Mandiant

SCOPING A THREAT HUNT WITH INTELLIGENCE

This article focuses on the cyber defense function: **Threat Hunting** – using intelligence to frame queries to find evidence of advanced persistent threat (APT) groups within your environment.

Read Time: 8-10 Min

If you haven't read our primer on the importance of the Threat Hunt Mission, we recommend you spend 6-8 minutes reviewing the article **Introduction to an Intelligence-led Threat Hunting Framework (https://docs.mandiant.com/home/introduction-to-an-intelligence-led-threat-hunting-framework)**.

To recap, Mandiant defines Threat Hunting as **the methodical, use-case driven, proactive identification of cyber threats within a computing environment or infrastructure**.

Objectives and Key Actions for Application

Cyber threat intelligence is required to help identify those threats we should be looking for, based on the likelihood and impact of those threats. In other words, organizations should be actively looking for those threats they expect to target them.

Threat Hunting complements security detection. Organizations often focus on reactive defense - identifying the presence of attacks through alert generation of their malicious activities by deployed security monitoring solutions. Threat Hunting is pro-active – actively searching through your environment to detect and isolate signs of malicious activity that have evaded existing security monitoring solutions. These two processes complement one another.

For this use case article, we'll be focusing on how to frame a Threat Hunt Query as part of the Assess phase of Mandiant's A4 Threat Hunting Framework:

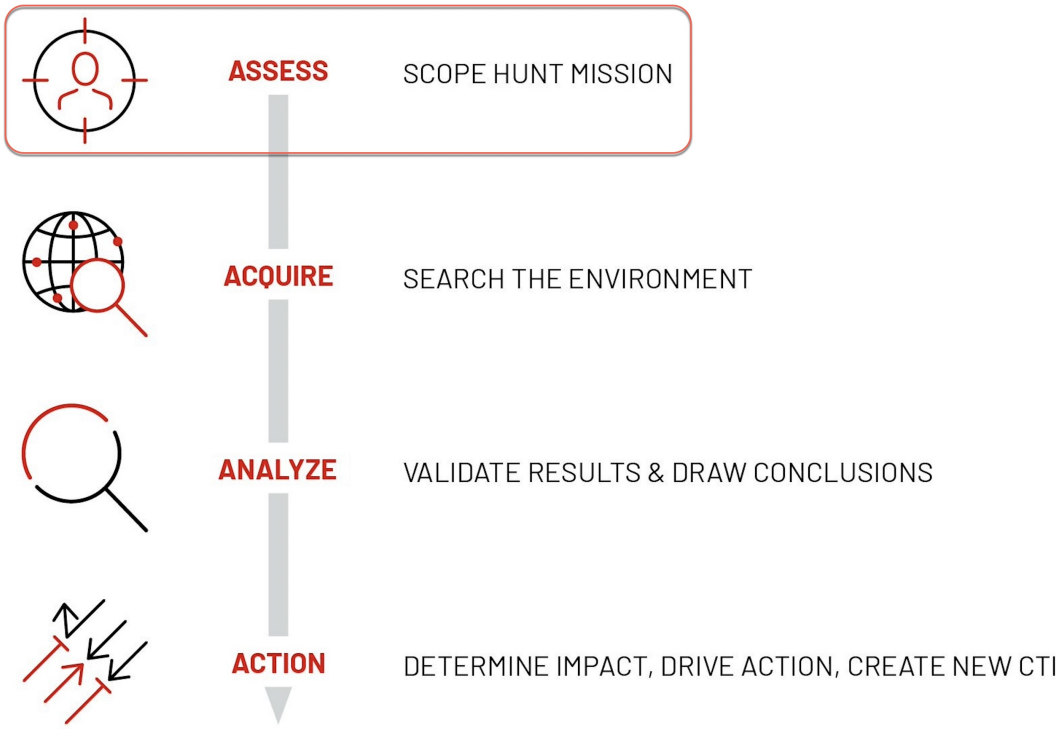

For a more detailed dive into the A4 Threat Hunting Framework, see **Introduction to an Intelligence-led Threat Hunting Framework (https://docs.mandiant.com/home/introduction-to-an-intelligence-led-threat-hunting-framework)**.

Why Scope?

Scoping is vital to ensure intelligence efforts aren't wasted. With finite resources and time available, analysts need to maximize their return on effort. Analysts need to hunt for those threats that are most likely to be present in your environment and have the potential for significant impact. Organizations should use cyber threat intelligence to define potential hunt targets. Ideally, the overall threat hunt mission should pull targets directly from the Cyber Threat Profile (CTP) and/or be prioritized from current or emerging significant trends identified by analysts focused on the threat landscape. To help align your Threat Hunt missions to the broader intelligence organization's CTP, see **Creating a Cyber Threat Profile (https://docs.mandiant.com/home/creating-a-cyber-threat-profile)**.

The scope of a Threat Hunt Query outlines the specific parameters under which analysts can conduct a defined search for specific conditions that may indicate adversary activity. These parameters need to balance between being too prescriptive and too broad—listing too far either way will either return no useable results or return too many results to accurately determine an adversary's presence.

The focus of Threat Hunt Missions are going to differ depending on your type of business, industry, geography, size, internet presence, etc.

Ultimately, scoping the Threat Hunt Query ensures that analysts don't go looking for something that is unlikely to have happened in the first place.

How to Scope a Threat Hunt

Like many intelligence-related activities, Threat Hunts should be scoped using a repeatable method to ensure consistent outcomes. Establishing a repeatable method for scoping Threat Hunts ensures the production of consistently measurable outcomes that, with all things being equal, results in the best chances to pro-actively discover previously hidden threats.

Think of Threat Hunt Missions as a science experiment with slightly changing variables. While different variables (adversary tactics, operating environment, tools) may alter, the way you frame, execute, and assess the results needs to remain consistent to ensure you (as the scientist) minimize your footprint on the results. Different variables should return different results, not different approaches to the same variables.

Likelihood and Impact

If you were an analyst for a major Australian bank, consider which adversary would be best to hunt for:

- An adversary that made a big splash in the open-source media for their significant attacks against US targets within the Defense industry
- An adversary that had targeted several Southeast Asian financial institutions over a period of 18 months using hard-to-detect tactics that are difficult for the open-source media to explain in simple terms

Analysts planning to conduct Threat Hunt Missions should have a good working knowledge of the operating environment (situational awareness), have access to guiding intelligence components such as the Intelligence Requirements and CTP, and be trained to construct hypotheses to test analytical assertions (in this case, the assertion will likely be: Has Threat Group *A* targeted my organization?).

While beyond the scope of this digest, analysts will also need familiarity with the operational environment (i.e. where the adversary would have attacked) and have the requisite access and tools that enable a fulsome search within this environment. This will be explored within a separate article on the Acquire phase of Threat Hunting.

Different Organizations, Different Threat Profiles

Imagine two different organizations – one, a high-profile bank in Australia, another a utilities provider in the Netherlands. Their respective CTP – who is likely to target these companies – will differ greatly. For example:

- An Australian bank may be targeted by globally based sophisticated financial threat groups backed by money mules on the ground
- A Dutch utilities provider may be targeted by agents of a foreign power, seeking access to security and safety controls that may disrupt productivity and supply

These two organizations may be threatened by similar threat actors. Both are critical infrastructure. However, the nature of their businesses means different threat groups would benefit from gaining access to their different assets. One threat group may just want to steal money. Another threat group may have espionage in mind. Different types of organizations will face different types of threats. A CTP needs to be calibrated to the needs of each organization and act as a guide for defense and security controls.

Analysts planning to conduct Threat Hunt Missions should possess the following:

- A strong working knowledge of the operating environment (situational awareness)
- Access to guiding intelligence components such as the Intelligence Requirements and CTP documents
- The expertise to construct hypotheses to test analytical assertions
- The access and tools necessary to conduct searches across the entire operating environment (**Note**: This will be explored in a separate article about the Acquire phase of Threat Hunting.)

Threat Hunt Query Process

Threat Profile and Intelligence Requirements

A threat hunt query cannot operate in isolation from the needs of your organization. Both the construction of the query and the subject (or target) of the query need to be aligned and relevant to broader business needs. To this end, the CTP

and Intelligence Requirements form the basis to identify the most significant threats based on likelihood and impact. These foundational intelligence components should be consulted by analysts constructing a Threat Hunt Query in the first instance.

This is especially true when starting this process. Threat Hunts should be conducted systematically and methodically over time, addressing each significant threat in turn. The CTP and Intelligence Requirements can help those undertaking Threat Hunt activities to prioritize an appropriate schedule.

This isn't to say that new, emerging, or immediate threats shouldn't be incorporated into Threat Hunts. When a new threat arises, the threats and risks within the CTP and Intelligence Requirements can serve as benchmarks to determine whether the new threat deserves higher prioritization.

Situational Awareness

Situational awareness is the working knowledge of the threat environment. Simply put, it's the analysts' combined knowledge of current threats, recent trends, and knowledge of adversaries and their motivations. Situational awareness is not necessarily an encyclopedic recall of every threat actor, their Tactics, Techniques, and Procedures (TTPs), their targets, capabilities, etc., but rather an overall awareness of what is currently happening, or recently happened, within the cyber threat world. Situational awareness comes from the day-to-day activities and tasks conducted by analysts but it also requires critical and curious thinking on behalf of the individual analyst. As an analyst, you need to cultivate an active interest not just in threat activities relevant to your business, but to the operating threat environment as a whole. Situational awareness will assist in forming relevant and useful hypotheses to start the Threat Hunt Query.

Additionally, efforts should be made to catalogue and summarize all threats faced by your organization. This summary will provide important historical context that helps inform the CTP, but will also help place new attack activity within a broader context relevant to your organization.

Historical Context

Over time, analysts develop an innate memory about attacks that impact their organization. Organizations should attempt to catalogue or summarize these activities for future, quick reference.

Knowing what has happened previously can inform how we handle current threats, often providing prompts to determine action. Prompts can include:

- How was the same threat handled previously?
- Is this the same threat or have elements changed?
- Is this repeated malware evidence of the same adversary returning?
- What lessons can we implement from the previous time this same threat occurred?

This last prompt in particular is useful in the context of Threat Hunt Queries.

Develop Hypothesis

A hypothesis is a working explanation based on limited evidence that helps frame or guide future investigation. In the context of Threat Hunt Queries, a hypothesis is the supposition that an adversary is hiding within our environment and we need to use our knowledge of their capabilities and methods to uncover their presence.

Our hypothesis should use the knowledge about the threat gathered from the foundational components and analysts working knowledge (as above). From there, we should ask ourselves:

- **What's the nature and type of the threat we wish to uncover?**

Is it something simple like a particular malware or tool that is used by attackers? Or is it an unusual pattern? Avoid making the hypothesis solely indicator-based as this can limit discovery.

- **Where will the activity likely occur?**

You should identify where the activity will likely occur based on the vector or the TTP.

- **Why do we care?**

What is the potential impact to the organization if the threat exists? The hypothesis needs to account for potential recommendations or courses of action if the activity is found, based on the significance of the threat.

Hypothesis Example

An example hypothesis structure:

<**Threat Actor**> verb <**capability**> preposition/verb <**infrastructure**> to achieve their <**objective**> against <**victim**>.

For instance:

Suspected APTx threat actors are exercising social-engineering tactics via email involving spoofed domains to achieve access into mission-critical systems within your organization's industry/sector.

Identify Target and Source Data

Target data is the operational and tactical intelligence that informs on adversaries' TTPs. In essence, how they conduct their attacks.

Source data is the information on your organization's infrastructure, both the data held within your organization and the information on how the infrastructure communicates, operates, and defends (otherwise known as network architecture).

If we were to generalize, a Threat Hunt Query seeks to compare the data used by adversaries against source data to determine the presence of potential malicious activity. Crucially however, analysts and defenders will need to interpret how target data has interacted with the source data to determine this presence: in many cases, the data used by adversaries (e.g., malware) have similar functions to other known malware but interact with infrastructure or are used in completely different ways from one another. It's this interaction, represented by anomalous artifacts or patterns within network traffic, that indicate the potential presence of an attacker.

The takeaway here is that analysts who execute the Threat Hunt Query need technical expertise that understands how adversaries operate and what that may potentially look like within the infrastructure of their organization. This requires good working knowledge of adversaries' capabilities and methods, and a deep understanding of how your organization's network and infrastructure works.

This is another reason why scoping is important—analysts cannot compare every threat against their infrastructure, especially when the likelihood of being targeted by certain adversaries is low. Prioritization of Threat Hunt Queries against high-risk threats is paramount to maximize your analytical skills.

Communicate Intent and Scope of Effort

Get into the habit of being explicit with the intent and scope of any Threat Hunt Query you construct and execute. Noting the intent and scope provides important parameters to your efforts and will serve to 'anchor' you when the execution of the Threat Hunt Query dregs up volumes of data. Clearly noting the intent and scope of the Threat Hunt Query will assist in keeping your efforts focused and manage stakeholder expectation on what the Threat Hunt Query is intended to achieve.

The intent also provides justification for the analytical effort being undertaken. Given the nature of the activity proposed, this justification will center around the risks posed by a specific adversary to your organization. Such a justification may be required if the results of your Threat Hunt Query return negative results. In this event, you must be careful to separate the two potential outcomes:

- That the identified risk posed by the adversary wasn't found
- That the identified risk posed by the adversary may still exist, but your query didn't return positive results (Note: this may be because of limitations to the search. See below.)

It can be difficult to determine which of these outcomes is accurate. In these instances, we recommend you focus on the risk posed by the adversary to business operations. The intent of the Threat Hunt Query should encapsulate this risk, as a mix of the likelihood of the adversary targeting your organization and the impact if they have targeted your organization and remain undetected, even after an initial Threat Hunt Query.

In the event of a high likelihood/high impact risk, a clearly worded intent can provide justification for further execution of the same (or slightly altered) Threat Hunt Query to ensure our initial results were accurate.

Determine Hunt Timeframe

Having a defined timeframe for the Threat Hunt Query provides a further parameter to properly scope the query. The timeframe should consider how long sources maintain high fidelity data and commonalities within attack patterns that occur with a significant time difference separating them.

For example, if malware is used prolifically within campaigns for a certain period, then goes dormant for 18 months before returning, analysts should be wary presuming the malware has the same intent and capabilities – the significant time lag between deployments should raise questions about why the malware disappeared and why it returned, and whether the same adversary is using the malware.

Other elements, such as resources available for the hunt, and changes to your organization's infrastructure over time should also be considered. The longer the timeframe of the Threat Hunt Query, the more data to sift through and analyze. For this reason, we recommend defining discrete timeframes pertinent to the threat and your organization.

Relevant Timeframes

Cyber threat intelligence can provide insight into determining timeframes from Threat Hunt Queries. Consider how the following information may impact your selection of timeframes:

- An adversary has a period of known or suspected dormancy
- An adversary is known or suspected to have pivoted to target organizations matching your organization's profile at a certain date
- An adversary mainly targets a certain remote access platform which your organization stopped using 6 months prior to the adversary's period of activity
- Your main source of logging data only retains PCAP records for 6 months before archiving

Identify Execution Resources and Constraints

Good intelligence analysis is honest regarding the potential limitations faced during analysis. Constructing a Threat Hunt Query is no different. Just as the Threat Hunt Query needs parameters regarding context, timeframe, and adversary capabilities, it also requires an honest appraisal of potential constraints that may limit visibility or results from the Threat Hunt Query. Noting the limitations, including those on resources available, access into infrastructure or data, and limited scoping of the query due to lack of adversary knowledge, ensures the results of the Threat Hunt Query are appropriately caveated.

Common constraints within Threat Hunt Queries are often limited accessibility into network infrastructure or organizational data . Stakeholders consuming the results of the query need to be made aware that results may have differed if these constraints were removed or altered. Noting these limitations may change how the stakeholders react to the results.

Not noting limitations may therefore result in stakeholders reacting to the results in inaccurate ways.

Establish Expected Outcomes

An integral part of any Threat Hunt Query is the expected outcome. The expected outcomes do not preclude the results of the Threat Hunt Query but rather craft recommendations or courses of action for stakeholders given a range of different potential outcomes. These outcomes are usually scaled in terms of impact to the organization, starting with what would need to occur if the Threat Hunt Query determined the adversary was still active within your organization's environment.

Noting several expected outcomes helps prevent over- or under-reaction in the event of adversary discovery (or absence). The expected outcomes should consider the likelihood of the adversary's presence when determining potential recommendations.

As always, recommendations need to be feasible and calibrated to the specific risk represented by the adversary. There's little point quarantining frequently used services if the risk posed by the adversary's presence is relatively low.

Expected outcomes will need to balance the tension between your organization's security and productivity.

Template

Developing a template that encapsulates the different aspects of creating a Threat Hunt Query will help to formalize the process and produce consistent outcomes with each different query. The elements discussed above should be placed into a template that helps to summarize the context, hypothesis, target data, timeframe, and expected outcomes of each planned Threat Hunt Query.

Templates can also be useful to plan schedules for future Threat Hunt Queries. Instead of having to construct the same Threat Hunt Query from scratch, analysts can update previous queries with any new information, making the process of running scheduled Threat Hunt Queries more efficient.