



SANS Institute

Information Security Reading Room

Hunting with Rigor: Quantifying the Breadth, Depth and Threat Intelligence Coverage of a Threat Hunt in Industrial Control System Environments

Dan Gunter

Copyright SANS Institute 2021. Author Retains Full Rights.

This paper is from the SANS Institute Reading Room site. Reposting is not permitted without express written permission.

Hunting with Rigor: Quantifying the Breadth, Depth and Threat Intelligence Coverage of a Threat Hunt in Industrial Control System Environments

GIAC (GSEC) Gold Certification

Author: Dan Gunter, dangunter@gmail.com

Advisor: Dr Johannes Ullrich

Accepted: July 6th 2018

Abstract

Threat hunting provides an organization a proactive opportunity to discover hidden attackers and to evaluate and improve the security posture of the environment. While existing research focuses on technical methods for threat hunting, a way to assess the rigor and completeness of threat hunting activities remains unexplored. This research examines several methods that can be implemented/used to calculate coverage of threat hunts. Coverage calculation methods include kill chain coverage, attacker tactic, technique and procedure coverage and threat intelligence coverage. This research also explores how to automate the calculation of threat hunt coverage. By following the process outlined by this research, analysts can ensure that planned threat hunts remain relevant to the overall goal of the hunt and that these hunts can maximize the chance of adversary detection success.

1. Introduction

The popularity of threat hunting as a form of proactive and reactive security has grown over the past few years. Threat hunting "is a focused and iterative approach to searching out, identifying and understanding adversaries that have entered the defender's networks" (Lee & Lee, 2017). While threat hunting continues to grow as an emerging trend, the corpus of knowledge remains sparse. To date, studies into threat hunting methodology have focused on the definition of hunting, the maturity of data collection programs, and specific approaches for hunting.

An area yet to be explored is the analytic depth and breadth of threat hunting. When organizations test software, a key metric gathered is code coverage. Code coverage is related to a measured percentage of the tested application's source code and is gathered/used to quantify software test coverage. The focus of this research is to provide automated methods to quantify the coverage of threat hunts in regard to the observed environment as well as to quantify threat intelligence-derived Tactics, Techniques and Procedures (TTP). The end goal is to help organizations quantify the coverage of hunting efforts, understand what an existing hunting program currently focuses on and areas where additional diversification might be required. The goal of this research is to provide several methods that analysts can use to evaluate a threat hunt.

2. Types of Threat Hunts

When considering the overall type of a threat hunt, the end goal plays a significant role in the overall classification of the type of the hunt. Hunt engagements start with a well-defined goal to uncover specific actors—which categorizes this type of hunt as a threat-focused. Similarly, an environmental hunt engagement focused on studying a particular subset of the overall environment from a purely technical perspective. Classifying the type of a threat hunt is as essential to rigor as the type of threat hunt changes the TTP and data sources required to conduct the hunt. A hunt might also start as an environmental hunt and spiral into a threat-focused hunt should any malicious activity be discovered.

Author Name. email@address

Both types of hunts are essential to an organization to maintain a comprehensive threat hunting program. Threat-focused hunts capitalize on known adversary behavior and can verify the presence of known attacker TTP in the environment. Environmental-focused threat hunts might choose a specific protocol or observable source and look for malicious behaviors not yet associated with attacker TTP. This research categorizes generic hunts focused on TTP not associated with a known attacker as being part of an environment-focused hunt. While the hunt does look at data related to a specific attack TTP, no context exists surrounding a known malicious actor. When hunting for sophisticated attackers, a threat-focused hunt that incorporates known intelligence about a specific attacker might be followed by an environmental hunt to look for potential progression in attacker TTP or unknown attacker TTP.

3. Elements of an Effective Threat Hunt

3.1. Threat Hunting Playbooks

A threat hunting playbook is a series of objective-driven tasks that lead an analyst through a particular analytic workflow. In the purest form, a playbook provides an analyst with a checklist of tasks to follow. Within the context of a threat-focused hunt, a hunting playbook might focus the threat hunter on very specific observables related to known attacker TTP. For environment-concentrated hunts, a hunting playbook might have a broader scope to uncover malicious activity within a more extensive set of data. Threat hunting playbooks might follow a format similar to incident response playbooks (Lamis, 2010). Incident response procedure provides pre-tested actions that enable responders to quickly neutralize attackers within the network (NIST, 2016). Threat hunt playbooks offer pre-tested actions for adversary discovery.

Rigor, in the context of threat hunting, can be defined as the degree of analytic thoroughness to achieve the defined end goal. Relevant to the overall rigor of the hunt, threat hunting playbooks ensure that threat hunts are repeatable and comprehensive. If an analyst hunts without a playbook, there is no guarantee that the analysis covers particular areas of interest. The playbook ensures the integrity of the hunt. A playbook should not, however, constrain the creativity and analytic mind of the threat hunter. While an analyst

should complete all steps of a given playbook to ensure the integrity and rigor of the hunt, the analyst should also be encouraged to explore beyond the playbook tasks. If an analyst discovers a novel or successful approach beyond the defined scope of the playbook, the analyst should be encouraged to update the playbook with the new approach.

3.2. Threat Intelligence

Threat intelligence provides one source of context for scoping the focus of a threat hunt through the study of a specific attacker's TTP. Sergio Caltagirone, Andrew Pendergast, and Christopher Getz proposed an intrusion activity model, termed the diamond model, in their article, "The Diamond Model of Intrusion Analysis" (2013). Events serve as a central component of the diamond model and capture the use of a capability or capabilities by an adversary over a given infrastructure against a victim. The four major components, adversary, capability, infrastructure, and victim enable analysts to develop a comprehensive picture of a given intrusion. While an individual diamond may represent a single adversary action, Caltagirone, Pendergast and Betz's research proposed the concept of an activity thread mapped to the cyber kill chain that represents both phases of a single attack as well as activities against other victims (Caltagirone, Pendergast, & Betz, 2013). The diamond model provides threat hunters with an approach to understand current activity groups and attacker TTP. By using the diamond model and produced intelligence, threat hunters can plan targeted threat hunts specifically to known attacker TTP.

3.3. Rigor

The calculation of rigor varies slightly between an environment-focused hunt and a threat-actor focused hunt. Within an environment hunt, rigor calculates the overall coverage of a planned hunt relative to all systems in the environment, the coverage of a hunt compared to assets critical to operation, or the coverage of a hunt compared to available data. For a threat-driven hunt, rigor calculates the coverage of a hunt against protocols an attacker is known to abuse, the overall relevance of threat intelligence to the environment, or the usefulness of collected information against potential observables needed to detect an adversary. With an environment-driven hunt, the rigor of the

conducted hunt compared to the overall environment while a threat-driven hunt calculates rigor of the hunt against known adversary TTP.

4. Modeling Attacker and Threat Hunting TTPs

A necessary part of calculating threat hunt coverage analysis requires analysts to translate attacker and threat hunter TTP into an analysis model. The following section will demonstrate a modeling approach for attacker and threat hunter action.

4.1. Modeling Attacker Actions

4.1.1. Tracking Actions Across Attack Stages

For an attack to be successful, an attacker must conduct a series of actions to prepare for the attack, gain a position within the target environment and conduct an action to achieve the end goal. Analysts often map the spectrum of necessary adversary actions within industrial environments to the ICS cyber kill chain. Michael Assante and Robert M. Lee introduced the ICS cyber kill chain model in a white paper titled, “The Industrial Control System Cyber Kill Chain” (Assante & Lee, 2015). This model is depicted below in Figure 1.

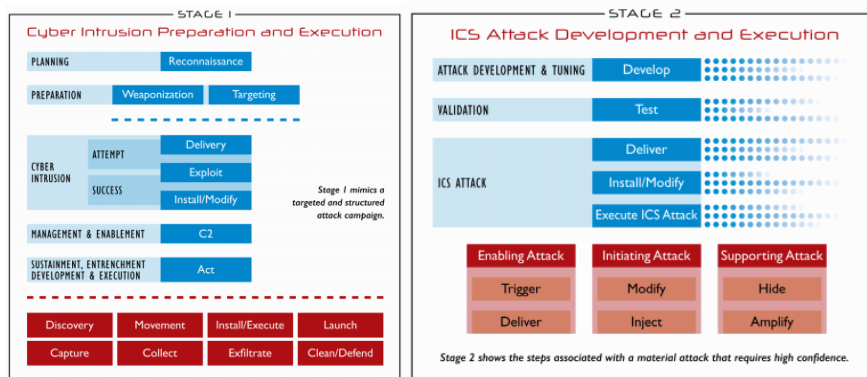


Figure 1: ICS Cyber Kill Chain Model

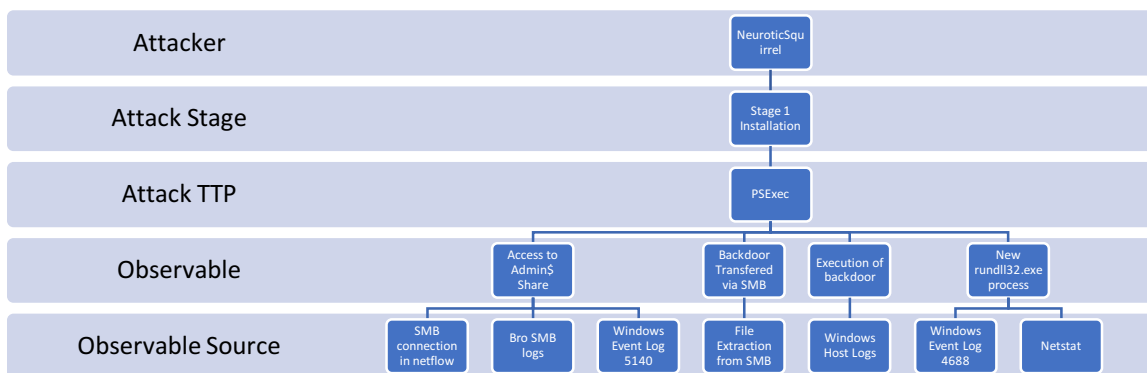
Stage one of the ICS cyber kill chain covers all of the reconnaissance, weaponization, delivery, exploitation and command and control associated with pivoting through the corporate IT network of a targeted ICS company. Additionally, phase one covers the pivot from the corporate IT network into the ICS or OT portions of the target network. An attacker enters phase two of the ICS kill chain when the attacker begins to

“specifically develop and test a capability that can meaningfully attack the ICS” (Assante & Lee, 2015).

Attackers employ TTP at each phase of the ICS kill chain to achieve the goals of the phase they are in and to prepare for the next phase. An attacker performing stage one reconnaissance might use nmap to port scan portions of a target network while an attacker within the install/modify stage of stage two might install a custom tool to hijack a valid communication process on an ICS machine.

4.1.2. Translating Kill Chain TTPs into Observable Characteristics

Understanding attacker TTP is essential but worthless unless the TTP is distilled into observables that threat hunters can discover/detect/etc. during a threat hunt. While Newton’s third law does not apply to network or host phenomena, this research proposes a similar corollary relevant to network and host phenomena that states, “for every attacker action there is a manifestation of the attacker’s action realized in network and host logs.” Consider an example where threat intelligence exists that states/support that a fictitious actor termed NeuroticSquirrel generally uses stolen user credentials and Metasploit’s PSEXec module to access a target machine to deliver malware remotely. This information alone is sufficient enough to generate observable characteristics. Knowing what observable characteristics an adversary action yields does, however, require domain knowledge. PSEXec is a remote Windows administration utility designed by Mark Russinovich that uses a Windows Server Message Block (SMB) file share to connect to the target machine using a share named Admin\$ (Maloney, 2013). Analysts can map observables tied to NeuroticSquirrel’s preferred use of Metasploit’s PSEXec module to deliver malware using the following model:



Author Name. email@address

Figure 2: NeuroticSquirrel TTP Observable Model

This research produced the model above by conducting technical analysis into the function of PSEXec. As shown in figure 2 above, the threat intelligence provided covered the attacker, attack stage and TTP portion of this model. Observables can be derived from understanding what the TTP does and the observable tier comes from understanding how the TTP operates on the network and host level. One of the methods from Metasploit's PSEXec module works by using the provided credentials to access the Admin\$ share, uploading an executable to the Admin\$ share, and creating a new process named rundll32.exe that eventually is injected with the attacker's shellcode. The outlined behavior is a simplified overview of PSEXec, but these steps provide an initial starting point for mapping observables to adversary TTP. Each step is an observable because host or network sources can be associated to validate the existence of them. Observables can be categorized as either host or network. This distinction is important when rigor calculations for observables come into play to indicate potential bias in host centric or network centric hunt approaches. Some observable categories will be empty if an observable has no host or network indicators. For example, a new process starting is not visible via network traffic. In this scenario, it is essential to understand that it is crucial to have data across both categories when possible. Finally, the observable sources are the actual host or network artifacts that have been analyzed to prove or disprove the presence of the observable in the environment.

Some observables might not be available due to operating system version while other observables might be overwritten by the attacker when the attacker attempts to cover tracks. Having a variety of observable sources increases the overall chance of observing a TTP. Additionally, attackers might utilize new TTP in subsequent attacks. While some observables might not be present in future attacks, an attacker would have to change all TTP for no observables to be present.

4.2. Modeling Threat Hunting Engagements

Similar to attack TTP, analysts can also model threat hunting TTP. The ICS kill chain is equally useful to threat hunters who model threat hunting TTP. Threat hunt TTP that focuses on countering a known adversary looks very similar to the related attacker

model regarding the observables a threat hunter plans. The similarity occurs because a hunt is essentially a focused effort to uncover the observables created by the attacker through the various kill chain steps.

A given threat hunt engagement consists of a set of hunting TTP. Each hunting TTP and playbook works across a variety of sources in both host and network traffic. A single TTP or playbook might only look across both host and network traffic or focus on one observable category. Within the observable categories, the chosen TTP uses one or many log types to prove or disprove a given hypothesis. A hypothesis is an analytic question made by the threat hunter targeted at uncovering adversary action that can be proven or disproven. Figure 3 below outlines the model for threat hunting engagement using a hypothetical hunt for the fake NeuroticSquirrel activity group.

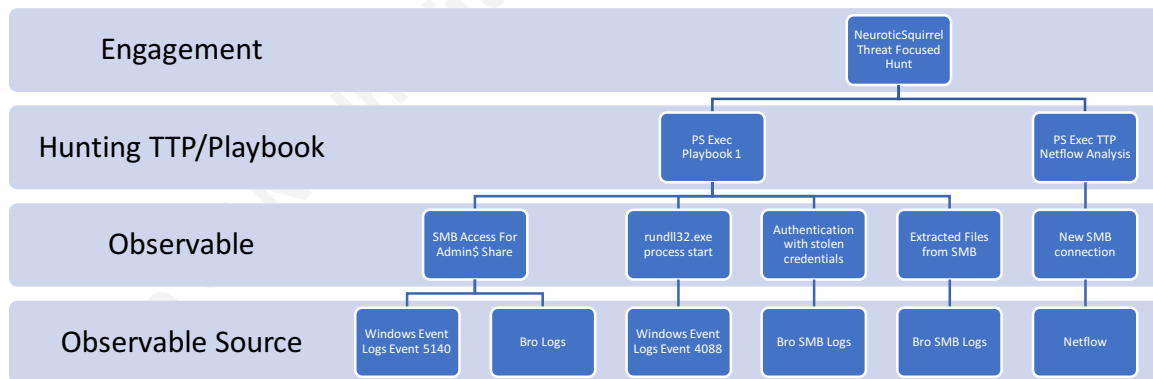


Figure 3: Hunting TTP Observable Model

In this particular example, the analysis focused on a threat-actor-focused hunt for NeuroticSquirrel's use of PSEXec. This research uses a threat hunting TTP that consists of a PSEXec playbook and a homegrown tool for analyzing NetFlow data. The PSEXec playbook looks across both host and network traffic. On the host side, the PSEXec playbook involves analysis of Windows event logs covering both SMB access to the Admin\$ share and referring to new processes creation related to the PSEXec process starting. On the network side, the PSEXec playbook looks at authentication logs from Bro IDS and searches for extracted binaries from SMB data. The planned hunt in this case study also includes a PSEXec TTP NetFlow analysis tool focused on the discovery of new

SMB connections between hosts. Calculation of the rigor of a hunt requires analysis of precisely what observables a hunt must include to be successful. While vendor products and open source tools can assist with/during the process of hunting, it is necessary to understand the exact capabilities and limits of the chosen tool.

4.3. Quantifying Rigor and Return on Investment

4.3.1. Coverage of Chosen Hunting TTP vs Known Attacker TTP

Under the attacker and threat hunter taxonomies presented in this research, the success condition of a hunt looks at how many of the known observables in the planned hunt coincide with a given attacker TTP. The Venn diagram below contains significant overlap between the attacker TTP observables and the observables analyzed in a given hunt.

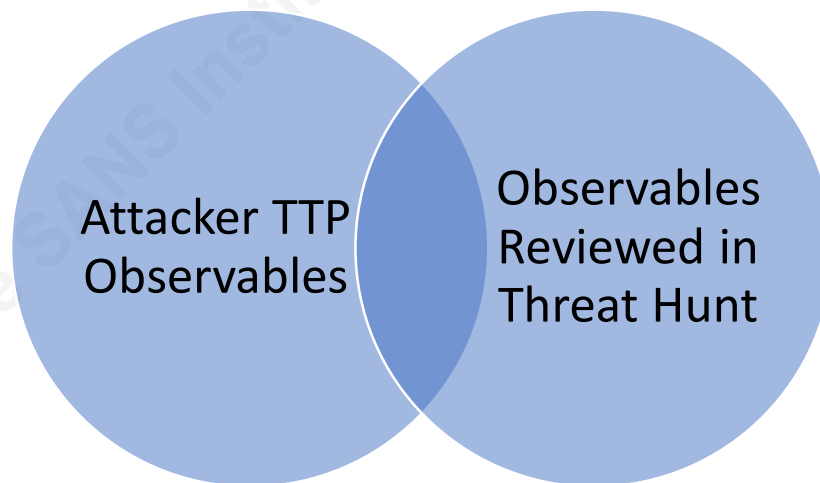


Figure 4: Overlap of Attacker TTP vs Observables Reviewed in Hunt

As defenders better understand malware, the number of attacker TTP observables will increase. The coverage calculation for subsequent threat hunts should account for the new subset of observables. For a threat hunt to reach the same calculation coverage, the threat hunt will need to include analytical techniques that cover the new observables.

4.3.2. Calculating Usage of Collected Data and Valuable Missing Data

At the collection level, analysis of data collection rigor includes the calculation of available data sources against the list of known data sources related to an observable. Consider the NeuroticSquirrel example where the attacker leverages PSEXEC.

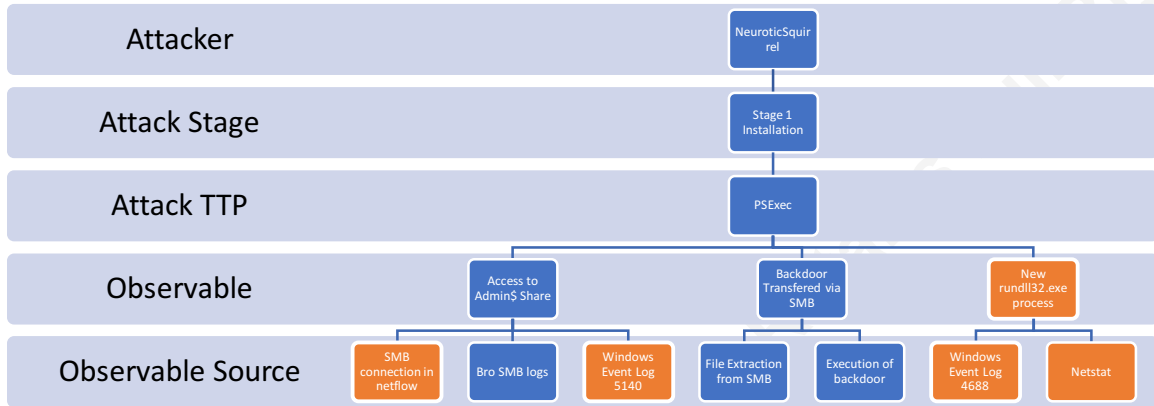


Figure 5: Observables Collected During Hunt

The orange boxes above correspond to observable sources not collected by the organization. Four out of seven, or 57% of, observable sources related to finding NeuroticSquirrel’s variant of PSEXec are not available for hunting. Additionally, the absence of both Windows Event Logs and netstat output has eliminated the ability to observe the start of the new rundll32.exe process.

4.3.3. Calculating Threat Intelligence Source Return on Investment

Threat intelligence serves a critical role in informing threat hunts. Another coverage calculation should include the return on investment (ROI) of threat intelligence sources. In the context of observables, threat intelligence should inform the definition of observables for an adversary’s attack TTP. Organizations can, therefore, track which threat intelligence sources have been most valuable to hunting efforts.

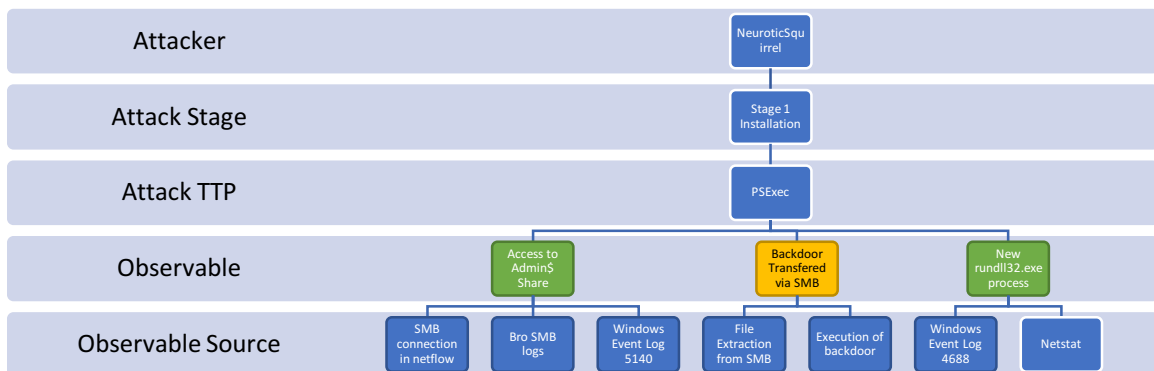


Figure 6: Intelligence Source Coverage in Threat Hunt

The diagram of the PSExec attack TTP depicts data from both threat intelligence source one (shown in green) and threat intelligence source two (shown in green). Threat intelligence source one accounts for 67% of knowledge about known NeuroticSquirrel's attack TTP while threat intelligence source two accounts for 33% of known attack TTP.

Threat intelligence return on investment should not stop at mere coverage calculation. As threat hunts are successful, return on investment for intelligence sources should keep track of which intelligence sources led to the discovery of the attacker. The discovery calculations should include all available threat intelligence sources that supported the discovery of the adversary. As a model of the attacker TTP continues to grow with more observables, the return on investment calculation will also expand.

4.3.4. Calculating Cyber Kill Chain and ATT&CK Framework Coverage

Lockheed Martin's Cyber Kill Chain and MITRE's ATT&CK matrix provide two additional models that are useful for rating the rigor of a threat hunt. A threat hunter might consider which of the seven steps a given attack-focused observable targets in Lockheed Martin's Cyber Kill Chain. Comprehensive threat intelligence will ideally provide knowledge of attack capabilities across as many kill chain steps as possible. Consider the following coverage for the previous NeuroticSquirrel hunt example in Figure 7 below.

Kill Chain Step	Threat Intelligence Available	Relevant Hunt TTP
Reconnaissance	Spear Phishing	
Weaponization	Metasploit Adobe PDF Embedded EXE module (CVE-2010-1240)	
Delivery	Spear phishing	
Exploitation	Metasploit Adobe PDF Embedded EXE module (CVE-2010-1240)	

Author Name. email@address

Installation	PSEXEC against publicly exposed SMB for delivery	Bro IDS SMB Log Hunt Windows Event Log Hunt Netflow Hunt
Command & Control	Unknown	
Action on Objectives	Unknown	

Source: https://www.lockheedmartin.com/content/dam/lockheed-martin/rms/documents/cyber/Gaining_the_Advantage_Cyber_Kill_Chain.pdf

Figure 7: Kill Chain & ATT&CK Coverage Table

Thus far, the planned hunt for NeuroticSquirrel only covers one stage of the Lockheed Martin Cyber Kill Chain. The table above shows other available threat intelligence that might support a targeted hunt against other kill chain stages. At the basic level, the planned hunt only covers 20% of the Cyber Kill Chain steps with known intelligence about attackers. This primary coverage can be useful when looking for overall trends between hunting engagements to see what analysis areas an organization favors across a series of hunts. Additionally, threat hunting teams might look at where the most adversary detection tends to take place. Past success should not preclude future hunts from looking in other areas but might indicate current analytic areas of strength and weakness.

MITRE's ATT&CK framework can also be used to provide more granularity for the attacker TTP targeted by the planned threat hunt. Similar to Lockheed Martin's model, coverage on MITRE's model looks at general coverage over the tactic and technique matrix regarding the planned hunt as well as where threat hunt success has occurred in the past.

5. Assessing Rigor and Coverage of Hunting Efforts

Under the attacker and threat hunter taxonomies presented in this research, the success condition of a hunt involves a high calculated overlap between observables studied by chosen hunt techniques compared to known attack observables associated with a given attacker TTP. The proposed methodology of this research employs JavaScript Object Notation (JSON) to represent the various components of attacker TTP and threat

Author Name. email@address

hunt TTP. The following subsections will examine the structure of each component and the relationship to other components for both attacker TTP and threat hunt TTP.

5.1. Modeling Attacker Context

Four JSON data structures were used to represent the capabilities of adversary groups. The data structures represented adversary group, attack TTP, observables, and sources. Each JSON data structure supports additional metadata to enable other rigor calculations. Diagram 8 below shows the hierarchy between the four JSON structures associated with the attacker.

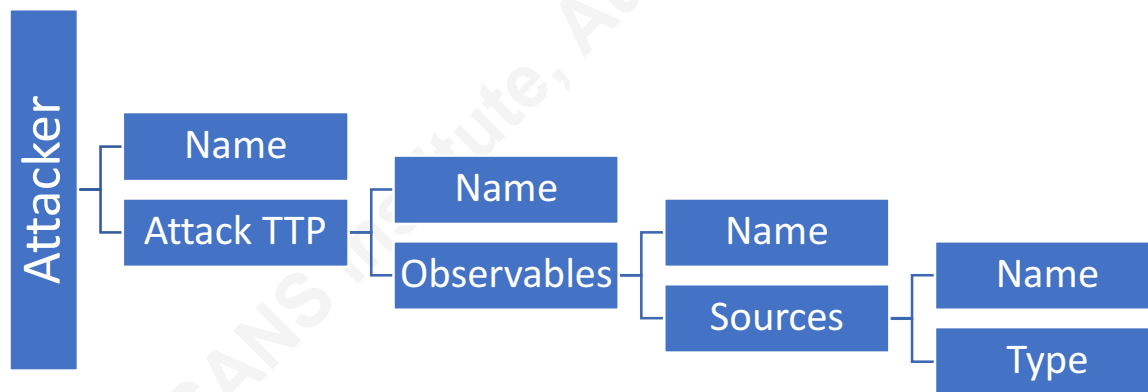


Figure 8: Attacker Data Structure Representation

5.1.1. Activity Group Representation

The root of the attacker JSON model consists of a name and a list of attack TTP. As threat intelligence yields further attack TTP, the attack TTP list expands with the new TTP. Attack TTP common between attacker groups will exist in the attack TTP list for both groups. For the NeuroticSquirrel example, only one attack TTP existed around the known use of PSEXec.

```

{
  "_id": "NeuroticSquirrel",
  "attack_otp": [
    "PSEXec (Metasploit)"
  ]
}
  
```

Figure 9: Activity Group Representation in JSON

5.1.2. Attack TTP Representation

The attack TTP data structure tracks known attacker tools and techniques, associated cyber kill chain or ATT&CK framework phases, and the corresponding observables. Updates to attack tools might lead to multiple attack TTP data structures as development continues on a particular attack tool or as defenders better understand how an attack tool works. The id field below contains the name of the attack TTP as seen in the activity group data structure. The observables list contains the different breadcrumbs the attacker generates in host or network logs through the execution of the TTP.

```
{
  "_id": "PSEXEC (Metasploit)",
  "kill chain stages": [
    "Stage 1 Installation"
  ],
  "ATT&CK framework techniques": [
    "Service Execution"
  ],
  "observables": [
    "Access to Admin$ Share",
    "Backdoor Transferred via SMB",
    "New rundll32.exe Process"
  ]
}
```

Figure 10: Attack TTP Representation in JSON

5.1.3. Observables Representation

Observables consist of a name and a list of possible data sources relevant to the attack TTP. The observable id corresponds to items in the attack TTP observables list. As defenders better understand the parent attack TTP, the observable list will grow with new detection opportunities. A commonality between one or more attack TTP representations can occur at this level. The list of observable sources should be as exhaustive as possible to account for as many potential observation areas for the attack TTP. The observable list allows organizations to determine what other potential data sources might be of value to detect adversary TTP.

```
{
  "_id": "New rundll32.exe Process",
  "observable source": [
```

```

    "Windows Event Log 4688",
    "Netstat"
  ]
}

```

Figure 11: Attack Observable Representation in JSON

5.1.4. Observable Sources Representation

The lowest data structure used to model attack TTP represents the observable sources for a given attack. The observable source data structure consists of the name of the observable source and metadata about if the source derives from the host or network information.

```

{
  "_id": "Windows Event Log 4688",
  "observable type": "host"
}

```

Figure 12: Attack Observable Source Representation in JSON

5.2. Modeling Defender Context

Playbooks serve as the root of the defender data structure. The playbook data structure contains a name, a set of steps for the analyst to follow and the observables associated with the playbook. The set of observables should update with the addition and removal of playbook steps. Observables might also need to be updated as playbook steps utilize new data sets.

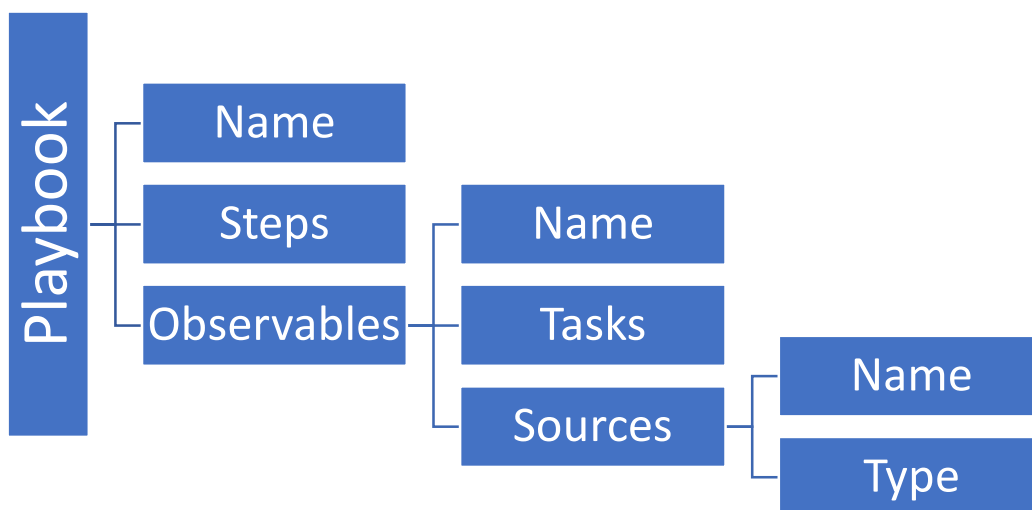


Figure 13: Playbook Data Structure Representation Heirarchy

The following JSON notation represents the playbook data structure. The id field contains the name of the playbook, while the steps list contains the list of steps that the analyst should follow during the threat hunt. The observables relate to the discovery potential attack TTP through the execution of the playbook. Note that the observables listed for this playbook are the same as the attack TTP representation observables.

```
{
  "_id": "Playbook for Metasploit PSEXec",
  "steps": [
    "Check for access to Admin$ share or other SMB shares",
    "Check Bro file extractions for transferred backdoors",
    "Check Windows Event Log record 4688 for unusual process creation"
  ]
  "observables": [
    "Access to Admin$ Share",
    "Backdoor Transferred via SMB",
    "New rundll32.exe Process"
  ]
}
```

Figure 14: Playbook Representation in JSON

5.3. Calculating Coverage

The JSON structures outlined in the previous section provide a basis for calculating analytic coverage. The relevance of a threat hunt playbook to an attacker TTP or an overall attacker can be identified/analyzed etc. by comparing playbook observables to the attacker TTP or attacker data structure. An organization can calculate the impact and quality of specific threat intelligence sources by tracking the threat intelligence source associated with the observable that led to attacker discovery. The outline approach affords the benefit of either adding metadata to the JSON structures or adding new JSON structures to represent both attack TTP and threat hunts.

6. Conclusion

The approach outlined by this research provides one method for calculating the rigor of a threat hunt using the concept of observables. Observables refer to the breadcrumbs left behind by the methods attackers use against a target. Defensive rigor

looks at how well available threat intelligence influenced threat hunt efforts. Additionally, the rigor and completeness of the threat hunt methods chosen provided insight into the comprehensiveness of chosen threat hunt analytics. Rigor seeks to both assess the current quality of threat hunts and also to provide opportunities for threat hunts to capitalize on untapped data sources with a high opportunity to detect attack TTP. Organizations that embrace analytic rigor will better be able to analyze strengths and weaknesses of past hunts and improve future hunts with lessons learned and new threat intelligence.

References

- Assante, M., & Lee, R. M. (2015, October). The Industrial Control System Cyber Kill Chain. Retrieved from SANS Reading Room: <https://www.sans.org/reading-room/whitepapers/ICS/industrial-control-system-cyber-kill-chain-36297>
- Bromiley, M. (2016, September). Threat Intelligence: What It Is, and How to Use It Effectively. Retrieved from SANS Reading Room: <https://www.sans.org/readingroom/whitepapers/analyst/threat-intelligence-is-effectively-37282>
- Farrell, A. E., Lave, L. B., & Morgan, G. (2002). Bolstering the Security of the Electric Power System. *Issues In Science & Technology*, 18(3), 49.
- Hosburgh, M. (2017, July 6). Offensive Intrusion Analysis: Uncovering Insiders with Threat Hunting and Active Defense. Retrieved from SANS Reading Room: <https://www.sans.org/readingroom/whitepapers/threathunting/offensive-intrusion-analysis-uncovering-insiders-threathunting-active-defense-37885>
- Kafol, C., & Bregar, A. (2017). Cyber Security – Building a Sustainable Protection. DAAAM International Scientific Book, 81-90. doi:10.2507/daaam.scibook.2017.07
- Lamis, Trevor. (2010). A forensic approach to incident response. In 2010 Information Security Curriculum Development Conference (InfoSecCD '10). ACM, New York, NY, USA, 177-185. DOI=<http://dx.doi.org/10.1145/1940941.1940975>
- Lee, R. M., & Bianco, D. (2016, August). Generating Hypotheses for Successful Threat Hunting. Retrieved from SANS Reading Room: <https://www.sans.org/readingroom/whitepapers/threathunting/generating-hypotheses-successful-threat-hunting-37172>
- Lee, R. M., & Lee, R. (2016, February). The Who, What, Where, When, Why and How of Effective Threat Hunting. Retrieved from SANS Reading Room: <https://www.sans.org/readingroom/whitepapers/analyst/who-what-where-when-effective-threat-hunting-36785>
- Lee, R., & Lee, R. M. (2017, April). The Hunter Strikes Back: The SANS 2017 Threat Hunting Survey. Retrieved from SANS Reading Room:

<https://www.sans.org/readingroom/whitepapers/analyst/hunter-strikes-back-2017-threat-hunting-survey-37760>

Long, M. C. (2016, July 11). Scalable Methods for Conducting Cyber Threat Hunt Operations. Retrieved from SANS Reading Room:

<https://www.sans.org/readingroom/whitepapers/detection/scalable-methods-conducting-cyber-threat-hunt-operations37090>

NIST. (2016, December) Guide for Cybersecurity Event Recovery. doi:

<https://doi.org/10.6028/NIST.SP.800-184>

Walker, C. (2017, July 6). Offensive Intrusion Analysis: Uncovering Insiders with Threat Hunting and Active Defense. Retrieved from SANS Reading Room:

<https://www.sans.org/readingroom/whitepapers/threathunting/offensive-intrusion-analysis-uncovering-insiders-threathunting-active-defense-37885>