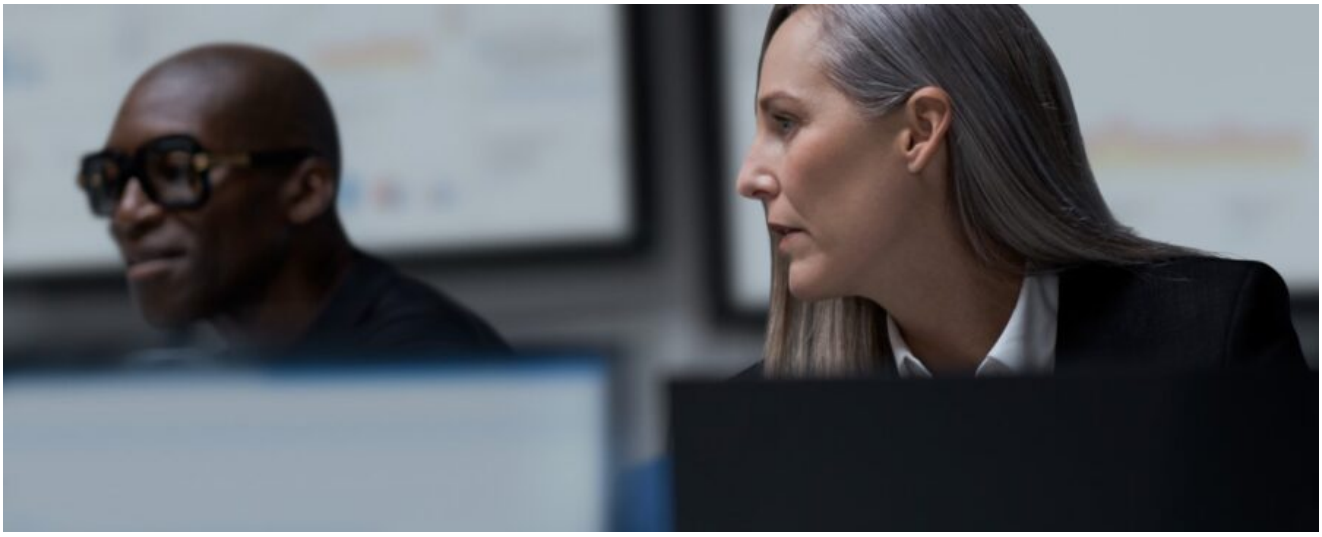

We discussed Microsoft Detection and Response Team's (DART) threat hunting principles in part 1 of The art and science behind Microsoft threat hunting blog series. In this follow-up post, we will talk about some general hunting strategies, frameworks, tools, and how Microsoft incident responders work with threat intelligence.

General hunting strategies

In DART, we follow a set of threat hunting strategies when our analysts start their investigations. These strategies serve as catalysts for our analysts to conduct deeper investigations. For the purposes of this blog, we are listing these strategies under the assumption that a compromise has been confirmed in the customer's environment.

Starting with IOCs ("known bads")

An incident response investigation is more manageable when you start off with an initial indicator of compromise (IOC) trigger, or a "known bad," to take you to any additional findings. We typically begin with data reduction techniques to limit the data we're looking at. One example is data stacking, which helps us filter and sort out forensic artifacts by indicator across the enterprise environment until we've determined that several machines across the same environment have been confirmed with that same IOC trigger. We then enter the hunting flow and rinse and repeat this process.

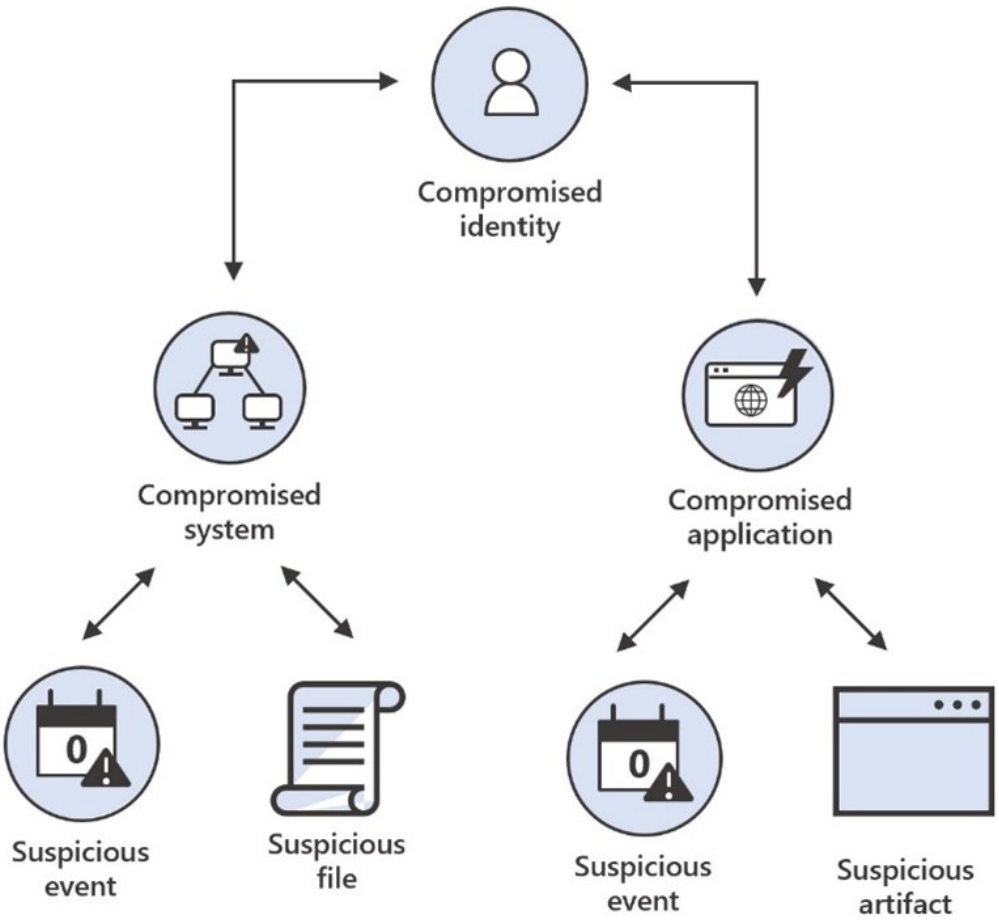

Figure 1: The hunting cycle starts with hunting for indicators or "known bads," ranging from the smallest unit of indicators to behavioral indicators that may define the actor.

Types of indicators can be classified into:

- **Atomic**—The smallest unit. For example, IP addresses, domain names, email addresses, and file names.
- **Computed**—Match multiple atomic indicators. For example, hashes and regular expressions.
- **Behavioral**—Patterns of adversary actions. For example, tactics, techniques, and procedures (TTPs) and demonstrated actor preferences (such as file paths, usernames, and tools).

Quick wins

Unfortunately, not everything we start out with is interrelated with the trigger IOC. Another hunting strategy we employ is to look for quick wins; in other words, looking for indicators of typical adversary behavior present in a customer environment. Some examples of quick wins

include typical actor techniques, actor specific TTPs, known threats, and verified IOCs. Identifying our quick wins is the most impactful to the customer, as it helps us formulate our attack narrative while guiding customers to keep the actor away from the environment.

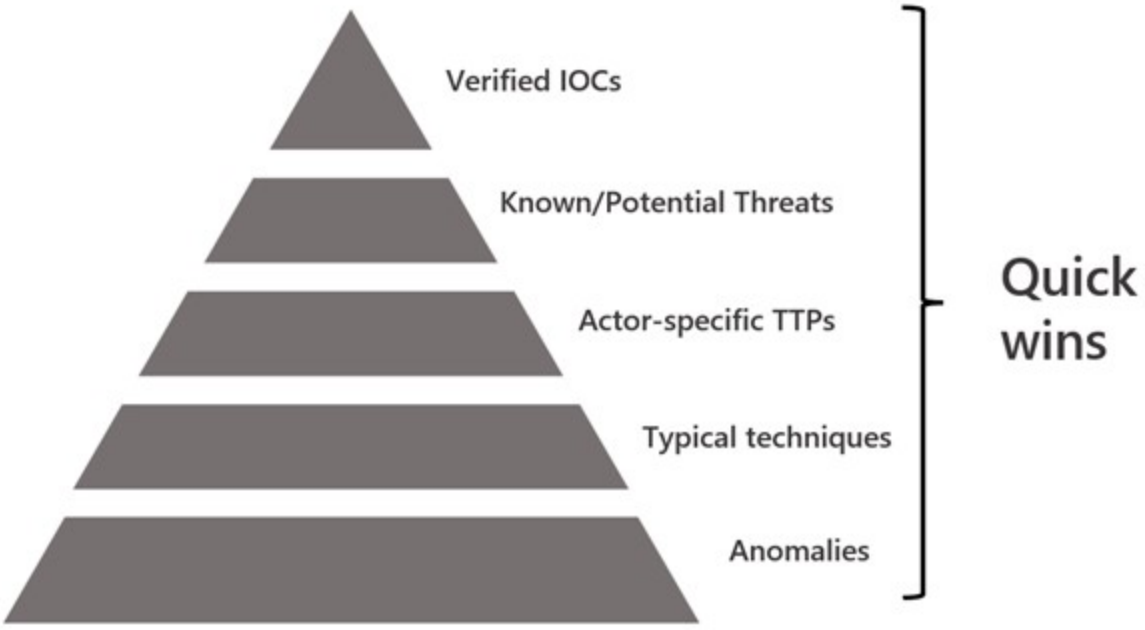

Figure 2: Hunting order of operations.

Anomaly-based hunting

If you're out of leads, another strategy to employ is pivoting to hunting for anomalies, which draws on information derived from our "known bads" and quick wins. We discussed anomalies in the first part of this series as part of understanding the customer data. Some techniques:

- **Define baselines**. Perform baseline comparisons for your dataset. Determine the usual versus the unusual in an environment.
- **Summarize your data and occurrences** and sort by indicator to find the outliers.
- **Clean the output**. We recommend formatting your data in favor of efficiency and accuracy to make the outliers and anomalies stand out.

Pure anomaly-based hunting may be performed concurrently with other hunting strategies on a customer engagement, depending on the data we're presented. This method is incredibly nuanced and requires seasoned experts to verify whether data patterns may encompass normal or "abnormal" behavior. This prevalence checking and data science approach is the most time consuming but can bear some of the most interesting evidence in an investigation. Case in point, we can detect new advanced persistent threat (APT) actor groups and campaigns with anomaly hunting, while they are rarely detected just by searching for the "known bads."

Tying it all together: The attack narrative

Stringing together our patterns of anomalous activity, factual data from quick wins, and analytical opinions must conclude with an attack narrative. In an incident response investigation, the MITRE ATT&CK framework serves as a foundation for adversary tactics and techniques based on real-world observations.

The MITRE framework helps us ensure that that we're looking at our hypothesis in a structured manner to enable us to tell a cohesive narrative to the customer that is rooted in our analysis. We aim to answer questions such as:

- How did the attacker gain access?
- What did they do once inside?
- Which accounts did they use?
- Which systems did they access?
- How and where did they persist?
- Was any data accessed or exfiltrated?
- And, most importantly, are they still in the environment?

Additionally, we want to answer questions surrounding threat actor intent to help tell a better story and build better defenses. Some common attack patterns from the MITRE framework are listed in Table 1.

Tactic	Techniques
Initial access	Phishing files
Execution	PowerShell and service execution
Persistence	Services installations, webshells, scheduled tasks, registry run keys
Defense evasion	Masquerading, obfuscation, Background Intelligent Transfer Service (BITS) jobs, signed executables
Credential access	Brute forcing, credential dumping
Discovery	Network share enumeration
Lateral movement	Overpass the hash, WinRM
Collection	Data staging
Command and control	Un/commonly used ports
Exfiltration	Data compression

Tactic	Techniques
Impact	Data encrypted for impact

Table 1: Common attack patterns from MITRE.

Threat hunting tools and methodology

To ensure maximum visibility of the attack chain, hunters use data sourced from proprietary incident response tooling for **point-in-time deep scanning** on endpoints, as well as bespoke forensic triage tools on devices of interest.

For point-in-time deep scanning, DART uses:

- Proprietary incident response tooling for Windows and Linux.
- Forensic triage tool on devices of interest.
- Microsoft Azure Active Directory (Azure AD) security and configuration assessment.

For continuous monitoring:

- Microsoft Sentinel—Provides centralized source of event logging. Uses machine learning and artificial intelligence.
- Microsoft Defender for Endpoint—For behavioral, process-level detection. Uses machine learning and artificial intelligence to quickly respond to threats while working side-by-side with third-party antivirus vendors.
- Microsoft Defender for Identity—For detection of common threats and analysis of authentication requests. It examines authentication requests to Azure AD from all operating systems and uses machine learning and artificial intelligence to quickly report many types of threats, such as pass-the-hash, golden and silver ticket, skeleton key, and many more.
- Microsoft Defender for Cloud Apps—Cloud Access Security Broker (CASB) that supports various deployment modes including log collection, API connectors, and reverse proxy. It provides rich visibility, control over data travel, and sophisticated analytics to identify and combat cyberthreats across all your Microsoft and third-party cloud services.

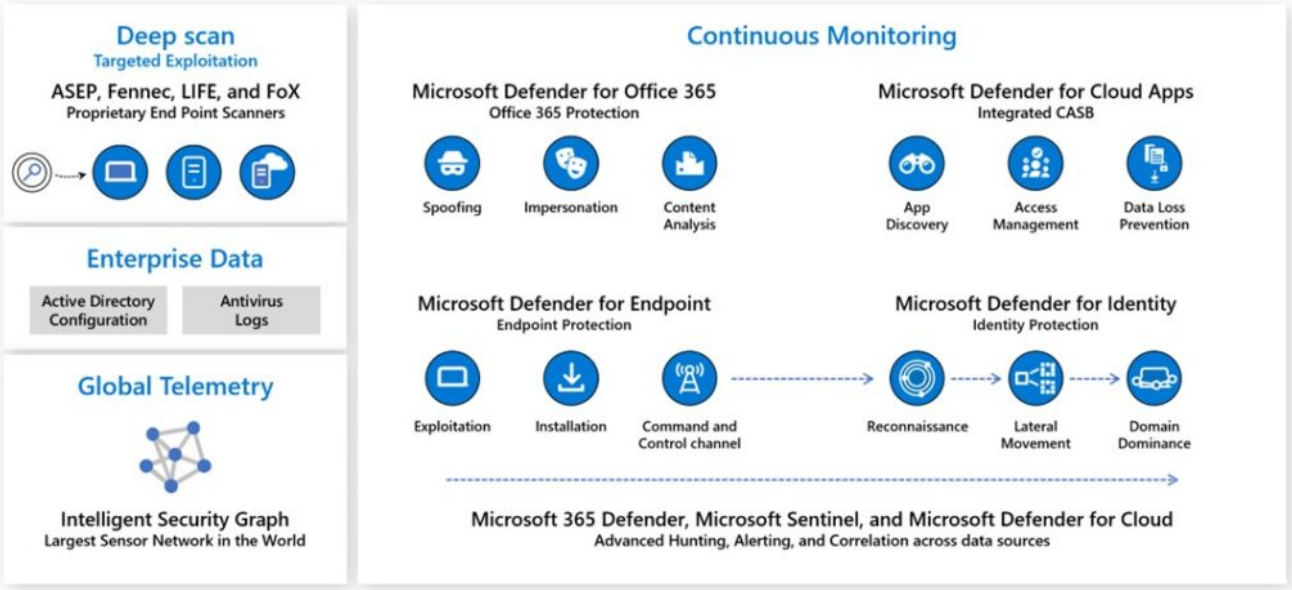

Figure 3 explains the products and services used to identify and monitor threats:

- Deep scan includes proprietary endpoint scanners such as ASEP, Fennec, LIFE, and FoX
- Enterprise data includes Active Directory Configuration and Antivirus logs.
- Global telemetry includes the Intelligent Security Graph, the largest sensor network in the world.

Continuous monitoring includes the following:

- Microsoft Defender for Office 365, which monitors spoofing impersonation, and content analysis.
- Microsoft Defender for Cloud Apps, which monitors app discovery, access management, and data loss prevention.
- Microsoft Defender for Endpoint, which monitors exploitation, installation, and command and control channel.
- Microsoft Defender for Identity, which monitors reconnaissance, lateral movement, and domain dominance.
- Microsoft 365 Defender, Microsoft Sentinel, and Microsoft Defender for Cloud, which include advanced hunting, alerting, and correlation across data sources.

In addition, we work with internal threat intelligence teams, like the Microsoft Threat Intelligence Center (MSTIC), to provide details from our hands-on experience with customer environments and going toe-to-toe with the threat actors. The information collected from these experiences, in turn, provides a trail of evidence to help threat teams and services conduct enriched threat intelligence and security analytics to ensure the security of our customers.

Contributing to threat intelligence innovation through openness and transparency

We give organizations and customers the visibility and relevance of security events by sharing our data from dynamic threat intelligence and our continued collaboration with the MSTIC team. This collaboration has proven successful in instances where Microsoft Security teams have actively tracked large-scale extortion campaigns targeted at multiple organizations, resulting in an industry-wide effort to understand and track the threat actor's tactics and targets.

The NOBELIUM incident in late 2021 was another instance of a large-scale cyberattack that launched a global hunting effort formed around MSTIC and Microsoft's team of global security experts. The threat actor targeted privileged accounts of service providers to move laterally in cloud environments, leveraging trusted relationships to gain access to downstream cloud service provider (CSP) customers. We engaged directly with affected customers to assist with incident response and drive detection and guidance around this activity. Through a successful partnership and continuous feedback loop, we have been able to improve our ability to minimize impact and protect customers over time.

The work we delivered in protecting customers against NOBELIUM attacks would not have been possible if not for the continuous hunting process and feedback loop with threat intelligence. We've crafted a symbiotic relationship that empowers threat hunters at DART to become better incident responders by looking at additional vectors seen in threat intelligence platforms.

Seeing events from a threat intelligence perspective, applying the art and science of threat hunting, and partnering with the security industry at large are all but signs of our commitment to growth and helping organizations stay protected from cyberattacks.

Learn more

To read about the latest attack methods and cybersecurity best practices from our investigations and engagements, visit the Microsoft Detection and Response Team (DART) blog series. To learn more about our specialized incident response and recovery services before, during, and after a cybersecurity crisis, visit Microsoft Security Services for Incident Response.

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.