FI-ISAC NL publication

TaHiTI

Threat Hunting Methodology

TaHiTI

•

A joint threat hunting methodology from the Dutch financial sector

•

Version 1.0
17.12.2018

Foreword

This threat hunting methodology was created as a joint effort between several Dutch financial institutions. The focus group operated as part of the Dutch financial institutes information sharing community (FI-ISAC). The goal of this cooperation was to create a joint understanding of threat hunting and a common approach to conducting threat hunting activity. This effort has resulted in the methodology described in this document: the **Ta**rgeted **H**unting **i**ntegrating **T**hreat **I**ntelligence (**TaHiTI**) methodology. This methodology has been created with a broad usage in mind: not only should it be valuable to the Dutch financial sector, but to any organization in any sector. Releasing this methodology and the accompanying practical resources to the public domain was part of the initial intent of the focus group and the setup of the methodology.

The methodology itself seeks to combine threat hunting and threat intelligence to provide a focused and risk-driven approach to threat hunting. Threat intelligence is used as a source for hunting investigations and is used throughout the investigation to further contextualize and enrich the hunt.

The **TaHiTI** methodology is supported by the ”MaGMa for threat hunting” tool, which allows hunters to document their results, structure the outcome of their hunting investigations and provide direction for growth of the threat hunting process. The tool can be downloaded from: https://www.betaalvereniging.nl/en/safety/tahiti/

This document and the MaGMa for threat hunting tool are released under the creative commons license and can be shared and adapted where required. As with any methodology: apply critical thinking, use what makes sense and adapt to fit your organization’s needs.

On behalf of the FI-ISAC focus group on threat hunting,
Happy hunting!

Rob van Os

Authors
Rob van Os, de Volksbank, *lead author*
Marcus Bakker, Rabobank*, co-author*
Ruben Bouman, ING / FinancialCERT
Martijn Docters van Leeuwen, Rabobank
Marco van der Kraan, Rabobank / FinancialCERT
Wesley Mentges, de Volksbank
Armand Piers, ABN AMRO Bank / FinancialCERT

Contents

1 Introduction

Threat hunting is a relatively new area of expertise. While the activity itself is not new, specific hunting tools, models and best practices have been developed in recent years. As with any new area, there is often confusion on what exactly comprises this activity. Good definitions are lacking, as are common approaches on how to perform such an activity. This document presents an approach for threat hunting that integrates with the threat intelligence process.

The 2017 SANS survey has indicated that only 4,6% of all companies engaging in threat hunting activities have adopted a published external methodology. Excluding outsourcing and companies that do not perform threat hunting, that leaves over 70% of organizations either using no methodology or a methodology that was created internally [[1]]. This shows a clear lack of availability of threat hunting methodologies that cover the entire process in a structured fashion. The 2018 SANS survey does state that wider sharing of best practices is expected [[2]].

Members of the Dutch financial sector that were conducting threat hunting activities have come to the same conclusion. In-house methodologies and hunting expertise were being developed separately. As such, the timing was right for a joint effort in creating a common understanding and common approach in threat hunting, as well as sharing best practices amongst each other. This document is the result of that effort and is shared publicly for those looking for a concrete approach on threat hunting.

To support the **TaHiTI** methodology, this document contains a template for documenting your threat hunting activities. Additionally, the "MaGMa for threat hunting" tool was created to document the results of hunting investigations. This tool provides insight into the performance of the threat hunting process, which can be used to provide focus for new investigations and improve the overall process performance.

The outline of this document is as follows:

- First, threat hunting is introduced. A definition is provided, and all aspects identified by the focus group are covered.
- Then, threat intelligence is introduced. Similar to threat hunting, this topic will be explained using a definition and an examination of all identified aspects and characteristics. While threat intelligence itself serves a wider purpose than merely threat hunting, the focus of this chapter will be on how threat intelligence can be used for decision making in threat hunting investigations.

[1] https://www.sans.org/reading-room/whitepapers/analyst/hunter-strikes-back-2017-threat-hunting-survey-37760

[2] https://www.sans.org/reading-room/whitepapers/analyst/2018-threat-hunting-survey-results-38600

- With the basic concepts underlying the methodology explained, **TaHiTI** is introduced. The methodology is explained in detail in a step-by-step fashion.
- Hereafter, metrics for threat hunting are covered.
- Finally, best practices are provided to guide the hunting process and the implementation of a hunting program.

This document also contains references to whitepapers and other resources used in the creation of this methodology and concludes with an annex containing a template to document hunting investigation and more information on the MaGMa for threat hunting tool.

2 Threat Hunting

This section outlines the concept of threat hunting by providing its definition, its purpose, its characteristics and the types of threat hunting investigations. The chapter concludes with hunting maturity models and the concept of the pyramid of pain.

2.1 Definition

Threat hunting in this document is defined as follows:
Threat hunting is the proactive effort of searching for signs of malicious activity in the IT infrastructure, both current and historical, that have evaded existing security defenses. This evasion of security defenses can be due to usage of new, improved or unknown attacker techniques, 0-day exploits or a lack of adequate detection technology within the organization. While incomplete or faulty configuration of detection technology or misinterpretation of security events by analysts during triage can be reasons for evasion as well, threat hunting assumes a properly running security monitoring process.

Besides defining what threat hunting is, the focus group also felt that it was important to note what threat hunting is not:

- It is not a form of pen testing, red teaming or purple teaming (although hunting activities could lead to insights on where to perform pen testing).
- It is not searching for IoCs (Indicator of Compromise) in the environment (although IoCs can be used in hunting activities). Note that other types of indicators (i.e. Indicators of Attack (IoAs)) are part of threat hunting. The difference between these types of indicators is explained in detail in a blog by Crowdstrike [[3]]. IoCs and IoAs are concepts in threat intelligence. Chapter 3 covers threat intelligence in more detail.
- It is not security monitoring (although output from hunting can be used to provide new detection mechanism that are followed up by security monitoring)
- It is not incident response (although hunting can lead to uncovering incidents, thus triggering the incident response process).
- It is not simply running a query in a tool (although automation and querying data is an important part of hunting activities). Simply put: if a tool can do it autonomously, it is not hunting. Threat hunters should use tools to support them in hunting investigations.
- It is not a process that has a guaranteed result. Not every hunt will uncover an attacker or lead to new detection mechanisms. This does not necessarily mean that attackers are not present. For example, the necessary data to perform the hunting investigation could be missing, or at the time of hunt the investigated IoA was not present. However,

[3] https://www.crowdstrike.com/blog/indicators-attack-vs-indicators-compromise/

hunting will always yield some secondary result, such as more insight into the infrastructure or identification of missing data.

- It is not easy to conduct. Threat hunting requires advanced knowledge of the environment and an excellent understanding of attacker capabilities. If traditional security monitoring is too challenging, threat hunting may be a bridge too far.

Note that hunting investigations and digging through data can have impact on privacy of the organization's employees. Thus, it is important to cooperate with the legal / privacy office in the organization during the setup of the hunting program and to ensure mandate for threat hunting activities.

2.2 Purpose

The main purpose of threat hunting is to reduce the time required to find traces of attackers that have already compromised the IT environment. By finding these traces as soon as possible, the impact of breaches to the organization can be minimized. The breach detection gap is an important concept in the context of this purpose.

Other benefits of threat hunting are:

- Identification of gaps in visibility necessary to detect and respond to a specific attacker TTP.
- Identification of gaps in detection.
- Development of new monitoring use cases and detection analytics.
- Uncovering new threats and TTPs that feedback to the threat intelligence process.
- Recommendations on new preventive measures.

2.2.1 Breach detection gap

As indicated, the goal of threat hunting is to decrease the gap between initial compromise by an attacker and the discovery of that attacker in the environment: the breach detection gap also known as *dwell time*. Figure 1 shows a timeline of an attack containing several key moments in time. The breach detection gap is the time between T=1 and T=2.

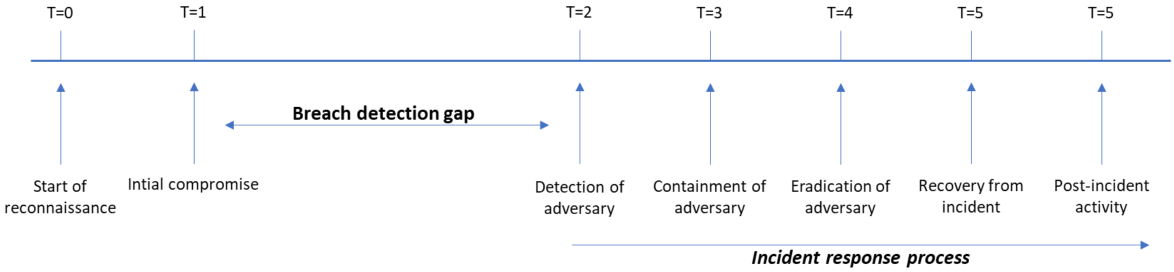

Figure 1: the breach detection gap

According to the latest Verizon Data Breach Investigation Report (DBIR), 68% of compromises went undetected for months [4]. Threat hunting plays an important role in reducing the breach detection gap. This is also evident in the SANS threat hunting survey, where improvements to incident response were mentioned as key improvements due to threat hunting activities. Threat hunting will aid to accelerate the detection of attackers by introducing new or improving existing detection mechanisms and thereby further closing the breach detection gap.

2.2.2 Dynamics of the breach detection gap

The breach detection gap stems from the ability of attackers to evade detection mechanisms. This is an important thing to realize. As detection capabilities continue to evolve and expand, cyber criminals will find new ways of evading these measures. Thus, over time, the Tactics, Techniques and Procedures (TTPs) of attackers will evolve to ensure that they can evade detection and operate unseen in an IT environment. TTPs can be defined as the tradecrafts of attackers.

There will always be a gap between what the organization is able to detect and the ability of a skilled attacker to avoid detection: the detection gap. Figure 2 shows the concept of the detection gap. Note that attacker capabilities will differ per attacker and detection capabilities will differ per organization. While attackers usually have capabilities to avoid detection, they may at some point trigger detection mechanisms. Either because these detection mechanisms evolved, or due to human error.

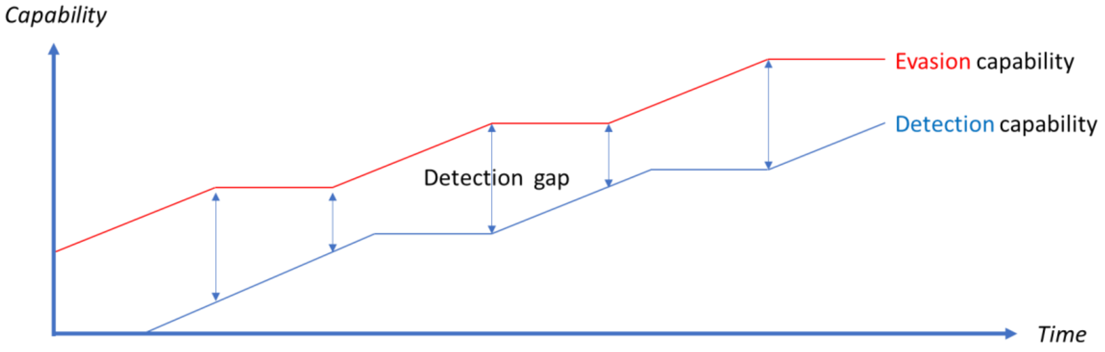

Figure 2: the detection gap

Note that the figure shows plateaus. These plateaus represent the cyber arms race where cyber criminals come up with new ways to compromise systems and evade detection,

[4] https://www.verizonenterprise.com/resources/reports/rp_DBIR_2018_Report_execsummary_en_xg.pdf

followed by the security industry closing the gap by developing new detection techniques or products.

A good threat hunting program aims to track malicious actor's TTPs and behavior and continuously reduce the breach detection gap. Particularly, threat hunting focuses on activity that would go undetected and thus continuously operates in the detection gap as seen in figure 2. Continuous insight into the state of detection mechanisms is required to avoid hunting for malicious activity that is already covered by traditional detection mechanisms. Use case management frameworks, such as MaGMa [[5]] can aid in such insight. As indicated before, mistakes during triage and security monitoring may also result in attackers going unnoticed. However, threat hunting is not meant as a control mechanism to check whether or not security analysts in the monitoring team are doing their job well. Threat hunting will focus on events outside the traditional detection capabilities, and may uncover missed or misinterpreted events during the hunt that can be used to improve detection and further train analysts.

Information on TTPs and actor capabilities is where threat intelligence comes in play. Threat intelligence can be used to determine TTPs of attackers. Thus, threat hunting uses threat intelligence in its process. In turn, threat hunting activities may uncover new TTPs that have not been identified or disclosed yet. Hence, threat hunting can provide unique insights into threat actor capabilities and generate threat intelligence. More on threat intelligence can be found in the chapter 3.

2.3 Characteristics

Several key characteristics of threat hunting were identified in the focus group sessions. Some of these characteristics were found to overlap with the research done by Anton Chuvakin (Gartner) [[6]], indicating that these are generally accepted characteristics:

- **Be proactive**. Threat hunters proactively search for indicators of malicious activity in the network instead of waiting from signals from traditional detection mechanisms to begin an investigation.
- **Assume the breach**. Being proactive does not make sense if you believe that prevention and detection mechanisms are sufficient to avoid breaches. Threat hunting assumes that there already was a breach and that it has not been identified yet.
- **Understand the attacker**. It is important to understand attacker motivations and mindset. These are important characteristics in determining how persistent and how capable an attacker is. Again, this is where threat hunting and threat intelligence meet, as this is a primary purpose of threat intelligence.

[5] https://www.betaalvereniging.nl/en/safety/magma/

[6] How to Hunt for Security Threats, Anton Chuvakin, Gartner

- **Detect the unknowns**. In a famous quote by Donald Rumsfeld [7], the concept of knowns and unknowns was introduced. Threat hunting focuses mostly on the known unknowns (finding traces of unknown attackers through known TTPs) and may uncover unknown unknowns in the process (finding traces of unknown attackers and previously unknown TTPs).
- **Creative and iterative process**. The threat hunting process is a creative process. The quality of the process heavily depends on the creativity, expertise and knowledge of the threat hunter conducting the hunting activity. It is an iterative process: threat hunting may lead to new insights and new hunting investigations; gathering information during the hunt may also lead to new assumptions about the current hunt. Chapter 4 on the **TaHiTI** methodology will explain this in more detail.
- **Data driven process**. Threat hunting requires data. Lots and lots of data. The SANS survey lists some preferred sources of information of threat hunting that include end point information, firewall logs, DNS logs, etc. [1]. The higher the quality of the data, the higher the likelihood of success of hunting investigations. The data should assist the hunter in the ability to perform investigations on hunting hypotheses and not complicate things by adding noise.
- **Based on hypotheses**. A common understanding between all resources that were used in the creation of this methodology was that hypotheses play a key role in the threat hunting process. For example, this has been described in more detail in publications by SANS [8] and RSA [9]. Given the fact that hypotheses play an important part in threat hunting activities and the lack of guidance on how to create effective hypotheses, the chapter on **TaHiTI** methodology will elaborate on this topic.
- **Team effort**. Hunting is a team effort. The hunting team uses a common approach and determines what to hunt for. The team will also prioritize hypotheses based on risk levels associated with the threat. Individual hunters will add their unique knowledge and skills to the team. The skills required for threat hunting can be found in multiple articles and publications, such as the Endgame threat hunting guide [10]. Usually, the skillset boils down to a few skills such as general security knowledge, IT environment knowledge, knowledge of analysis techniques, knowledge of attacker techniques and good communication skills.

2.4 Types of threat hunting

[7] http://archive.defense.gov/Transcripts/Transcript.aspx?TranscriptID=2636

[8] https://www.sans.org/reading-room/whitepapers/threats/generating-hypotheses-successful-threat-hunting-37172

[9] https://www.rsa.com/en-us/blog/2017-07/hypothesis-in-threat-hunting

[10] https://www.endgame.com/resource/white-paper/endgame-guide-threat-hunting-practitioners-edition

Threat hunting comes in different shapes and sizes. For example, Sqrrl lists 5 types of threat hunting [11]: data-driven, intelligence-driven, entity-driven, TTP-driven and hybrid hunting (a combination of 2 or more threat hunting types). For the sake of simplicity, in this document we only differentiate 2 types of threat hunting: structured hunting versus unstructured hunting.

2.4.1 Structured hunting

Structured hunting is hunting based on hypotheses: a hypothesis is created (the details on how to create hypotheses will be outlined in chapter 4), the hunting activity is scoped and subsequently performed. Looking at the hunting types defined by Sqrrl, **TaHiTI** is in essence an intelligence-driven methodology. However, since it revolves mostly around TTPs (TTP-driven) and can be potentially triggered by- or scoped for- crown jewels (entity-driven), these hunting types are also covered.

2.4.2 Unstructured hunting

Unstructured hunting is data-driven hunting. Potentially malicious activity can be detected by a hunter who is simply digging through available data looking for anomalies. This type of threat hunting does not start with a hypothesis, does not follow a predetermined path and is thus considered unstructured. Since **TaHiTI** is based on structured hunting, unstructured hunting is not in scope of this document, only as a source for triggers to start structured hunting. It must be noted that unstructured hunting requires a lot of effort and is much less likely to yield valuable results.

2.5 Pyramid of Pain

The pyramid of pain [12] is an important and elegant concept that can be used in threat hunting and threat intelligence. The pyramid addresses how difficult it is for attackers to change certain characteristics of their attack. At the same time, it also shows how difficult it is for organizations to find these characteristics. Finding a file with a certain hash value is easy, but uncovering illegitimate use of PowerShell in an organization where PowerShell is commonly used poses an entirely different challenge. Similarly, it is trivial for attackers to generate a new file with a different hash, but much harder to move or modify an attacker technique to evade detection. Figure 3 shows the pyramid of pain.

[11] https://sqrrl.com/5-types-threat-hunting/

[12] http://detect-respond.blogspot.com/2013/03/the-pyramid-of-pain.html

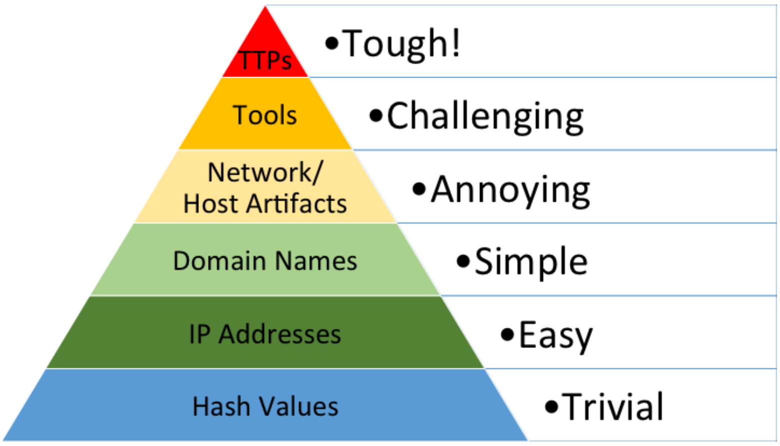

Figure 3: Pyramid of Pain (source: David Bianco, detect-respond blog). ***TaHiTI*** *focuses on the top 3 layers of the pyramid.*

The pyramid of pain connects threat hunting to threat intelligence. Threat intelligence provides relevant information on attackers on all layers of the pyramid. Threat hunting with the **TaHiTI** methodology will focus on the top 3 layers (but may use the lower 3 layers nonetheless) of the pyramid. TTPs are the top focus of threat hunting investigations, and has been widely covered in this chapter.

2.6 Hunting maturity

One of the most referenced resources in hunting publications is the hunting maturity model (HMM) published by David Bianco [13]. This hunting maturity model focuses on data collection, procedures, hypothesis creation, tools & techniques for hypothesis testing, pattern & TTP detection and analytics automation. The first 2 maturity levels, as explained in the model, are not threat hunting by the definition in this document. The first level (HMM0) is mostly ad-hoc and relies on automated alerting, while the second level (HMM1) only adds searching for IoCs (something that was explicitly excluded in the first paragraph of this chapter). Hunting at the third maturity level (HMM2) starts to use a structured approach and is considered hunting in the context of this document as long as the organization moves beyond hunting for simple IoCs.

The hunting maturity model can be combined with the pyramid of pain and the **TaHiTI** methodology to provide an overview of where each hunting maturity level acts in the pyramid and how **TaHiTI** should be positioned in the HMM. This overview is shown in figure 4. The hunting reference model by Sqrrl was used to map hunting activities to maturity levels [14].

[13] http://detect-respond.blogspot.com/2015/10/a-simple-hunting-maturity-model.html

[14] https://sqrrl.com/the-threat-hunting-reference-model-part-3-the-hunt-matrix/

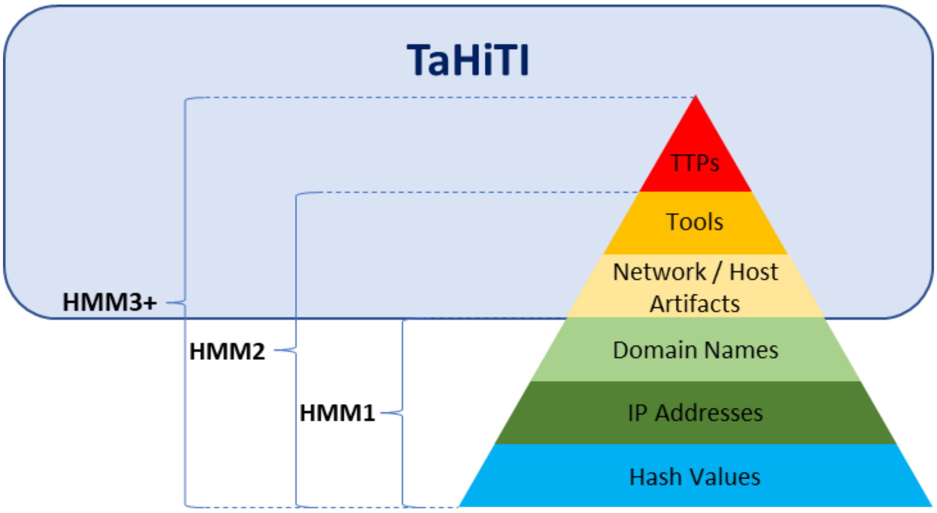

Figure 4: pyramid of pain, hunting maturity and the TaHiTI methodology

The lower layers of the Pyramid of Pain may be used in hunting investigations performed using the **TaHiTI** methodology, but the main focus will be on the top 3 layers. Does this mean the threat hunting process has to be at a very high maturity level before even engaging in threat hunting activities? Not necessarily. There are different maturity levels in which to apply the **TaHiTI** methodology. As more hunting investigations are performed, the process itself will be optimized:

- Gaps in visibility are be identified and resolved.
- New analysis techniques are utilized, optimized and implemented (which technique to apply in which situation).
- Procedures and reports (containing hunting results) are standardized.
- Hypothesis creation are refined and standardized.
- Data collection is standardized.
- Data analysis techniques become more refined.
- A threat hunting "platform" is leveraged. This does not necessarily imply an off-the-shelf platform, but can also be a custom set of tools and techniques combined with a data lake.

The above are all signs of increasing maturity and capability for the threat hunting team that make the team more effective. As with any maturity and capability model, there are other factors that play a role as well, such as organization and team dynamics. The Threat Hunting Team Maturity Model [15] provides a wider view of maturity and capability as it considers people, process and technology as well, and can be an augmentation to the HMM.

[15] https://www.happythreathunting.com/single-post/2017/10/29/Threat-Hunting-Team-Maturity-Model

3 Threat Intelligence

This chapter covers threat intelligence and focuses on how it relates to threat hunting. As **TaHiTI** is not a threat intelligence methodology, only the basics of threat intelligence are covered. Note that within the context of this document, the term threat intelligence is used to describe *cyber* threat intelligence. While cyber threat intelligence is derived from traditional threat intelligence (for example, as conducted by the military), there are different challenges and aspects to each field.

3.1 Definition

As with threat hunting, it helps to have a clear definition of threat intelligence. We define threat intelligence as follows:

Threat intelligence is the process of gathering, processing and dissemination of information about threats and attackers. The goal of threat intelligence is to contextualize the information and to deliver actionable information that can be used in the decision-making process.

The threat intelligence process puts information from the outside world into the organizational perspective and, if possible, advises on how to proceed. This requires determination of risk, impact and possibly mitigating measures from intelligence information. Threat intelligence also provides insight into how attackers operate, their motivation, the sectors and geographic locations they operate in and the level of capability they possess.

3.2 Purpose

The purpose of threat intelligence is to provide the organization insight into the threats they are facing. The threat intelligence process should yield actionable information about potential attackers, including their means, motive and opportunity. The process should also provide information on how the attackers operate, so that the organization can use this vital information to defend against these attackers, and to build detection mechanisms in their infrastructure. This can be a direct result from the threat intelligence process, but also from threat hunting activities triggered by threat intelligence.

3.3 Characteristics

Some of the characteristics that are important to threat intelligence are:

- **External and internal**. While external information is the most likely source of threat intelligence, internal information will play a role as well. Threat intelligence may come from the threat hunting process, the security incident response process, or from departments within the organization that have detailed information on internal risks linked to attack scenarios.
- **Tactical, operational and strategical**. A distinction is often made between these 3 levels of threat intelligence. Tactical threat intelligence provides information that can be used in security monitoring systems. For example: IP addresses associated with

command & control servers, malware hash values and known bad domain names. Tactical threat intelligence is preferably automatically verified. Operational threat intelligence is aimed at information higher up the pyramid of pain. Finally, strategical intelligence focuses on motivation, trends and forecasts, and is used for strategic decision making within the company [[16]].

- **Cooperation**. Threat intelligence has much more value once you start cooperating with a threat intelligence partner and actively taking part in threat intelligence communities. This is also the idea of the information sharing and analysis communities (ISACs). Such communities can be a source of threat intelligence, but can also be used to verify whether or not attacks are seen at your organization alone (and are thus targeted), or at other organizations as well. Helping your peers to be more secure can help the sector as a whole to become more secure.
- **Iterative process**. Processing threat intelligence information may lead to new insights about threat intelligence previously processed. For example, information on a new campaign may be connected to earlier campaigns from similar attack groups. Such information can be used to enrich current knowledge on attacker TTPs with additional information uncovered previously.

Actionable intelligence was also mentioned as a characteristic. The next paragraph focuses on this in more detail.

3.3.1 Actionable intelligence

Lately, the focus of threat intelligence has shifted from delivering intelligence feeds with a large number of IoCs to 'actionable intelligence'. The difference between intelligence and actionable intelligence is the same as the difference between data and useful information. Without proper context, threat intelligence is meaningless. IP addresses by themselves have little value, but once they can be connected to an attack campaign, and other indicators from that same campaign can also be found, context is created. This context provides direction for refined hunting investigations. Information on threats is rarely actionable as-is, it is made actionable through the organizations threat intelligence process by making sure it is:

- **Stakeholder-focused**. If the stakeholder is a technical team, the information is entirely different than when it is aimed at management level for strategic decision making.
- **Usable**. Threat intelligence must be usable to the organization. Thus, understanding the context of the IT environment and knowing what advice is technically and functionally feasible is important.
- **Credible**. The information itself should be credible and come from a reliable source. These are the 2 elements of the information reliability system as used by military

[16] https://securityintelligence.com/security-intelligence-at-the-strategic-operational-and-tactical-levels/

intelligence [17]. Keep in mind that some intelligence providers may copy each other's information. So, the same piece of information that was initially only found in a single source, can be found in multiple sources over time. This does not make the information more reliable, as there is still only one 'true' source. It is difficult, or even impossible to make this distinction. High quality intelligence providers will mostly generate their own intelligence, so they have much higher credibility. Note that even though the information itself may be credible, improper analysis of information (for example, due to bias) may actually reduce credibility. Thus, analysis on information conducted within the threat intelligence process must be done accurately to avoid incorrect outcomes.

- **Clear & concise**. This last part focuses on data quality attributes, such as completeness, relevance, timeliness and accuracy. Many other data quality attributes can be introduced into the process. For example, research on data quality has uncovered as much as 20 aggregated attributes [18]. Each organization should choose which attributes are deemed important to the threat intelligence process.

3.4 The relationship between threat hunting and threat intelligence

There is a clear relationship between threat hunting and threat intelligence. This has become apparent in the previous chapter, as some concepts in threat hunting are difficult to explain without basic knowledge of threat intelligence. For the **TaHiTI** methodology, 3 concrete elements of the relationship between threat intelligence and threat hunting are especially important:

- Intelligence as a starting point for hunting.
- Intelligence for contextualizing and driving the hunt.
- Hunting to generate intelligence.

3.4.1 Intelligence as a starting point for hunting

As threat intelligence provides us with a lot of information on attackers and their capabilities, it can be a major source for engaging in hunting activities. For example, a threat intelligence report describing an attacker group (such as the report on APT-1 [19]) and their distinct capabilities should be of great interest. If that attacker group also operates in your organization's sector and is also geographically relevant, the threat it poses may be significant. The threat intelligence process can trigger the threat hunting process based on this information and provide relevant context on the threat. Note that relevance of the actor group should not be the only factor to consider as it may provide a too narrow view for threat intelligence and threat hunters. TTPs can be shared by different actor groups, and actor groups

[17] https://fas.org/irp/doddir/army/fm2-22-3.pdf

[18] http://courses.washington.edu/geog482/resource/14_Beyond_Accuracy.pdf

[19] https://www.fireeye.com/content/dam/fireeye-www/services/pdfs/mandiant-apt1-report.pdf

can switch to other sectors at any time, so the threat intelligence team should keep a view of the threat actor landscape.

While systems can cover the low-level IoCs (such as IP addresses) automatically, hunting activities are required to determine whether an attacker has left any traces in the environment by looking at the attacker TTPs.

3.4.2 Intelligence for contextualizing and driving the hunt

During hunting investigations, threat intelligence can be used for contextualization of findings. For example, a certain TTP may be uncovered during the threat hunting process. Using threat intelligence, that information may be used to find related TTPs (for example, using the MITRE ATT&CK framework [[20]]) or additional information on that TTP. This can subsequently be used to further drive the hunt. This process is called pivoting and may lead to additional hunting activities or refinement of the active hunt. For the **TaHiTI** methodology, this interaction between threat intelligence and threat hunting is especially important. Context from threat intelligence may lead to extending the scope of the hunt, adding new data to the hunt, refining the hunting hypothesis or generating ideas for subsequent hunts. This interaction has also been described by Sqrrl, with the addition of incident response [[21]]. The full interaction between threat hunting and other security processes within **TaHiTI** can be found in chapter 4.

3.4.3 Hunting to generate threat intelligence

As mentioned earlier in paragraph 2.2, threat hunting can be a source for threat intelligence. Hunting investigations may uncover previously unknown TTPs for attackers. This information can be used in the threat intelligence process to build an attacker profile. All such information can subsequently be shared with peers in threat intelligence communities, providing them with information regarding the uncovered threat. If these peers start their own hunting investigations based on this new TTP, they may uncover additional indicators that can be shared with the threat intelligence community. This way, a more complete picture of attacker capabilities and TTPs can be built in a community effort. Within an active threat intelligence ecosystem, the sum is greater than the whole of its parts.

With the basics of threat hunting and threat intelligence and their interactions covered, the **TaHiTI** methodology will now be discussed. Some of its elements have been mentioned already, but will be covered in more detail in the next chapter.

[20] https://medium.com/mitre-attack/finding-related-att-ck-techniques-f1a4e8dfe2b6

[21] https://sqrrl.com/a-framework-for-cyber-threat-hunting-part-2-advanced-persistent-defense/

4 TaHiTI

This chapter explains the **TaHiTI** methodology by outlining the process and subsequently focusing on the individual elements. As indicated in the foreword of this document, **TaHiTI** stands for **Ta**rgeted **H**unting **i**ntegrating **T**hreat **I**ntelligence. *Targeted* because the methodology uses hypotheses to drive hunting activities. This means threat hunting is conducted with a specific goal in mind. *Integrating threat intelligence* because threat intelligence is a major source of threat hunting hypotheses, and is used to enrich and contextualize hunting activities. Lastly, threat intelligence may also be generated as a result of hunting activities.

4.1 The TaHiTI process overview

Figure 5 provides and overview of the **TaHiTI** process, its 3 phases: Initiate, hunt and finalize. The process has 6 steps in total.

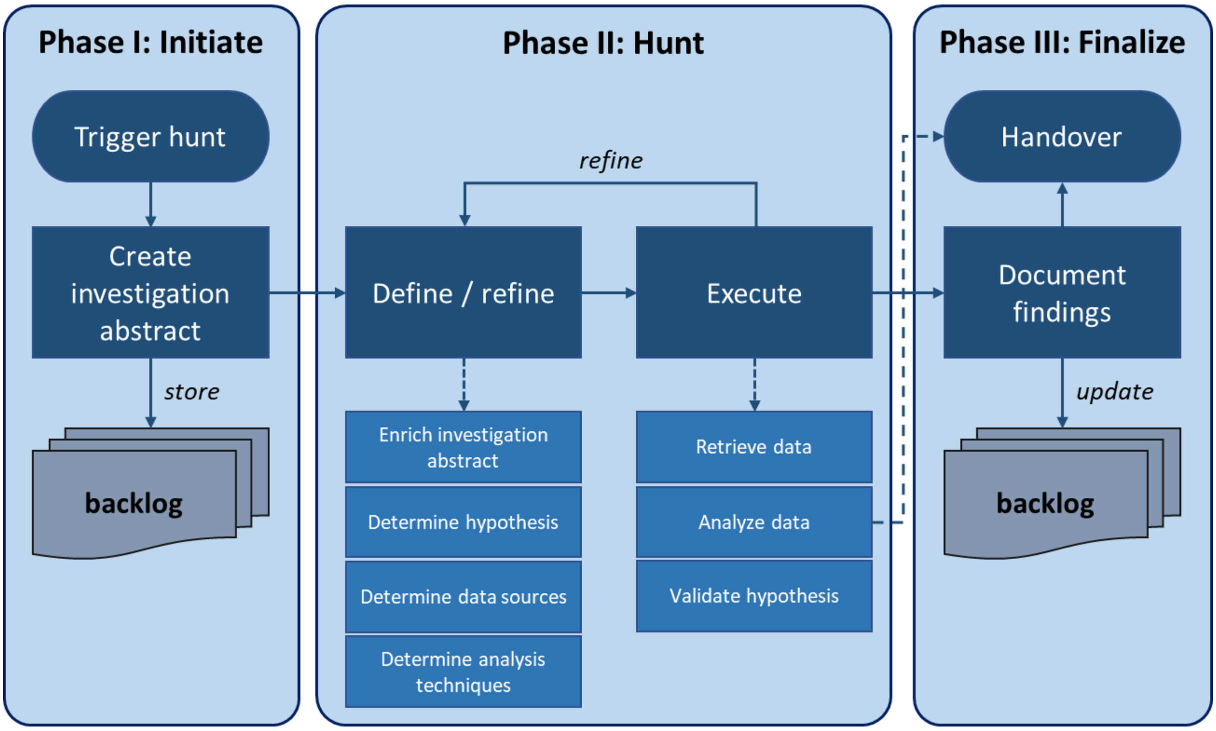

*Figure 5: the **TaHiTI** process*

4.2 Phase 1: Initiate

The initiation phase is where the input for threat hunting is processed. First, there is an initial trigger to initiate the hunting process. Next, the trigger is converted to an abstract of the hunting investigation and stored on the hunting backlog.

4.2.1 Step 1: Trigger hunt

The threat hunting process can be triggered from several processes. All triggers are shown in figure 6. An important thing to notice is that the processes that could potentially provide triggers to start the hunt strongly overlap with the processes that receive output from the investigation (figure 7). This feedback loop supports the threat hunting characteristic 'iterative process', as described in chapter 2. When executed well, hunting can act as an accelerator for improvement of these other processes.

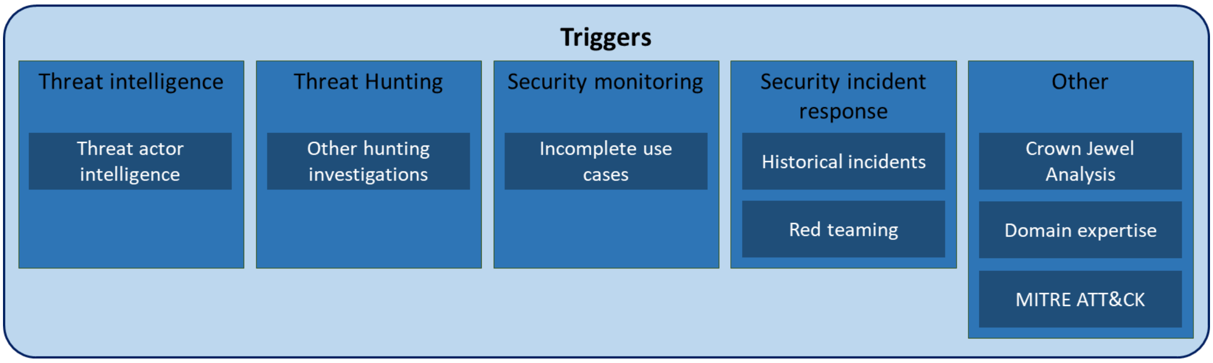

Figure 6: hunting triggers

Threat intelligence

Threat intelligence is a major source for hunting investigations. As show in figure 4, **TaHiTI** focuses on threat intelligence in the top layers of the Pyramid of Pain. This does not mean that information from lower layers is not used in threat hunting activities. It means that the information from the lower layers in the Pyramid of Pain will not trigger the threat hunting process.

Threat hunting

Threat hunting itself can trigger additional hunts as hunting investigations themselves may lead to new insights that require additional hunting investigations. Again, this demonstrates the iterative nature of the process.

Security monitoring

Security monitoring can trigger the threat hunting process through:

- **Incomplete use cases.** Some use cases may be incomplete and thus leave room for attackers to avoid detection. Such insights may come from red teaming or security incident response, but can also come from reviews of the use case framework. For example, the MaGMa use case framework provides a means for reviewing and identifying gaps in security monitoring. These gaps are candidates for hunting investigations. More on MaGMa and threat hunting can be found in Annex B.

Security incident response

Security incident response can trigger the hunting process through:

- **Historical incidents.** Historical incidents can be used for hunting investigations as well. Usually, during the post-mortem phase of an incident, actions will be taken to avoid such incidents from recurring. However, in some cases the root cause may be impossible to resolve. These cases are the best candidates for outlining hunting investigations as the method that was employed has been proven to be successful against the organization.

- **Red teaming.** Red teaming is an effective way of determining how well preventative and detective measures are able to defend against determined attackers. Red teaming will test technology (technical measures that prevent attackers from gaining access) as well as people (through social engineering) and processes (mainly the effectiveness of security monitoring and incident response). As such, it can provide valuable insights into what TTPs are effective against the target organization, and which tools can successfully be used to gain access and move laterally throughout the infrastructure.

Other sources

Other triggers for outlining new hunting investigations are domain expertise and crown jewel analysis. Crown jewel analysis as input for threat hunting is also discussed in the SANS paper "The Who, What, Where, When, Why and How of effective Threat Hunting" [[22]]. First, the organization evaluates its crown jewels and determines potential ways to compromise them. These methods of compromise can then be used to create hunting hypotheses. Attack tree analysis [[23]] can be part of determining attack vectors in the crown jewel analysis.

The MITRE ATT&CK framework [[24]] can be used as input for potential attack vectors and techniques, and contains a wealth of information for any hunter. The framework also provides suggestions for detection, which is valuable for both hunting and security monitoring. Note that this is not the primary purpose of the framework and should be treated as guidance for monitoring only.

Common sense, hunter experience, domain expertise and gut feeling. These are all elements that make threat hunting a creative process. Great threat hunters do not rely only on the previously mentioned triggers only, but over time develop a sense of what is important and what is not. This sense will also help hunters prioritize and select the most relevant hunts to execute.

4.2.2 Step 2: Create investigation abstract

When a trigger is received, the hunting team creates a hunting investigation abstract. This abstract does not include all details, but is a basic description of the investigation. Most of the

[22] https://www.sans.org/reading-room/whitepapers/analyst/who-what-where-when-effective-threat-hunting-36785

[23] https://www.sans.org/reading-room/whitepapers/threats/scalable-methods-conducting-cyber-threat-hunt-operations-37090

[24] https://attack.mitre.org/

information will be refined and updated in a later stage, when the hunt is selected for execution.

The following information is documented in the abstract:

- **Date.** The creation date of this abstract.
- **Initial hypothesis.** An initial hypothesis is added to the abstract. In a later stage (when the abstract is selected for the next hunt), this initial hypothesis is refined to create a definitive hypothesis.
- **Trigger.** What was the trigger for this abstract? Any information available should be attached to the abstract. Examples are URLs to online resources, or references to an internal security incident ticketing system.
- **Hunt priority.** The hunt can be prioritized based on the threat level. As such, hunting abstracts dealing with active campaigns and actors that are targeting the organizations sector should have a much high priority. Many other factors can contribute to hunting priority, including: existing detection mechanisms, existing preventative measures, use of TTPs in the wild, etc. The organization should use a risk-driven approach to prioritizing hunting investigations.

After the abstract is created, it is stored on the hunting backlog. This backlog does not need to be a complex tool. Simple collaboration tools such as Microsoft SharePoint or JIRA can suffice. The most important thing is that the backlog provides the hunting team with the required insight to select the most relevant abstract for the next hunt.

Maintenance of the hunting backlog is also required. Priorities will change over time, for example as the popularity of TTPs amongst attackers decreases or standard detection mechanisms are put into place for specific TTPs. Thus, regularly reviewing the backlog will ensure that the selection process for hunting investigations remains effective.

4.3 Phase 2: Hunt

The second phase of the hunt is where the actual investigation takes place. There are 2 activities in this phase. The first activity in the hunting phase is called define / refine. The second activity is called 'execute' and is the actual conduction of the hunt.

4.3.1 Step 3: Define / refine

The 3rd step of the process is called 'define / refine'. 'Define' as the details for the hunt are defined and made more concrete. 'Refine' as these details may need to be changed during the hunt as new evidence is uncovered, or problems are encountered. During this step, the abstract is turned into an investigation by refining and adding information. Some new elements are added, such as required data sources, data analysis techniques and scope. Most importantly, a hypothesis is created that drives the hunt. These activities are described below.

Enrich investigation abstract

The abstract that has been selected is now turned into a concrete hunting investigation. The following information is added to the abstract:

- **MITRE reference.** If possible, references to a technique from the MITRE ATT&CK framework can be added to the abstract.
- **Threat intelligence.** In this stage, potential actors associated with the attack technique under investigation should be identified and their capabilities and motives should be known. This provides the hunters with information on how determined the attacker is, and how well they will likely be able to hide their presence and attack traces. Additionally, this enrichment can be used to find other closely related attack techniques that might be used in the attack (and thus uncovered during the hunt). The MITRE ATT&CK framework [25] can be used this in this process. Additionally, the MITRE ATT&CK navigator [26], can be a useful resource as it associates attack techniques to APT groups. To determine which APT groups are relevant for your sector and organizational type, the APT threat tracking overview is a good starting point [27]. Note that, depending on the hunting investigation, this step may not always be feasible or necessary. Also, this information can be added to the investigation at a later time.
- **Hunt classification.** The hunt can be classified using the cyber kill chain [28] or any other suitable classification method.
- **Required resources.** An estimation of the required resources, such as time spent by hunters, cooperation with other departments, additional technical resources, etc.
- **Refined hypothesis.** The next paragraph will go into the details of hypothesis creation.

Determine hypothesis

Generating the hypothesis that will drive the hunt is an important step in the hunting process. A badly defined hypothesis will likely lead to no results or even worse, wrong results and thus wrong conclusions and recommendations to the organization.

A good hypothesis has the following characteristics [29]:

[25] https://medium.com/mitre-attack/finding-related-att-ck-techniques-f1a4e8dfe2b6

[26] https://mitre.github.io/attack-navigator/enterprise/

[27] http://apt.threattracking.com

[28] https://www.lockheedmartin.com/en-us/capabilities/cyber/cyber-kill-chain.html

[29] C.R. Kothari, Research Methodology, 2nd edition

- **Clear and precise**. This is especially important to avoid discussions on what exactly is being hunted for. Clearly defined hypotheses will also make investigation much easier.
- **Testable**. A hypothesis that cannot be properly investigated will never yield a useful result from the hunting investigation. For example, a hypothesis that deals with the absence of information (for example: "*Attackers are removing evidence of their presence on systems*") is impossible to test properly. You can test for specific signs of evidence removal (clearing audit logs), but if no evidence of such signs is found, the hypothesis can still not be validated. There is no way to conclude that attackers are not performing this activity just because the clear signs are not there.
- **Limited in scope and specific**. A hypothesis must be limited in scope. For example, the hypothesis: "*Attackers are exfiltrating data*" is much too wide. There would not be a clear starting point. A more specific hypothesis would be "*Attackers are exfiltrating data through existing infrastructure*". But even this hypothesis is not properly scoped. The hunt should be executed on a concrete and specific hypothesis, such as "*Attackers are using covert channels based on DNS to exfiltrate data*". This limits the scope of the hunt to DNS traffic specifically.
- **Consistent with most known facts**. Basically, this means that the hypothesis should build on information already available regarding attackers and behavior. Threat intelligence is an important source of facts, where quality of the threat intelligence (as discussed in the previous chapter) is very important.
- **Testable within a reasonable amount of time**. Time is a precious resource. Hunting investigations should be conducted within a reasonable amount of time. An agile approach with time boxing and regular evaluation moments to determine whether or not to continue the investigation is a good way to avoid overspending time on a single hunt. Determining what is 'reasonable' is also something that needs to be discussed with the threat hunting stakeholders during the setup of the threat hunting process.

More on the usage of hypotheses in threat hunting can be found in a SANS paper [[8]] and a blog from RSA [[30]].

Determine data sources

With a clear and precise hypothesis, it should be relatively simple to determine which data sources are required for the hunting investigation. When determining the data sources, it should also be determined if access to the data source is already in place, if the format of the data is usable in its current state and whether or not the data contains sufficient detail to investigate the hypothesis. All data should be collected in a central data lake to allow efficient hunting and analytics capabilities.

[30] https://www.rsa.com/en-us/blog/2017-07/hypothesis-in-threat-hunting

Determine analysis techniques

This step in the process may be difficult at first when starting with hunting investigations. Likely, the hunters will not have a large array of data analysis techniques at their disposal. As the hunting team matures and becomes more capable, hunters will learn additional data analysis techniques. The hypothesis, data sources and scope are the main elements for the decision which data analysis techniques to apply.

As indicated, this may be hard to determine beforehand without having any hunting experience. In such a case, the analysis techniques will be applied 'on-the-fly' during the execution of the hunt in the data analysis step. Even experienced hunters may change analysis techniques during investigations if the need arises. Refinements will take place continuously during hunting investigations.

4.3.2 Step 4: Execute

With define / refine activity completed, the 'execute' activity can be started. During the 'execute' step of the hunt phase, data is retrieved and analyzed. Analysis of data will yield results that can be used in the last step of this phase: validation of the hypothesis.

Retrieve data

In the define / refine step, the data sources and data requirements have been identified. The retrieve activity is not required for all investigations and depends on the data available in the data lake. If missing information is identified, this information should be retrieved from the relevant data sources. Alternatively, the information can also be exported from the data lake into a separate environment for further analysis.

Analyze data

After the data has been collected, it needs to be analyzed to obtain the required results. Existing hunting documentation [[31]][[32]][[33]][[34]] lists a number of data analysis techniques. Some of these technique, such as querying, are simple and easy to perform. Other techniques, such as clustering (of which there are over 100 published possible algorithms [[35]]), are more difficult to understand and require some basic understanding of statistics to use them properly. Some analysis techniques can be applied manually by analysts, while other require some form of

[31] https://www.corvil.com/blog/2015/data-visualization-and-linked-data-analysis-of-electronic-trading-activity

[32] https://www.linkedin.com/pulse/four-common-threat-hunting-techniques-sample-hunts-ely-kahn

[33] https://www.threathunting.net/files/huntpedia.pdf

[34] https://www.blackhat.com/presentations/bh-usa-07/Del_Moral_Talabis/Whitepaper/bh-usa-07-del_moral_talabis-WP.pdf

[35] https://en.wikipedia.org/wiki/Cluster_analysis#Algorithms

machine learning. Hunting platforms that contain analysis techniques and visualizations can be leveraged to simplify analysis.

In the data analysis step, the hunting team may find omissions introduced in the define stage. At this point, the hunters will refine the initial investigation. This is an iterative process that is repeated until the investigation is optimized. Refinements can be done to hypotheses, scope, selected data sources and analysis techniques.

Different levels of sophistication regarding data analysis and data analysis tools exist. An analyst using Excel, especially with pivot tables, is a powerful data analysis tool. Automation in Excel is problematic, as is dealing with large data sets, so additional tools are required. Programming languages such as R and Python (with specific libraries enabled) provide frameworks for data analysis for hunters willing to create their own analysis toolbox. The book 'Data-Driven Security' [36] is a great starting point for data analysis techniques and explains the basics of analytics and the usage of Python (in particular the NumPy and pandas libraries) as well as R. If your organization uses a threat hunting platform, data analytics as well as other techniques such as machine learning will likely be embedded in the platform.

Data analysis is mostly the domain of data scientists. However, data scientists are not threat hunters, just like hunters are not data scientists. If there are data scientists available within the organization, they could be consulted to aid the hunting process.

Example

In our example hypothesis, where the hunting team is investigating data exfiltration through DNS covert channels, some the following analysis techniques can be considered. Note that this is an illustrative example, not a complete overview:

- **Simple querying**. Querying for specific records, such as DNS TXT records, a type of DNS record that is often used in data exfiltration [37].
- **Stack counting**. Determining the total bytes sent and received per source and destination of the DNS query. The top sources and destinations are targets for further analysis.
- **Request / response ratio**. Determine the ratio between the total bytes sent and received per source and destination of the DNS query. Similar to the above, but now focusing on abnormalities in the ratio. With large requests and small responses, the ratio will be quite different from requests and responses of approximately the same size.

[36] J. Jacobs, B. Rudis; Data-Drive Security; 2014

[37] https://www.icann.org/news/blog/what-is-a-dns-covert-channel

- **Statistical analysis**. For example, determine the average DNS request and response size and subsequently search for standard deviations lager than 2 (negative and positive) to identify abnormal requests.
- **Clustering**. Creating time-based clusters to identify bursts of activity within certain time periods. These bursts may indicate moments in time where exfiltration has taken place.
- **Grouping**. Potential exfiltration attempts that have been identified can be grouped for further analysis. For example, that group could consist of servers performing abnormal requests. The hunting team should look for other patterns common to that group, such as common accounts, common connections, etc.

These are just examples of potential analysis techniques that can be used. It is the task of the hunting team to think about potential analysis techniques before engaging in hunting activities (see step 3d) to optimize the hunting investigation.

Threat intelligence integration

When performing data analysis, threat intelligence can be used to add context to investigations. The need for this depends on the investigation. When the threat hunting team finds matches on specific TTPs, further analysis into that TTP must be performed. If possible, this activity should be conducted in collaboration with the threat intelligence team. Such analysis may provide information on possible threat actors, their methods and capabilities, technical infrastructure and other victims of the same actor (these are the 4 features of the diamond model of intrusion analysis [[38]]). This information can subsequently be used to extend the hunting investigation to find additional malicious activity. This provides the hunter with a more complete overview of the compromise that has taken place. This process of ‘pivoting’ has previously been covered in paragraph 3.4.2. Note that the incident response team should be informed of any confirmed or suspected breach. Threat hunters may continue looking for additional compromised systems while the incident response team initiates its containment, eradication and recovery processes.

Validate hypothesis

The final activity of the ‘hunt’ phase is hypothesis validation. When the hunting investigation is finished, the hypothesis must be validated. There are 3 possible outcomes from the validation step:

1. **Hypothesis proven**. The analyzed data provides proof that the hypothesis is true. In this case, a security incident is uncovered.

[38] http://www.activeresponse.org/wp-content/uploads/2013/07/diamond.pdf

2. **Hypothesis disproven**. The analysis of the data provides proof that the hypothesis is false (no security incident). This may prove to be difficult and risky. A threat hunter must convince himself thoroughly that something is truly not there before rejecting a hypothesis. Note that the investigation is still valuable, as other findings and improvements can result from hunting investigations, such as: identification of gaps in detection, development of new detection analytics and recommendations on new preventative measures.
3. **Inconclusive**. After data analysis, there is still insufficient information to either prove or disprove the hypothesis. This could be due to insufficient data, application of incorrect analysis techniques, too narrow scope or quite simply an incorrect hypothesis. The hunter can cycle back to the first step (define / refine) to change some of the parameters of the hunt and repeat the execution

4.4 Phase 3: Finalize

In the last phase, the hunt is finalized and results are handed over to one or more processes.

4.4.1 Step 5: Document findings

The threat hunting team must process the results from the execution step and document findings. This documentation must cover the most important results of the hunt, and the conclusions drawn based on those results. The documentation may also have recommendations. Recommendations may include improvements to preventative measures (from simple configuration changes to architectural changes), recommendations for logging (additional sources, additional details, etc.), recommendations for security monitoring use cases and process recommendations (improvements in vulnerability or configuration management). Finally, the document should have a 'lessons learned' sections that covers how the hunt has helped the hunters to improve. Lessons learned could also be that the hunters have gained valuable insight into parts of the infrastructure. Such insights may ultimately lead to new hunting activities and make subsequent hunts more efficient.

The completed document on the hunting investigation should be disseminated to the appropriate stakeholders. Such stakeholders include SOC managers, risk managers, (chief) information security officers and other teams involved in cyber defense. Note that threat hunting reports may contain sensitive data. Therefore, apply the need-to-know principle where required.

A supporting tool is released together with this methodology: MaGMa for threat hunting. This tool allows hunters to document some aspects their finding. These aspects are then used to create insight into performance of the hunting process. Organizations that are already using MaGMa Use Case Framework will find that it also simplifies the integration between security monitoring and threat hunting. More on this integration can be found in annex B.

Step 6: Update hunting backlog

When the hunt is completed, the hunting backlog is updated. The results are entered into the system, along with the execution data of the hunt and the follow-up. This information can be used in a later stage to determine if a new investigation on the same hypothesis should be conducted.

4.4.2 Step 6: Handover

The final step is handover to other processes. Potential processes that can receive input from the hunting investigation are security incident response, security monitoring, threat intelligence, vulnerability management and others.

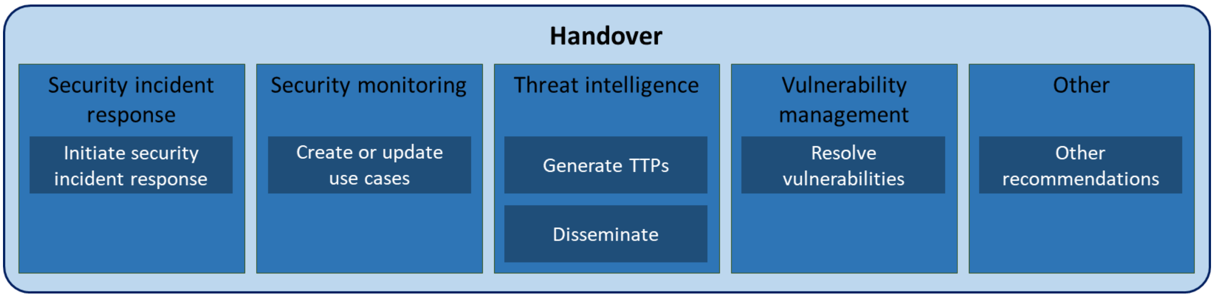

Figure 7: processes triggered by threat hunting investigations

Security incident response

During the analysis step in the threat hunting investigation, indications of malicious activity in the network may be uncovered. In such a case, a hand over must be done to the security incident response process. Threat hunters should not become part of the incident response team, but can support the team by sharing their findings and the analysis techniques used to uncover the attackers. This collaboration between the two teams can prove to be valuable in accelerating the incident response process and quickly containing and eradicating the cause of the incident.

Note that in smaller organizations, there may not be a separate threat hunting team or a separate security incident response team. The person performing hunting activities may also be part of the security incident response when required. Of course, incident response always takes precedence over hunting activities. Being part of both teams can be blocking in achieving a higher maturity level, but it can also be easier to accelerate the incident response as expert knowledge of both processes resides within the same person.

Security monitoring

The threat hunting team will hand over the recommendations for monitoring to the security monitoring team. New use cases may be identified during the hunt that need to be implemented. Additionally, existing use cases may need to be updated or the effectiveness of that use case may need to be increased. These use cases can be based on the findings and detection analytics of the hunters and implemented in the security monitoring tools.

Threat intelligence

The hunt may have uncovered a new TTP that was previously unknown. The findings from the threat hunting team are handed over to the threat intelligence process, where the information is converted into actionable intelligence and subsequently disseminated within the organization as well as outside the organization.

Vulnerability management

A threat hunting investigation may uncover weaknesses in the infrastructure or applications. The threat hunting team can handover these findings to the vulnerability management team to resolve these vulnerabilities before they are exploited.

Other

The threat hunting investigation may yield findings or recommendations for other teams within the organization. For example, recommendations for security architects or teams responsible for hardening the organizations workstations.

4.5 Investigation bias

Any type of investigation can be subjected to bias. Bias can cause hunters to overlook certain evidence and to misinterpret results, leading to wrong conclusions and subsequent actions. Many types of bias exist, including:

- **Confirmation bias**. This type of bias occurs when threat hunters are looking for ways to prove the hypothesis and ignore facts that are inconsistent with the hypothesis.
- **Anchoring bias**. This type of bias occurs when hunters ignore new information and keep focusing on the information they have received before. The **TaHiTI** process is iterative, so hunters should evaluate and refine continuously to avoid this type of bias.
- **Availability bias**. This type of bias occurs when data is investigated based on availability and not based on the data that is required for full insight and thorough investigation.

To avoid bias, it is important to ensure objectivity throughout the investigation as much as possible. A properly defined hypothesis that is clear and concise is a first step in ensuring objectivity. Proper scoping and execution of the 'define' activity is equally important. Keeping an open mind while conducting hunting investigating is another. The dynamics of the threat hunting team play an important role as well. Individual team members should challenge each other when interpreting results. Many resources (some examples: [[39]][[40]][[41]]) are available regarding information on bias and avoiding bias. For very mature teams, advanced techniques such as Analysis of Competing Hypotheses [[42]] can be introduced in certain hunts to increase objectivity. Note that this particular technique is difficult to execute and takes time to do properly.

[39] https://baselinesupport.campuslabs.com/hc/en-us/articles/204305695-Avoiding-bias-in-qualitative-data-analysis

[40] http://info.marshall.usc.edu/faculty/critthink/Supplemental%20Material/Reducing%20Bias.pdf

[41] https://baselinesupport.campuslabs.com/hc/en-us/articles/204305695-Avoiding-bias-in-qualitative-data-analysis

[42] https://www.cia.gov/library/center-for-the-study-of-intelligence/csi-publications/books-and-monographs/psychology-of-intelligence-analysis/art11.html

5 Metrics

Metrics are important to determine the efficiency and effectiveness of the threat hunting process and show its added value to the organization. There are 2 basic types of metrics: quantitative (numbers) and qualitative (value). The focus should be on how threat hunting adds value to the organization, so careful selection of metrics is required.

Several resources, such as the endgame [10] and Sqrrl [43] threat hunting guides and the SANS surveys [1][2] list possible metrics for threat hunting. The following is a list of metrics that are indicators of the value added by the threat hunting process:

- **The dwell time of the findings**: since threat hunting should reduce dwell time (see paragraph 2.2.1), this should be reported for any compromise uncovered in threat hunting. A trend line provides additional information of the progression.
- **Incident response**: number of incidents triggered by the threat hunting process. When traces of attackers are found, a handover is performed to the incident response team.
- **Security monitoring**: number of added and updated use cases. Hunting investigations can lead to new insights about use cases and missing detection mechanisms.
- **Threat intelligence**: new threat intelligence created during the threat hunting process. This is a qualitative metric, as it is hard to express this in numbers that make sense in terms of process quality.
- **Security recommendations**: new preventative measures suggested in threat hunting reports. This is also a qualitative metric.
- **Vulnerability management**: number of vulnerabilities or misconfigurations uncovered. While this is not a primary purpose of threat hunting, these improvements can be side-effects of investigations.
- **Other quality indicators**. For example:
  - Knowledge gained by threat hunters.
  - Visibility gained by the threat hunters.
  - New analysis techniques learned by the threat hunters.
  - New data sources added to the data lake.

When defining metrics for threat hunting, it is important to start out with the goal of the process and then determining useful metrics. While defining metrics that are indications of quality is harder than simply providing numbers, it is well worth the effort. Some of the above metrics have been embedded in the MaGMa for threat hunting tool.

[43] https://sqrrl.com/media/Your-Practical-Guide-to-Threat-Hunting.pdf

6 Best practices for threat hunting

The **TaHiTI** methodology is created with best practices on threat hunting in mind. Some concrete best practices are described hereafter.

Automate where possible

Automation should be applied where possible to make the life of a threat hunter easier and to allow the team to be more productive. While many tasks cannot be automated completely, partial automation to support those tasks should be considered.

Build a threat hunting toolkit

A threat hunting toolkit will help hunters to conduct their activities in a standardized way. The toolbox will likely start out small and grow as required. Some sort of analytics platform and integrations to APIs (for example for threat intelligence and enrichments of findings) are the basis for threat hunting. Many open source tools for threat hunting are available. A good starting point on additional tools can be found on GitHub [[44]].

Cherish your stakeholders

Not every threat hunting investigation will yield very visible results. Because of that, support for the threat hunting team may decrease over time. Ensure that your stakeholders are aware of the purpose and difficulty of the task and that lack of concrete results should not be confused with lack of progress.

Keep track of failed hunts

Threat hunting, especially for teams new to the task, is difficult at first. So even when hunting investigations are not successful for whatever reason, keep track of those failed hunts. Learn and improve from them, and seek to understand why the hunt has failed. Failed hunts are candidates for new hunting investigations once the obstacles that have caused their failure are removed.

Data, data and then some more data

Hunting investigations require data. Threat hunters should have access to a data lake in which relevant data is stored. Data sources such as firewalls, proxies, DNS queries and servers / workstations are vital for successful hunting. As indicated in this document, quality of the data is another factor that is equally important.

[44] https://github.com/topics/threat-hunting

Use a dedicated team

If possible, use a dedicated team of threat hunters that is properly trained and educated. Threat hunting is one of those activities that is often on the lower end of priorities. Thus, there is a risk that hunting investigations are delayed or reduced greatly in scope. If a dedicated team is not possible, a time-boxing approach may be valuable, where team members reserve a dedicated portion of their time for threat hunting activities. This will still mean that hunting investigations take longer, but at least they will be completed.

Work together with other security teams

As shown in the **TaHiTI** process, other processes can trigger the threat hunting process. Vice versa, the threat hunting process can trigger other processes as well and should have a tight integration with threat intelligence. The threat hunting team should cooperate closely with other teams involved in cyber defense for maximum impact. Investing in this relationship will take time in the beginning, but will save time for the team in the long run and benefit the company.

Invest in analysis techniques

As indicated, data analysis is mostly the domain of data scientists. However, less complex data analysis techniques can be used by any threat hunter. Invest in data analysis techniques to build an analysis toolbox for the team. This will allow more efficient analysis. Where possible, ask data scientists to share their insights and recommendations.

Learn from your mistakes and triumphs

Every hunting investigation should have a report. And every report should have a section on lessons learned. But lessons learned do not only apply to elements that require improvement. Successful hunts should be analyzed as well to determine why they were more successful than other hunts.

Share best practices

If you are aware of other teams that are performing threat hunting activities, reach out to them and share your best practices. Even when those teams are less mature and may not be able to offer any valuable insights in return. In due time, they likely will.

Use standards where possible

There are many standards and frameworks available in cyber security. So, do not reinvent the wheel. Build further on the efforts of others and use their insights to improve your own. But as with anything: apply critical thinking. Use what is relevant and discard what is not to keep your focus.

Avoid bias

As indicated in the document, avoiding bias is important to ensure objective and high-quality reports. While avoiding bias may be difficult (as people are prone to bias in general), critical thinking, a formalized methodology and proper team dynamics are essential ingredients of objectivity.

7 Conclusion

In this document, a threat hunting methodology has been introduced. This methodology integrates threat hunting and threat intelligence and provides a clear step-by-step process that hunters can follow to conduct structured hunting investigations. Take these last considerations into account:

1. Carefully select, prioritize and document your input (triggers).
2. Execute hunting with care and apply critical thinking continuously.
3. Use hunting output to drive other security processes and mature and evolve the hunting process itself.

As with any methodology, not all of it may be required for every single hunt. In some cases, hunting investigations will be broad and look at different aspects of a complex TTP. In other cases, hunting investigations may be narrow and scope at a single specific aspect. Some hunts will benefit from a very formal approach, while others may not. Because each hunt is different, investigations will have different requirements. It is up to the organization to apply the methodology in a flexible way that allows the hunters the freedom to hunt in a standardized and efficient manner, without introducing unnecessary overhead. The threat hunting team should consider which elements are required before initiating a hunt, while retaining the flexibility to apply changes where required.

7.1 A final word

This methodology is released under the common criteria with the intention that others can build on it and improve it. The same applies to the accompanying MaGMa for threat hunting tool. So, do not hesitate to provide feedback and do not hesitate to share your own findings and experiences. Creating a cyber security society is not anyone's responsibility, it's everyone's responsibility.

Annex A: Hunting investigation template

General information	
Date	<Date>
Created By	*<Hunter initials>*
Last execution date	*<Date>*
Hypothesis & trigger	
(Initial) Hypothesis	*<Input your initial hypothesis. Will be refined later>*
Hypothesis status	*<Initial / refined>*
Trigger	*<What triggered the creation of this abstract?>*
Reference	*<Reference to the trigger>*
Priority	*<Priority level of the abstract>*
Threat intelligence	
MITRE Reference	*<Reference to attack techniques from MITRE ATT&CK>*
Possible actors	*<Any actors that use these techniques>*
Possible motivations	*<Possible motivations>*
Other TTPs	*<Other TTPs associated with this actor group>*
Active campaign?	*<Is there an active campaign in which these techniques are used?>*
Actor capability	*<High, medium, low>*
Classification & Resources	
Classification	*<Step in the cyber kill chain>*
Estimated resources	*<Rough estimation of time and resources required>*

Annex B: MaGMa for threat hunting

The MaGMa Use Case Framework (UCF) was released in 2017 as a methodology for use case management and a framework for the documentation of use cases. An adaptation of the MaGMa UCF was made to allow for the documentation of threat hunting findings. All the elements mentioned in the 'Hunt' phase of the **TaHiTI** methodology (hypothesis, analysis techniques, scope, data sources) as well as some of the metrics from chapter 5 have been embedded in the tool. By adding hunting results to the tool, insight is created in the performance of the process and focus of threat hunting investigations.

MaGMa for threat hunting uses the same basic setup as MaGMa: an L1 layer based on the cyber kill chain (other threats from MaGMa UCF were removed), and L2 layer that provides a high-level overview of attack techniques related to each step of the cyber kill chain, and an L3 layer for detailed results. For each completed hunt, the hunters add their findings to the L3 layer. The aggregated results are then automatically calculated for L2 and L1, by consistent usage of identifiers in the tool. If a hunting investigation cannot be properly mapped to L2, new attack types can be added. New elements can be added to L1 as well, if desired.

By default, the tool aggregates information at the L1 level and provides quarterly statistics for hunting time spent on each L2 use case (below left) and the trend in average dwell time (below right).

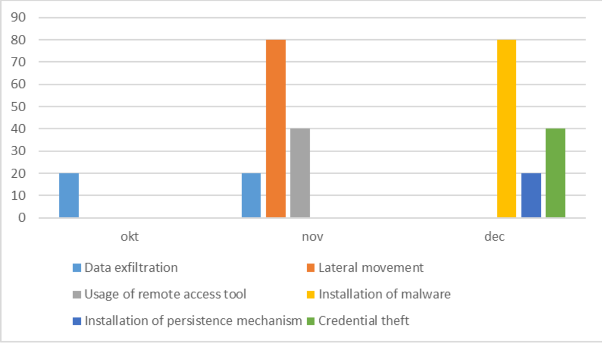

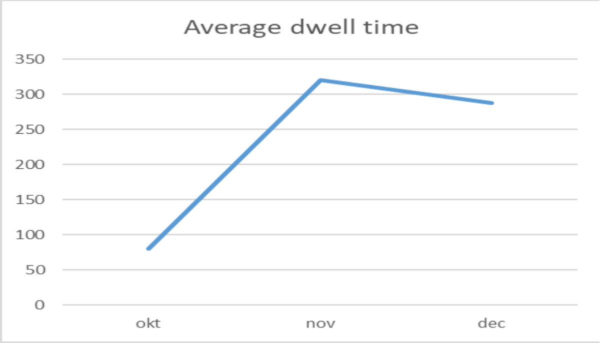

Any useful metrics can easily be added to the tool, or existing metrics can be modified. Also, any other useful statistics and trends can be added to provide the organization with the right information to further increase the performance and maturity of the threat hunting process.

Note that while this tool can be used separately from MaGMa UCF, organizations that are already using MaGMa for their use cases management will be able to more easily integrate threat hunting and security monitoring processes. This is due to the common language between these teams and the fact that the L1 and L2 layers are the same in these tools.

TaHiTI

Threat Hunting Methodology

FI-ISAC NL Publication

©2018