



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"SIEM with Tactical Analytics (Security 555)"
at <http://www.giac.org/registration/gcda>

Applying the Scientific Method to Threat Hunting

GIAC (GCDA) Gold Certification

Author: Jeremy Kerwin, jeremy.kerwin@thalesgroup.com.au

Advisor: Rajat Ravinder Varuni

Accepted: May 28th 2020

Abstract

Threat hunting is a proactive approach to discover attackers within an organization. Without the use of a repeatable framework, the practice of threat hunting is challenging and time-consuming for an analyst. The scientific method, used in fields such as medicine and physics is a repeatable methodology that can be applied to threat hunting to detect threats to an organization.

1. Introduction

Traditionally within Security Operations Centers (SOC), the detection of a threat to an organization is usually reactive. A rule is configured within the Security Information Event Management (SIEM) System to detect a malicious action on a network, and the SOC Analyst waits for an alert to trigger to identify that there may be something malicious present. As attackers become more sophisticated in their attacks and their covert techniques to avoid detection, a SOC's use of the reactive alert approach in attempting to detect malicious actions reliably may not be considered the best approach.

A proactive approach of threat hunting, can be used by SOC Analysts to determine if an attacker compromises an environment. Many within the Information Security community consider threat hunting to be an essential step in detecting adversaries and forms part of a complete security program. (Brenton, 2020)

As a SOC Analyst, you may have decided that you want to start threat hunting, but you're unsure of where and how to begin the process. Threat hunting can be a confusing and challenging process for inexperienced analysts to conduct. Many resources online simply state, "Ask a question, now go and threat hunt." Merely stating, 'Go and threat hunt' doesn't make things easy for analysts who require a structured process to conduct their threat hunts. Developing a hypothesis forms a core part of this process. (Lee & Bianco, 2016, p. xx)

The scientific method is a methodology that has been in existence for hundreds of years within scientific fields like Medicine or Physics. It proposes a series of steps that a researcher can perform to reach a conclusion of a hypothesis based on a series of experiments conducted with certain information available at the time of the test. The scientific method has been studied quite extensively and has a profound philosophical aspect to it. Science Buddies ("Steps of the scientific method," 2012) outlines what practical steps of the method.

The following pages describe what threat hunting is, what the scientific method is and how it can be used as a methodology by members of a SOC threat hunting team.

Jeremy Kerwin, jeremy.kerwin@thalesgroup.com.au

These results can then be replicated and shared to ensure consistent results, remove confirmation bias and increase confidence in the security of a computer network.

2. Threat Hunting and the Scientific Method

2.1. What is threat hunting?

The most basic definition of threat hunting is that it is the proactive approach of searching and finding threats within an organization's network that may go undetected for a long time due to weaknesses in traditional reactive detection systems and techniques.

Threat hunting emerged considered by many in the Information Security industry as a necessary practice within the operations of a security team. Threat hunting complements existing tools and methods to help reduce the detection time of an attack, identify gaps of detection capability within a network, helping to stay ahead of rapidly changing threats and providing a mechanism to respond to these threats quickly. Threat hunting doesn't replace existing systems or processes within a security team; it adds to what is already in place to form a complete security picture of an organization.

One of the goals that threat hunting aims to reduce, or remove, is the perceived confirmation bias by organizations that consider themselves secure because they've implemented a new detection system and haven't seen any alerts generated. Attackers have become increasingly more capable of avoiding reactive detection methods and increasing the amount of time between compromise and detection; this is the 'dwell time.' Threat hunting aims to bring that window of opportunity for an attacker down to as low as possible and prevents them from the chance to return.

The process of threat hunting involves a manual, step by step approach of using the knowledge of threats, attackers and techniques to form a question or 'hypothesis' and then using the information stored within systems like Security Information Event Management (SIEM) systems to attempt to prove that hypothesis correct or not. The threat hunt, or 'hunt' itself can be conducted within a team of many analysts or by a single analyst who shares their results with others for confirmation.

Jeremy Kerwin, jeremy.kerwin@thalesgroup.com.au

For many security analysts, hunting can be a daunting and challenging task to undertake, especially if it is something that they have not experienced. Many analysts are more used to looking at their SIEM and waiting for an alert to trigger so they can then begin their Incident Response Procedures. It also doesn't help that many resources online that describe threat hunting make vague statements like 'Come up with a hypothesis, then hunt for it.'

There are many ways to start the process of developing a hypothesis. The following are merely suggestions on how to gather resources that guide an analyst. Use known techniques, tactics, and procedures (TTP) of known adversary groups as the guiding principles in developing a threat hunting hypothesis.

Determine who or what is the biggest threat to an organization at the beginning of the threat hunt. Is your organization more concerned about nation-state attackers, crime syndicates, or even malicious insiders looking at conducting corporate espionage? Frameworks such as MITRE ATT&CK ("ATT&CK Matrix for Enterprise," n.d.) are a great resource to research what adversary groups might fit your threat profile and then provide information about the known TTPs that this adversary group is known to use to compromise a network.

The information of TTPs combined with the collected data within a SIEM system from various sources within an organization forms the dataset that a 'hunter' uses to conduct the hunt.

Results and conclusions documented during the threat hunting process can be used in subsequent hunts or shared with other organizations or colleagues within the Information Security community to assist them in their threat hunting activities.

Threat hunting always adopts an 'Assume Breach' mentality. It's not a case of if an organization gets breached or compromised but, when. Threat hunting allows a security analyst defender to think like an attacker and ask questions like 'How would I compromise this network?', 'How would I avoid detection?' 'Am I able to detect the attacker?' 'If I can't detect the attacker, what data do I need, or what has to change to allow me to detect them?'

Jeremy Kerwin, jeremy.kerwin@thalesgroup.com.au

2.2. What is the scientific method?

The scientific method is a process of experimentation testing a hypothesis to answer questions. Its use can be traced back for hundreds of years in the advancement of science in fields such as medicine, biology, chemistry, and physics.

It is a process that takes a statement or question about something a researcher has either observed or thought up, applies research and experimentation. Then the results are recorded and reported in scientific papers or studies. The goal of the process is to either prove or disprove the hypothesis.

While the application of the method may differ slightly across different fields, according to an article from Science Buddies ("Steps of the scientific method," 2012) the process consists of the following six actions:

1. Ask a question.
2. Perform background research.
3. Construct a hypothesis.
4. Test the hypothesis through experiments.
5. Analyze the results of the experiments and conclude.
6. Report the results.

The hypothesis, which forms a critical component of the entire method, is a statement that a researcher makes based upon their experience and researches knowledge to seek answers to a question. The question can be broad or very specific to what you are attempting to explore. Being testable and able to produce a result is the mark of a successful hypothesis. Use deductive and inductive reasoning when researching. Deductive reasoning uses a valid premise to reach a correct logical conclusion, while inductive reasoning takes the opposite approach of achieving a false end. The experiments should include an experimental group and a control group. The control group compares against the test. By comparing the results against a known control, you can compare and determine if the experiment was a success or failure.

One of the goals of applying experimentation to the scientific method is to remove or reduce human confirmation bias and instead to rely on empirical evidence in

Jeremy Kerwin, jeremy.kerwin@thalesgroup.com.au

conclusions. Not observing the trust of a situation is an example of confirmation bias. An example of confirmation bias in the information security world is assuming a compromise hasn't occurred because you haven't been alerted to it.

Depending on the results of the experimentation, the original hypothesis may require refinement, alteration, or outright rejection. Through the continual process of applying the scientific method to testing, the results of previous applications of the process improves through the acquisition of new information. In some cases, the results of the scientific method may uncover new and different questions.

A hypothesis that becomes well supported and accepted as correct through the application of the scientific method is referred to as a Theory and comes in the form of Published Studies in Scientific journals. Some well-known theories that have had the scientific method applied to include 'Einstein's Theory of General Relativity,' 'Theory of Evolution,' or 'Theory of Quantum Mechanics.'

Through repeated applications of the scientific method process, the results build up a body of accumulated knowledge. The reproducibility of published studies is a core tenant of the scientific method. If others cannot reproduce a study, it is deemed to be inaccurate and discarded as part of the body of scientific knowledge. Once the process is complete, the results help refine the original hypothesis, feed into the research of a different study, prompt new questions, or disprove it so the researcher can discard that premise and work on a different approach.

2.3. What problem are we attempting to solve?

Threat hunting is a hard, time consuming, and in most cases, a manual process that requires advanced knowledge of an organization's network. Automation helps make the process more comfortable, but it is always a predominately manual process.

Security Operations Centre (SOC) teams are often understaffed and overworked and are always under pressure to detect threats to a network. These threats are continually

evolving, and traditional reactive security tools might only raise an alert well after an adversary has been in place for some time.

Analysts within a SOC team may have different methods of threat hunting, which may result in inconsistent findings or not be able to replicate another team member's hunt results. Different approaches may result in threats not being detected.

To attempt to solve this problem, using a methodology that is repeatable and testable and in use for many years in the scientific field, helps in developing a standard hunting practice within a SOC team without having to re-invent the wheel. A series of repeatable steps that all SOC analysts can follow while threat hunting allows them to share the results and conclusions with other team members to verify the results. A standard process helps achieve a consensus within the SOC team about threat hunts that indicate if an incident has taken place or if it was a false positive, assisting an organization to increase their level on the hunting maturity model.

2.4. Apply the method to a hunt.

Take each step of the scientific method and apply them to the practice of threat hunting. Not all of the steps within the scientific method apply. Much like with Science, the process is flexible to meet the needs of the hunt. The technique isn't rigid and strict, but more of a guidance that if an analyst were to follow as best they can, it gives an excellent foundation to base their activities on in a repeatable manner.

2.4.1. Ask a question.

The first step in a threat hunt based on the scientific method is asking a question. This first step may seem simple, and in many cases a statement as 'Are we breached?' suffices. The item could also be more specific and targeted such as, 'Has a threat actor penetrated our network and are they performing lateral movement across the network using default tools available on Windows 10'. Both types of questions represent their unique challenges in being able to answer them through a hunt.

The security operations team asks the questions. They seek guidance from external sources as management or compliance teams, CERTs, colleagues within the community, news articles, or other organizations.

Not all questions require answers, and it's up to the threat hunter to determine if it is a valid one. Whether the issue is worthy of the time and effort to hunt for is a crucial consideration. Before moving onto the next step, take the time to determine if the question you are asking is required to be more specific. Perhaps it could be broken down into multiple smaller hunts that each yield a result that contributes a conclusion to the original question.

2.4.2. Do background research

Once the question has been asked and deemed valuable to continue the hunt, the next step in the process is to conduct background research. Apart from the actual experimentation or hunt step, this step has the potential to be one of the most time-consuming.

The purpose of performing research is to be able to determine what sources of information may be required, and helps guide the creation of the hypothesis. The hypothesis forms the basis of the threat hunting activity.

Take this opportunity to investigate the capabilities of your existing tools like your Security Information Event Management (SIEM) system. They may already provide you with what you need to complete the hunt.

An analyst can draw upon their own experience as a source for research or refer to findings from previous hunts. Another source for research is published frameworks, commercial and free threat intelligence services. Blogs, news articles, or social media form another source to leverage during the research phase. Discussions with peers within the industry and other organizations are another source to utilize.

The research phase produces several outputs. It might include technical information about how solutions and technologies work or what type of event data is generated and in what form it would look. One crucial output to attempt to produce are

Jeremy Kerwin, jeremy.kerwin@thalesgroup.com.au

the appropriate conditions that need to exist for an event to take place successfully. It is essential to understand because if you don't know what indicators to look for during a hunt, you may never reveal an attack.

2.4.3. Construct a hypothesis

This step can be considered the first phase of the actual threat hunting process. When it comes to constructing a hypothesis, many resources on the Internet make a statement that simply states, 'Construct a hypothesis and then go hunting.' It's a broad statement, doesn't provide any sort of guidance about how to go about doing that, and many people may not know what it means.

A hypothesis is a statement based on observations which is assumed to be accurate based on the outcomes of the first two steps. It is considered to be the first step of the actual threat hunting process and informs the activities of the hunt. The hypothesis guides you in which direction the threat hunt should take and helps to eliminate any non-value paths of approach.

When creating a hypothesis to be used on a threat hunt, it may be broad or specific and should conform to a similar format as "If condition A exists and condition B exists then result C must be true."

Below is an example of a hypothesis for a threat hunt:

"Reports indicate that a nation-state threat actor is using increasingly advanced phishing emails to attempt deployment of a newly discovered variant of the Emotet Malware family. Based on background research, the threat actors appear to be using the same infrastructure as previous campaigns. Indicators of IP addresses and file hashes of malware samples are available. If searches of these indicators against our SIEM and EDR platforms reveal positive results, it could indicate a breach."

Attempt to predict the outcome of the subsequent experiments. Predications must be easy to measure. In the context of a cyber threat hunt, it might be as simple as finding the stated indicators associated with your background research in your environment.

A hypothesis must be falsifiable, meaning that during a hunt, a possible outcome of an experiment is that it conflicts with the prediction. A conflicted result doesn't say that the hypothesis is incorrect. It may merely mean that based on the information at hand and the testing available that this particular hypothesis was not provable by reasonable means.

2.4.4. Testing the hypothesis

The goals of this stage of the process are relatively simple. Do I have the tools and data sets to either prove or disprove my hypothesis? What experiments do I need to conduct to reach that conclusion? Use this phase of the hunt to take the background research to search through your existing SIEM systems.

This stage is the actual hunt. This stage is where an analyst takes all the outputs of the previous steps and starts to go hunting to find (or not find) an attacker in an environment. It tests the robustness of an organizations security systems and reveals any shortcomings in the data collection coverage.

While the goals of this stage might seem simple, the practical application of those goals is challenging due to the time-consuming and challenging nature this step entails. Challenges present themselves to a threat hunter in this stage, and they need to be able to know how to address them and move forward in the hunt.

Firstly, take an inventory of the tools and systems that are available at an analyst's disposal and what their hunting capabilities are. Is the use of any additional tools required?

If you are utilizing a SIEM platform, assess the current status of the platform. Are you ingesting the right forms of data needed on the hunt? Are there any gaps in the monitoring coverage? Do you require to engage with various IT teams to be able to resolve those gaps?

The findings from this step can inform fixes in coverage gaps discovered. Is your SIEM receiving the correct logs from the right systems? Information is used within the

final report to explain why a specific technique may not be able to be detected within the environment because the SIEM system was missing an essential piece of the puzzle.

You may be dealing with a large quantity of data-design experiments with that in mind and attempt to narrow the focus to yield the best results. In the various science fields, you are designing experiments to either prove or disprove your hypothesis. In threat hunting, this step is attempting to do the same.

When designing experiments, reference the Tactics, Techniques, and Procedures (TTPs) that are common to most attackers contained in the MITRE ATT&CK Framework. The Framework provides resources on what detection methods are available and the dataset required within your security systems to be able to detect the technique accurately.

Information discovered from this step, either through using existing SIEM tools or additional tools, can be used to improve existing detection mechanisms. Using manual techniques, tool-based workflows and analytics, a hunter then aims to uncover the specific patterns or anomalies in a threat hunt.

The development of techniques to be reproduced by another member of the hunt team is essential. A reproducible process ensures that the results are consistent over previous hunts. The use of simulation tools such as MITRE's Caldera or Red Canary Atomic Red Team can assist with finding the gaps within detection coverage and then provide the opportunity to correct those gaps.

There isn't a setlist of step by step instructions or point and clicks tests and searches within a SIEM tool that would work with everyone as all organizations are different, with different security systems, different detection capabilities and a different threat profile. It's essential to keep in mind that adopting the concepts and principles from frameworks and what analysts have performed previously help guide your hunting experiments.

Compare the results to a control group not subject to experiments. In threat hunting, knowledge of what would be considered baseline within your environment is the

Jeremy Kerwin, jeremy.kerwin@thalesgroup.com.au

control group. For example, host A talks to host B over Port 12345 on the hour every hour and transfers 200KB of data per communication. If the discovery of that communication during your hunt for C2 traffic, you could conclude that it isn't necessarily malicious because you've identified this communication as usual or as the 'control group' for your environment.

How long should the hunt last? If your organization is mature in its threat hunting practices, lasting no more than a week wouldn't be out of the ordinary. If you're only just starting on your threat hunting program, then expect to plan for a hunt that might last for a few weeks.

2.4.5. Analyze the data and draw a conclusion

The threat hunt is now complete. Compare a conclusion against your original hypothesis? You can never be 100% sure that the outcome is accurate; however, getting as close as possible is the ultimate goal.

With the tools and data at your disposal and the research conducted, was the hunt able to be adequately completed? Was there some piece of the puzzle missing, that if it was available, may have influenced the conclusion reached?

The analysis results of the experimentation (hunting) phase have the potential of generating new indicators and threat intelligence that your existing tools can utilize, potentially shared with others within your team or other organizations to assist them in their threat hunting efforts.

Once you've analyzed the results of the experiments and drawn a conclusion, share those conclusions with other members of the security operations team. Are they able to reproduce the results from your hunt and arrive at the same end?

When using the scientific method in science, the replication of results is crucial. By being able to reproduce the results repeatably, it improves the confidence that the original hypothesis was correct and adds to the whole body of knowledge. The same applies when using the scientific method as a threat hunting methodology. Reviewing the results of a

Jeremy Kerwin, jeremy.kerwin@thalesgroup.com.au

completed hunt by another team member, removes the cognitive bias or blind spots that an individual may have about the results.

2.4.6. Communicate the results. Was the hypothesis correct?

At this final step, it's time to report your outcomes of the threat hunt. When you publish your results, ask the following kinds of questions:

- Who is the audience for the report? What is essential to include? Will the report be just internal to the security operations team, or will there be a broader audience such as management or compliance? Including management in your report findings may help drive organization change that assists with resolving threats.
- Has the threat hunt and analysis of the results concluded that the original hypothesis was correct? What are the next steps?
- If your hypothesis was correct, does that mean you have found an active threat in your environment? Do you need to start any incident response procedures? Has the hunt found something malicious?

If, during the hunt and the analysis of the results, it shows that the hypothesis stated at the start was incorrect, it doesn't mean that the effort was not worth it. Use the limitations and negative results of the hunt to drive input into any subsequent pursuits. Did you find you were missing a specific dataset? Perhaps a takeaway is to work on how to get that dataset available for following hunts.

2.5. Example hunt using the scientific method.

Alice is a Security Analyst working in the Security Operations Centre (SOC) of a large corporation. Her daily tasks usually involve monitoring the SIEM for any alerts that might indicate malicious activity.

Her manager recently tasked Alice with conducting threat hunting activities. By applying the scientific method in her threat hunt, it might look like the following.

2.5.1. Ask a question

Alice had been reading about in the news about web shells being a popular method for attackers to compromise systems. Alice asks the question:

"Are there web shells on our external web servers?"

2.5.2. Background research

The first step that Alice takes in her research is to perform an Internet search for the search term 'web shell.' One of the first results returned is a security advisory from the Australian Cyber Security Centre (ACSC) on web shells ("Detect and prevent web shell malware," 2020). This advisory was a joint publication from the Australian Signals Directorate (ASD) and the United States National Security Agency (NSA).

The advisory gives an overview of what web shells are and why they are considered a severe threat. The attachment to the bulletin goes into details of detection methods to detect web shells.

Alice also consults the MITRE ATT&CK framework on web shells ("Web shell, technique T1100 - Enterprise | MITRE ATT&CK®," 2019) for additional information, including what threat actors use web shells to target the organization's in the same industry as Alice's.

To assist the detection of a web shell, Alice concludes requiring several data sources, either in the SIEM or through other means. A known good copy of the web server file system to compare is required; this could come from a known good backup, an installer package, or a cloned server. Access logs from the webserver show a history of who has accessed the file server Firewall and network traffic logs show the network communication that has occurred between the web server and client. Endpoint logs from client workstations provide details on any suspicious processes that are attempting communications to the web server.

Alice confirms that all data sources required are available, except for the endpoint logs from client workstations.

2.5.3. Hypothesis

Alice proposes the following hypothesis

"The external web server may be compromised by a web shell. Comparing the directory structure of the web application to a known good copy may reveal a web shell installed. Analyzing any connection attempts to this server could also reveal evidence of a web shell."

2.5.4. Experiment

Using the information from the completed research, Alice compares the web files on the external against a known image of the web server using the example scripts that came with the ASD, NSA guidance on web shells. Also, Alice conducts a series of searches within the SIEM platform to analyze the connections attempts made to the web server to determine if any logs look suspicious.

2.5.5. Analyze data and conclusion

The results of the first experiment, comparing the file directory to a known image, returned that all there was one discrepancy on the public web server with a single PHP file added to the list. Upon inspecting the contents of this rogue file, it contains the following PHP code. `<?php @eval($_POST['password']);>`

Alice refers to the research and finds that according to the MITRE ATT&CK framework web shell technique, it might indicate the presence of the 'China Chopper' web shell, a popular web shell among many threat actors.

The results of analyzing the web server traffic logs in the SIEM platform reveals no suspicious connections made to the web server.

Alice draws the following conclusion:

"An unknown attacker has compromised the external web server and placed a web shell onto the file system. There does not appear to be any suspicious connection attempts made to the web server. The lack of suspicious network traffic may indicate no active use for exploitation."

2.5.6. Report the results.

Alice writes up the results of her hunt activities, concludes that her hypothesis was correct, and documents her findings into a report that she delivers to the rest of her team to review.

Her colleague, Bob, decides to replicate the experiments that Alice conducted to determine if he gets the same results. Bob runs the same file system comparison that Alice ran and received the same discrepancy that Alice found. Bob also did not find any suspicious connection attempts in the web server access logs.

Since both Bob and Alice reached the same conclusion, it is reasonable to assume at this point with the evidence presently available that the hypothesis was indeed correct and found a potentially compromised host.

Alice can now initiate the Security Operations Centre (SOC) Incident Response procedure and contain the compromised host.

3. Conclusion

The scientific method is a philosophy based upon a testable, repeatable process that produces results that confirm or deny a hypothesis. Threat hunting also relies on a testable, repeatable process to confirm or deny the existence of an attacker on a network. Threat hunting is challenging, but the benefits have far-reaching consequences for an organization. The structure and repeatability of the scientific method, can be applied to threat hunting to achieve the goal of making the process a little less challenging for an analyst. There are many resources available online to assist the structure of a hunt (Gunter & Seitz, 2020) that take a similar approach to the scientific method.

References

- ATT&CK Matrix for Enterprise. (n.d.). MITRE ATT&CK®. <https://attack.mitre.org>
- Brenton, C. (2020, March 10). What Is Threat Hunting and Why Is It So Important? Active Countermeasures. <https://www.activecountermeasures.com/what-is-threat-hunting-and-why-is-it-so-important-video-blog/>
- Detect and prevent web shell malware. (2020, April 23). Australian Cyber Security Centre (ACSC). <https://www.cyber.gov.au/advice/detect-and-prevent-web-shell-malware>
- Gunter, D., & Seitz, M. (2020). A Practical Model for Conducting Cyber Threat Hunting. <https://www.sans.org/reading-room/whitepapers/threathunting/practical-model-conducting-cyber-threat-hunting-38710>
- Lee, R. M., & Bianco, D. (2016). Generating Hypotheses for Successful Threat Hunting, 13. <https://www.sans.org/reading-room/whitepapers/threats/generating-hypotheses-successful-threat-hunting-37172>
- Steps of the scientific method. (2012, June 7). Science Buddies. <https://www.sciencebuddies.org/science-fair-projects/science-fair/steps-of-the-scientific-method>
- Web shell, technique T1100 - Enterprise | MITRE ATT&CK®. (2019). MITRE ATT&CK®. <https://attack.mitre.org/techniques/T1100/>