Improve your SOC: SOAR or Threat Hunting or both?

Author: Rukhsar Khan
rkhan@rukhsarkhan.de

***Abstract* – Based on the sophistication and constant change of the threat landscape in the cyber space, many mature organizations have identified the necessity to improve the detection, analysis and response capabilities of their Security Operation Center (SOC). Currently, security analysts are often engaged with trivial copy-paste or other annoying low-level tasks rather than gaining a deep understanding of the modus operandi of relevant threat actors and preparing the organization to defend against the risk they pose to it.**

Before we can answer the question, whether a Security Orchestration, Automation and Response (SOAR) solution, the introduction of Threat Hunting or both would be the right course to take in order to improve the SOC, we first need to understand how a SOC is currently operating.

Current SOC working model

In order to **prepare** against cyber threats, a SOC defines use cases for the various threats identified as part of the overall threat landscape. The detection part of the use cases is implemented in a Security Information & Event Management (SIEM) system whereas the Incident Response (IR) part is covered by corresponding runbooks. The SIEM continuously monitors the environment and in case a SIEM rule is triggered, an alarm is generated. This kicks in the IR process in which a runbook gives clear instructions to the security analyst on how to respond to a specific kind of incident. The IR process is well-defined by the SANS institute in its *6-step IR model* [SANS 6-step IR] shown in Figure 1.

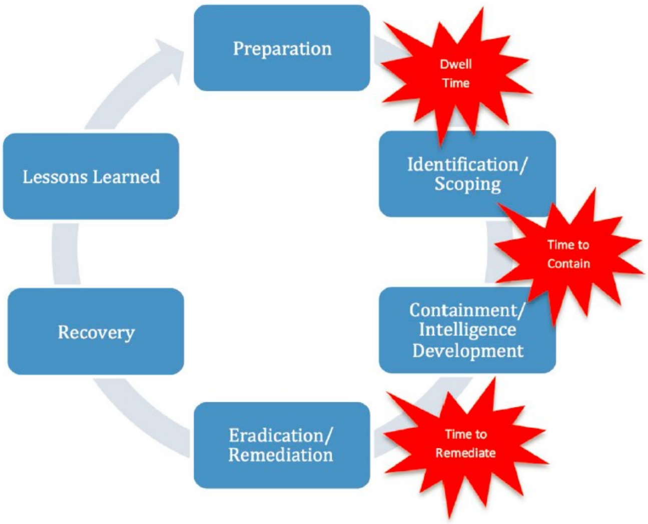

Figure 1 - SANS 6-step IR model

After the initial **preparation** phase, which is described above, the second phase is intended to **identify** and **scope** out an incident.

So, when e.g. a malware incident takes place in a basic SOC organization, the implemented use case could look as follows:

After the Anti-Virus (AV) **identifies** that malware was dropped on an endpoint, the SIEM system triggers an alarm. As part of the **scoping** process, the corresponding runbook instructs the analyst to **verify** the potential malware by

1. checking the malware hash on Virustotal,
2. checking if other systems are affected and
3. checking whether the malware was executed.

In case the malware was "confirmed" by Virustotal and execution evidence was found, next runbook instructions that would be part of the **containment** phase advise the security analyst to

1. isolate the endpoint from the network and
2. block the destination IP address on the proxy which was identified to belong to a C2 server.

Immediately after a successful containment, the next runbook instruction advises the analyst or incident responder to re-image the endpoint. This would reflect the **eradication, remediation and recovery** phases in a single step. Once the re-imaging process is completed and the incident documented as part of the "lessons learned" phase, final instructions are to close the incident.

Another example would be the monitoring of privileged user activity. If an alarm is triggered (**identification**) after a user executed commands with elevated privileges on a server, as part of the **scoping** process, the runbook instructs the SOC analyst to reach out to the server admin and check with him/her whether this was a legitimate action.

It should be noted that a basic SOC organization does merely do a verification as part of the scoping process.

Improve your SOC with SOAR

The main goal of a SOAR tool is to **automate analyst routine tasks** like copy-paste and manual processing (collection, extraction, normalization, formatting) for data enrichment. It also helps to **recover broken media**. E.g. a SOC organization might be using disconnected systems like an IT Service Management ticketing system, a SIEM, email, phone, instant messaging, a wiki and many other tools in order to cover all aspects of the overall IR process.

By leveraging the SOAR as a central hub for IR it allows to **integrate** with all of the SOC systems and tools listed above. This helps in **orchestrating** the complete IR process from within a single platform.

With respect to SOC improvement, the main benefit that comes from the deployment of a SOAR platform is of **quantitative** nature, meaning more incidents can be processed in less time. I.e., through its orchestration and automation features as well as integration capabilities, **a SOAR platform helps an organization to free up valuable time and resources giving the SOC analysts more time for concentrating on their main task**, which is namely

> **"understanding the modus operandi of advanced adversaries and preparing to defend the organization against them"**.

SANS 6-step IR model ≠ SANS 6-step IR model

An advanced SOC organization that is conducting Threat Hunting would also build upon the SANS 6-step IR model, however, the level of implementation would be much higher.

The **preparation** phase begins by consuming Cyber Threat Intelligence from various open source and/or commercial sources. The main goal in this phase is to determine relevant threat actors for the industry and region an organization is located in. Once determined, the organization needs to learn about their capabilities, i.e. their attack ecosystem, their applied attack tools and malware, their Tactics, Techniques and Procedures (TTPs), their current Indicators of Compromise (IoCs), etc.

An excellent model that puts adversaries and their capabilities into a relationship with a victim (organization) and its infrastructure is *The Diamond Model of Intrusion Analysis* [Diamond Model].

These can be documented in an adversary emulation plan [MITRE Adversary emulation plan] which helps the organization to prepare and test the own environment for defending against these threat actors. It's also important to prepare the organization for defending against unknown threat actors.

Improve your SOC with Threat Hunting methodology

In contrast to the quantitative SOC improvement achieved with a SOAR platform, Threat Hunting is of **qualitative** nature. It raises the SOC maturity from a basic to an advanced level by moving from mere **verification** to real **analysis**. Also, a basic SOC is providing **reactive** monitoring which means that the IR process is only kicked in if an alarm is generated whereas an advanced SOC leveraging Threat Hunting methodology has a **proactive** approach.

Various clues can lead to initiate a hunt engagement, however, in this article we will stick to a single procedure, namely *proactively searching for signs of intrusion by the threat actors determined as part of the preparation phase*.

In section *Current SOC working model* we saw how a basic SOC use case for a malware scenario could look like and how it corresponds to the SANS 6-step IR model. Now we can see how an advanced SOC would do the **scoping** based on Threat Hunting methodology.

The first thing to note is that a hunt engagement is always initiated under the assumption that a breach has already taken place.

The very first step in a new hunt engagement is the formulation of a hypothesis [The ThreatHunting Project]. The hypothesis sets the focus of the hunt engagement. An example hypothesis could look as follows:

APT28 has compromised our environment by using an exploit for CVE-YYY-XXX against one of our vulnerable Internet-facing servers and thus gained an initial foothold.

Now, this hypothesis needs to get broken down into smaller, testable elements. Testable elements could be IoCs, TTPs, tools, malware, exploit code for vulnerabilities, assets, etc. Depending on the capabilities of the defense ecosystem and the type of data collected from a defending organization’s infrastructure, a varying list of testable elements could be produced. E.g. our hypothesis could be broken down into the following testable elements:

APT28 has compromised our environment by using an exploit for CVE-YYY-XXX against one of our vulnerable Internet-facing servers and thus gained an initial foothold.

1. APT28 – we might have Threat Intelligence in place that serves with hourly updated IoC lists currently known to be leveraged by APT28 which allows us to test our environment against these.
2. Exploit for CVE-YYY-XXX – our Threat Intelligence might also provide us search patterns or malware hashes for testing on exploit code being leveraged in our environment.
3. Internet-facing servers – these might be grouped together in an asset list.
4. Initial foothold – we might have prepared our environment for being able to test on certain TTPs known to be used by APT28 in order to gain initial foothold.

In order to remain focused at all times of the analysis, accompanying investigative questions do prevail. The above testable elements could answer the following corresponding investigative questions:

1. Which IoCs are currently used by APT28 and thus found in our environment?
2. Which of the malware known to incorporate exploit code for CVE-YYY-XXX has been found in our environment?
3. Which of the vulnerable Internet-facing servers have been identified as being exploited by exploit code for CVE-YYY-XXX?
4. Which of the TTPs known to be leveraged by APT28 have been tested positive?

Every element that is tested positive represents a part of the puzzle and thus becomes part of a so-called **activity thread** [Diamond Model] as depicted in Figure 2.

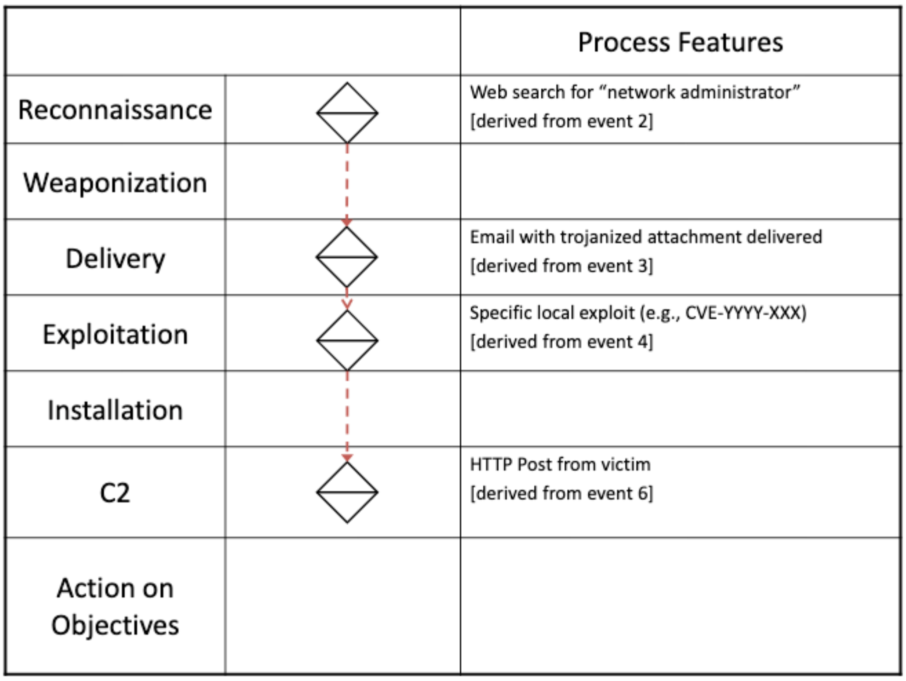

		Process Features
Reconnaissance	◇	Web search for "network administrator" [derived from event 2]
Weaponization		
Delivery	◇	Email with trojanized attachment delivered [derived from event 3]
Exploitation	◇	Specific local exploit (e.g., CVE-YYYY-XXX) [derived from event 4]
Installation		
C2	◇	HTTP Post from victim [derived from event 6]
Action on Objectives		

Figure 2 - Activity thread

This activity thread is extracted from *The Diamond Model of Intrusion Analysis* and thus does not reflect the results of our hypothesis. However, it shows very nicely how such an activity thread can be built based upon positive tested elements of a hypothesis and by linking them together.

Although looking pretty straight-forward theoretically, in practice this can be a real challenge. Every production infrastructure has a lot of noise. Hypothesis testing in a real-world environment would therefore cause many false positives. E.g. many TTP tests would result in being positive because normal administrator activity does also fall into this category. Eliminating false positives and linking together positively tested hypothesis elements which are relevant to a specific attack campaign is the main challenge of a hunt engagement.

Once identified, a defending organization uncovers an ongoing or historic compromise. This methodology also allows us to understand in which stage of the attack lifecycle [MITRE ATT&CK Matrix for Enterprise] an attacker has been at a given point in time and how he/she has moved laterally.

False positive reduction during a hunt engagement is a crucial aspect and can be facilitated by previous verifications done during the scoping phase of a basic SOC use case. This is why the results of a basic SOC verification are highly valuable and should be taken as input for a hunt engagement. In return, the results from the hunt engagement also need to flow back into the original SOC use case as described in the MaGMa use case framework [MaGMa] as follows:

> **“hunting activities may lead to new insights about threats and security monitoring and are therefore input in the lifecycle management for use cases.”**

Conclusion

The level of the SANS 6-step IR model implementation reflects the maturity level of a SOC organization. Basic SOC organizations without Threat Hunting do mere **verifications** as part of the scoping phase. Advanced SOC organizations that are performing Threat Hunting are able to do **extensive analysis** as part of this phase.

While the results from a basic SOC verification might just reflect the tip of the iceberg, they are still very valuable to a hunt engagement as they help enormously in false positive reduction.

With respect to the SOC maturity level, the outcome from deploying a SOAR platform is of quantitative nature whereas the introduction of Threat Hunting methodology helps an organization to improve its quality, thus raising the SOC maturity from a basic to an advanced level.

Therefore, the original question can be answered as follows:

In order to improve your SOC, start with deploying a SOAR platform if you haven’t already done so. This helps you in freeing up valuable time and resources so that your analysts can concentrate on the main task, which is gaining a deep understanding of the modus operandi of relevant threat actors and preparing the organization to defend against the risk they pose to it.

Bibliography

[SANS 6-step IR] It's Awfully Noisy Out There: Results of the 2018 SANS Incident Response Survey

[Diamond Model] http://www.activeresponse.org/the-diamond-model/

[MITRE Adversary emulation plan] https://attack.mitre.org/resources/adversary-emulation-plans/

[The ThreatHunting Project] http://www.threathunting.net/sqrll-archive

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

[MITRE ATT&CK Matrix for Enterprise] https://attack.mitre.org/matrices/enterprise/

[MaGMa] https://www.linkedin.com/pulse/magma-use-case-framework-released-today-rob-van-os