Structured threat hunting: One way Microsoft Threat Experts prioritizes customer defense

microsoft.com/security/blog/2021/12/02/structured-threat-hunting-one-way-microsoft-threat-experts-prioritizes-customer-defense/

December 2, 2021

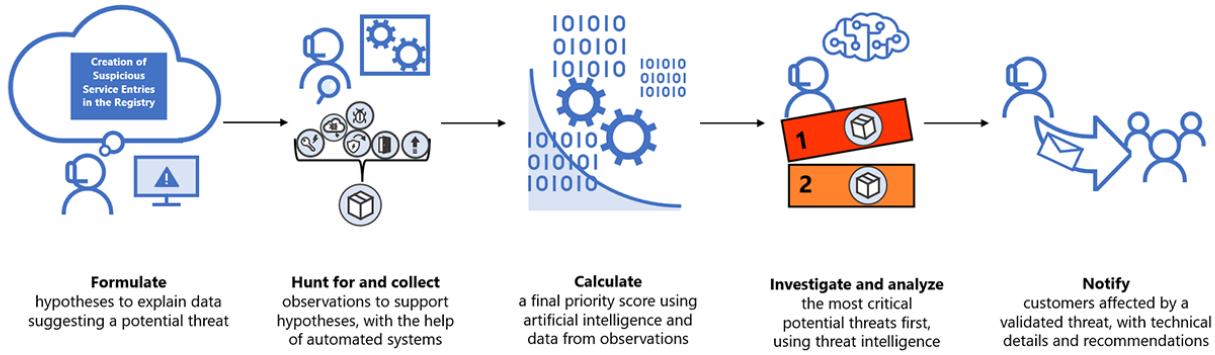

***Update [5/9/2022]**: In line with the recently announced expansion into a new service category called **Microsoft Security Experts**, we're introducing the availability of Microsoft Defender Experts for Hunting for public preview. Defender Experts for Hunting is for customers who have a robust security operations center but want Microsoft to help them proactively hunt for threats across Microsoft Defender data, including endpoints, Office 365, cloud applications, and identity.*

Today's threat landscape is incredibly fast-paced. New campaigns surface all the time, and the amount of damage that they can cause is not always immediately apparent. Security operations centers (SOCs) must be equipped with the tools and insight to identify and resolve potentially high-impact threats before attackers set up persistence mechanisms, exfiltrate data, or deploy payloads such as ransomware.

Every day at Microsoft, threat hunters work alongside advanced systems to analyze billions of signals, looking for threats that might affect customers. Due to the sheer volume of data, we're meticulous about surfacing threats that customers need to be notified about as quickly and accurately as possible. This helps ensure that customers can respond to the most critical threats in their environments through our products and services, such as Microsoft Threat Experts, a managed threat hunting service that provides expert-level monitoring and

analysis. With Microsoft Threat Experts, customers get Targeted Attack Notifications, which are designed to identify the most important risks and provide technical information, as well as hunting and mitigation guidance. Customers can also consult with our analysts through Experts on Demand.

Microsoft Threat Experts allows organizations to collaborate with Microsoft analysts to benefit from their expertise in tackling critical incidents. This collaborative relationship also gives Microsoft analysts the opportunity to gain invaluable insight into real-world threats, how attackers operate inside enterprise networks, and how security operations teams function. This creates an environment conducive to mutual learning and innovation, which helps improve our processes, protections, and services.

One way we support this close collaboration with customers is through a structured approach to threat hunting. Human analysts are augmented by AI in the search for potential threats to our customers. AI helps our human analysts tease out which events in our data require closer examination. Humans drive the initial hunt, then validate the AI's findings, and provide deeper analysis and context for each potential threat. By combining what humans and AI are each individually good at, we're able to process data at speed and scale. The process ultimately decides which potential threats to address first.

Our strategy is designed to evaluate impact and escalate potential threats for investigation, based on how damaging the potential threat would be if it was found to be valid. Our strategy is also designed for speed: due to the highly time-sensitive nature of the threat response, our security analysts need to be confident that the most dangerous potential threats are analyzed first.

This process starts with Microsoft analysts formulating hypotheses to explain suspicious behavior discovered within our data. If the hypotheses pass our initial quality checks, we perform an automated hunt for and collect observations that could ultimately confirm or deny our suspicions.

Once we've gathered more evidence, the observations are grouped into potential threats and run through a variety of computations to evaluate the possible impact. One of the key figures we use for evaluating potential threats is the amount of diversity we see across our observations. A more diverse set of observations indicates that a potential threat is more likely to have a broad impact.

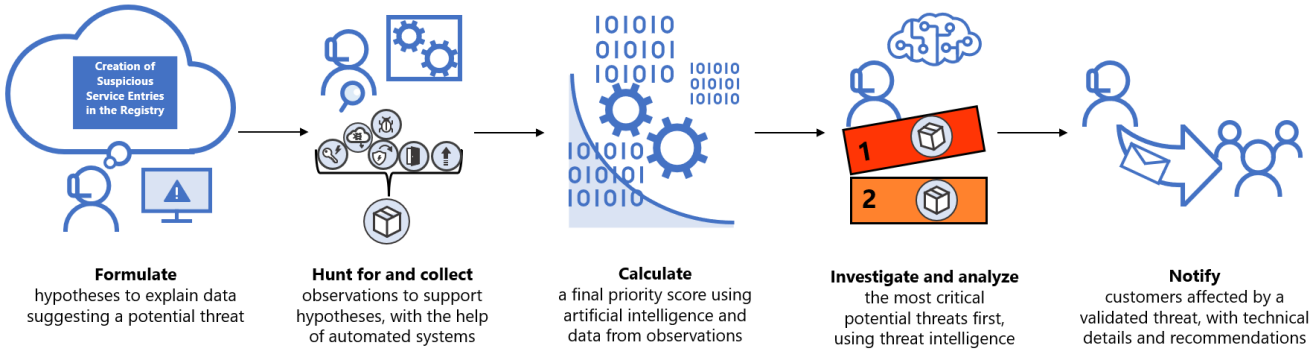

Figure 1. Visual overview of the prioritization process.

The computations on the impact of the potential threat are then combined, to calculate a final priority score. The priority score is used to sort potential threats according to how likely they are to require urgent attention. Potential threats that are more likely to have a devastating impact are given higher priority scores, so that they can be addressed more quickly.

Seasoned threat experts investigate and analyze the potential threat once it is ready. If the threat is found valid, our analysts conduct a deep-dive investigation, gathering the information our Microsoft Threat Experts customers need to keep themselves safe. They collect technical details and develop security recommendations. Affected customers are immediately notified with targeted attack notifications, containing detailed information on what the threat is, and what they can do to defend themselves. These steps are summarized in Figure 1.

Our process uses automated hunting and AI to speed up the decision-making, while humans define the observables, adjust and tune the parameters, perform triage, and craft the targeted attack notifications sent to affected customers. A deeper look at our process will show how analysts work closely alongside automation to help protect customers.

Hunting for potential threats

As the first step in our process, threat hunters formulate a hypothesis around data related to a potential threat, such as, “The attacker remotely executed code by exploiting a vulnerability in a system process.” After validating the soundness of the hypothesis by measuring the signal-to noise ratio and using known data sets to ensure that the accuracy is within acceptable limits, the hypothesis is modeled in our hunting systems, which automatically perform data collection, correlation, and enrichment.

These automated systems collect observations through our telemetry, from multiple devices and often from different stages of an attack. Each observation represents a single instance of a hypothesis – in our example, here is the system process, here is what the system process

did, and here are the arguments that were passed to it. Examining the observations associated with the hypothesis helps analysts determine if that instance of the hypothesis is valid.

Observations are then grouped into potential threats that represent our best current understanding of a collection of observations. If these observations truly reflect malicious behavior, the picture they paint is that of an attack on a customer's computing infrastructure. Looking at data associated with these observations can help us understand which potential threats may wreak the most havoc, and prioritize the investigation of more critical threats.

Developing priorities

If a potential threat is malicious, it is more likely to cause critical damage if it involves many devices inside an organization and is associated with many distinct attack stages. A confirmed threat that involves a single observation and is confined to an isolated machine is less likely to have the same level of impact.

Similarly, a potential threat is more likely to have more impact if it involves observations supporting many different hunting hypotheses. If the potential threat is malicious, a greater variety of hypotheses reflects a broad-ranging attack that hits many parts of the organization's infrastructure in many different ways.

Therefore, a potential threat with more diverse hypotheses is more likely to have a big impact on our customers, and is prioritized for investigation by Microsoft Threat Experts.

Using diversity to prioritize potential threats

How can we measure the diversity of different aspects of a potential threat? Microsoft Threat Experts uses entropy, borrowed from information theory. Entropy measures how many bits (or yes/no questions) it would take, on average, to identify a random element of a set if we know the elements of the set.

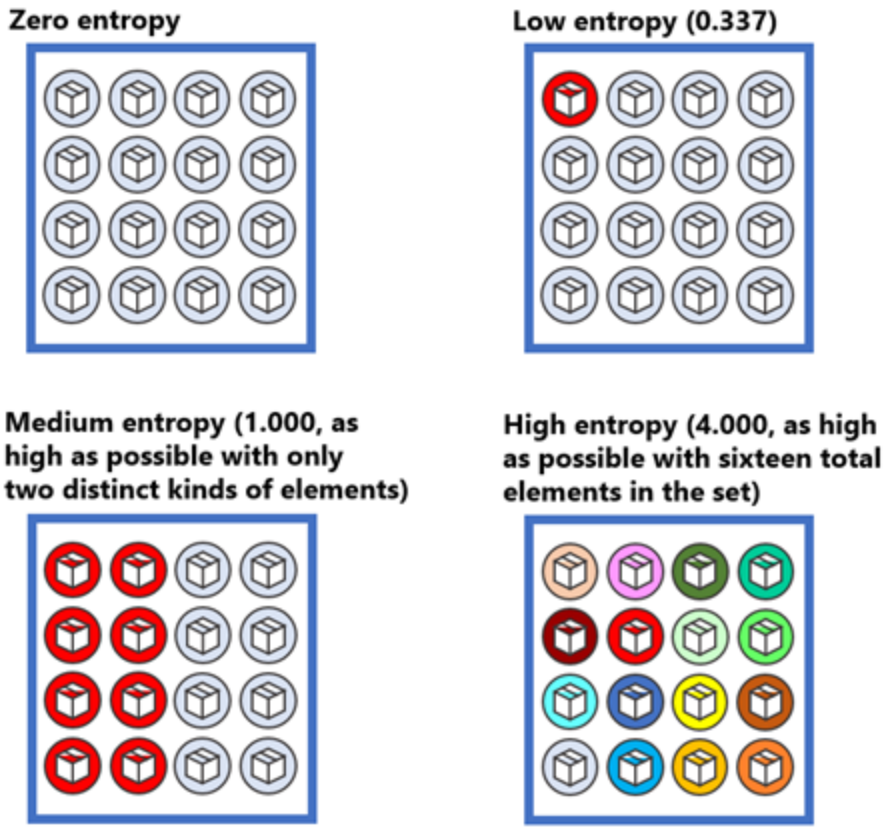

Figure 2. A visual representation of entropy, demonstrating how sets with many distinct elements have higher entropy than sets with fewer distinct elements.

A set containing elements that are all alike has zero entropy. Meanwhile, a set with more distinct elements, or the distinct elements in more equal proportions, will have higher entropy.

We use entropy to help decide which potential threats need swift attention, by calculating it for several attributes that the observations in each potential threat might have. For example, we calculate the entropy for the hypotheses associated with observations for each potential threat. We also calculate the entropy for the MITRE techniques associated with the observations in the potential threat.

After the AI automatically calculates an entropy value for each of these categories, and values for other factors, such as the severity of the hypothesis for each of the observations involved, it combines these values together to create an overall priority score for the threat.

Calculating the final priority score

Since our final score takes into consideration many different kinds of information about the potential threat, we need a way to combine these values. To do this, we convert each the results of each calculation into a p-value.

A p-value represents the percentage of potential threats we'd expect to have a certain value or larger in that category. For example, if only 5% of the MITRE technique entropy values were larger than the value from our calculation, then the p-value for our potential threat's hypothesis entropy would be 0.05.

We do this same p-value conversion using other numbers we've generated from the same potential threat, some based on entropy and others not. We then combine all of these p-values into a final priority score.

Validating the potential threat

The final prioritization score is used to sort potential threats, so that the most critical potential threats are analyzed the most quickly. When a potential threat is ready, a dedicated team of security experts look over the results and perform deep analysis to determine the threat's validity.

The analysts closely investigate suspicious activity related to the potential threat, using information from possibly-affected devices and networks. They search for unexpected events on the device timeline, indicators of compromise, and other evidence that something malicious may have occurred.

If the threat is validated and the activity is found to be malicious, our experts set to work to determine the full scope and protect customers. Related activity surrounding the validated threat is tracked down, technical details about the systems affected are gathered, and the team develops security recommendations to defend against the threat.

Microsoft Threat Experts: Helping defenders help themselves

Whenever we discover and validate a critical threat in the environment of a Microsoft Threat Experts customer, we immediately send out a targeted attack notification. These are special alerts that provide deep context about the threat, with details specific to each customer's unique environment.

Targeted attack notifications aim to help SOCs formulate a response before a critical attack can wreak havoc on their network. They help highlight the most critical threats, and clearly identify the ones that need aggressive handling. They also contain key technical details, which can assist SOCs in swiftly handling an ongoing threat. Customers receive a timeline of observed events in their organization, as well as advanced hunting queries for surfacing threat activities. These details can aid in understanding the attack flow and discovering the scope of the threat.

A case study in prioritization

Our process can fill the gaps to realize the true scope of a potential threat. Recently, there was a potential threat involving pre-ransomware activity detected by our machine learning. Outside of the context provided by Microsoft Threat Experts, it appeared as if this was just another instance of a widespread Qakbot campaign. It did not initially appear that this threat was any more dangerous than similar instances of the same campaign.

However, the Microsoft Threat Experts prioritization process had pieced together the evidence to suggest that the threat could represent one of the most impactful attacks seen that month. The potential threat was associated with a diverse collection of hunting hypotheses, as measured by entropy, and thus held a high priority for active investigation by Microsoft Threat Experts. Since the process evaluated it to be very high priority, Microsoft Threat Experts quickly assessed the potential threat.

The potential threat was found not only to be valid, but to be every bit as dangerous as suggested: although the techniques were largely typical of the ongoing Qakbot campaign, there was an especially swift progression to reconnaissance and credential theft. In addition, because of the method the attacker used to launch the backdoor, the activity was detected and alerted on, but not fully remediated.

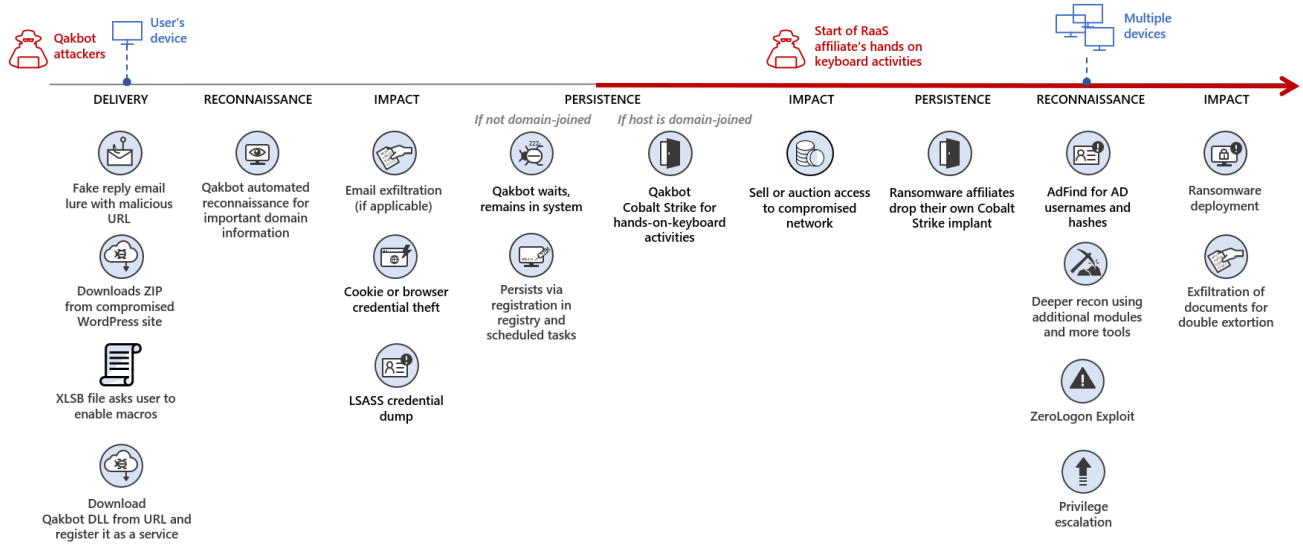

Figure 3. The flow of a more typical Qakbot infection. In this atypical case, human operators began late stage, hands-on-keyboard reconnaissance soon after initial entry.

This sort of progression was not normally associated with the flow of an initial entry campaign, but with human operators attempting lateral movement. The attacker seemed to be rushing to find out as much as they could, to deal out maximum damage. Furthermore, Microsoft hunters had identified these same techniques being used in ransomware attacks, strongly indicating that the operators might soon move to ransom the organization's devices. Qakbot has been known to facilitate ransomware-as-a-service (RaaS) activity. In the RaaS

model, a RaaS operator works with affiliates and provides tools for launching ransomware attacks. The affiliates deploy ransomware payloads by purchasing access to networks with existing malware infections like Qakbot.

Analysts at Microsoft Threat Experts immediately alerted the organization and provided advice on how to deal with the validated threat. Security researchers and experts at Microsoft Threat Intelligence Center (MSTIC) and Microsoft Detection and Response Team (DART) provided further help to prevent the attack from escalating. This aided the defenders at the targeted organization as they moved to remediate the threat. The collaboration between Microsoft and the targeted organization ultimately succeeded in stopping the attack. In spite of the attackers achieving broad lateral movement across the organization and nearly achieving their objective, the organization was not ransomed and were able to recover.

Empower your organization

In this example, and many others, notifications from Microsoft Threat Experts can be invaluable to decreasing the damage threats pose to organizations. Thoroughly warned of the threat at hand, organizations can take immediate action to remediate the threat using information from the notification, and perform further analysis with the suite of investigation tools available through Microsoft Defender for Endpoint.

Targeted attack notifications aren't the only kind of assistance available through Microsoft Threat Experts. Organizations can also send inquiries to our analysts, through purchasing Experts on Demand. By consulting with analysts through Experts on Demand, customers can get more context on an alert, gain clarity on the root cause of an incident, or receive personalized guidance on how to protect their organization.

Through Microsoft Threat Experts, SOCs are empowered to act quickly and decisively. Learn how your organization can get expert level monitoring and analysis through Microsoft Threat Experts.