Fingerprinting False Positives: How to Better Integrate Continuous Improvement into Security Monitoring

DESIREE SACHER, Finanz Informatik GmbH and Co KG

Security incidents created during the security continuous monitoring process (ISCM, NIST 800-137 [1]) can be a valuable source for improving company processes and an efficient reference for budget planning if used correctly. The idea presented in this article was created from years of experience in monitoring NIST 800-53 [2] controls, among others, where often the processes in place did not address the root cause of the problem but rather added suppressions to the detection rule. This article discusses how security incidents can be tracked to focus on company improvements and strategically create statistical values. It is a suggestion for a structured resolution approach, inspired by the Post-Incident Activity—Lessons Learned phase by NIST 800-61 [3].

CCS Concepts: • **Security and privacy** → **Formal security models**; • **Software and its engineering** → *Formal methods*;

Additional Key Words and Phrases: Security monitoring, continuous improvement, SOC, CDC, security operation center, incident, quality, rule

ACM Reference format:

Desiree Sacher. 2020. Fingerprinting False Positives: How to Better Integrate Continuous Improvement into Security Monitoring. *Digit. Threat.: Res. Pract.* 1, 1, Article 7 (March 2020), 7 pages.
<https://doi.org/10.1145/3370084>

1 INTRODUCTION

For a long time, security-related incidents have been categorized into true positives, when an alert actually detected a malicious event, and false positives, when no malicious activity could be identified. *False positive* is the term most often used to express that all of the safety measures worked and no harm was done to the entrusted infrastructure. When reminded of what our current security measures have grown out of—a company lifecycle of policies, specifications, and configuration guidelines—this view makes sense. In today’s world, where security processes need constant improvement as attack scenarios change faster than ever, this view is obsolete.

When looking at security measures through the SOC perspective, it becomes clear that false positives become an antiquated measure, not supporting technical changes often needed in companies. At worst, the statistics created from security incident tracking even reassure management that everything is well and no improvements need to be done. Usually at this stage, the security professionals turned to penetration tests to illustrate technical gaps in the security of products and services. Penetration tests are still a helpful instrument, but they do not always show cultural barriers or give a well-investigated basis for discussion. In addition, they can be very expensive, and thus not every company can afford them.

What this article suggests is to use security incidents as a means to track and illustrate improvements and adjustments of the security infrastructure. SOC’s are most efficient when every incident or event can be turned

Author’s addresses: D. Sacher, Finanz Informatik GmbH & Co KG, Nevinghoff 25, 48147 Münster; email: d.sacher@gmail.com.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2020 Copyright held by the owner/author(s).

2576-5337/2020/03-ART7

<https://doi.org/10.1145/3370084>

Table 1. Overview Resolution Categories

	English Version	Deutsche Version
(A)	Announced administrative/user action	Kommunizierte administrative/Benutzer Tätigkeit
(B)	Unannounced administrative/user action	Unangekündigte administrative/Benutzer Tätigkeit
(C)	Log management rule configuration error	Log-Management-Regel-Konfigurationsfehler
(D)	Detection device/rule configuration error	Sensor-/Endpunkt-Regel-Konfigurationsfehler
(E)	Bad IOC/rule pattern value	Schlechter IOC/Regel-Vergleichswert
(F)	Test alert	Test-Alarm
(G)	Confirmed attack with IR actions	Bestätigter Angriff mit IR-Massnahme
(H)	Confirmed attack attempt without IR actions	Bestätigter Angriffsversuch ohne IR-Massnahme

into a specific action to improve overall security of a company; therefore, all events should be categorized as such.

2 MATERIALS AND METHODS

Throughout my work over the past 3 years, where I was part of new SOC's that were built up, I have promoted tracking of security incidents and implemented the resolution categories displayed in Table 1. Every category is illustrated with an example. Section 2.1.2 explains what this category in numbers can tell the reader and why it is of importance. Section 2.1.3 provides a quick outlook on what positive changes can be introduced when interpreting the numbers of such a statistic.

The resolution categories were defined for SOC monitoring rulesets, also called *use cases*, usually configured in a SIEM. Security incidents created by use cases defined as part of the information security continuous monitoring process, often are the primary source of information for SOC's, as generally this process is best documented and visible to the company. This value is not intended to replace other categorization but should be instead be used as an additional field to improve SOC quality by focusing on what caused the wrong alert instead of just suppressing part of the detection rule. Figure 1 can assist in finding the right categorization.

The categorisation should be applied after successful analysis of the incident where identification of root cause for the event has been finished, along with the diagnosed actual impact of the event.

2.1 Announced Administrative/User Actions

2.1.1 Examples.

- Detected port scan can be correlated to a previously communicated penetration test.
- Support connection with administrative privileges was detected on a user device with default privilege.

2.1.2 Description/Significance. The process to communicate administrative activities or special user actions was in place and working correctly. Internal sensors are working and detecting privileged or irregular behavior.

2.1.3 Benefits. Regular statistics can be created about tested security measures and processes. Log settings or suppressions for known activities can be optimized.

2.2 Unannounced Administrative/User Actions

2.2.1 Example.

- A network scan was performed by a network engineer while troubleshooting a problem.

2.2.2 Description/Significance. Internal sensors have detected privileged or user activity, which was not previously communicated. It can also reflect improper usage behavior. This illustrates a problem with internal

Fig. 1. Decision matrix for applicability categorization.

communication channels or processes. This category, when evaluated, illustrates time and effort by the SOC that was spent following up on issues that actually were not a problem. With growing maturity of an SOC, as much time as possible should be spent on actual technical analysis or the optimization of the detection capabilities.

2.2.3 Benefits. This category reflects how well the SOC is included in surrounding IT processes. It can also be used as an indicator for awareness throughout the company, as reaction from the SOC emphasizes that such behavior is monitored, which might not have been well known yet.

2.3 Log Management Rule Configuration Error

2.3.1 Example.

- Analysis of alerts for failed logins shows a misconfiguration in the central log management system rule, where the algorithm for counting the failed attempts is wrong.
- Analysis of alerts for command and control traffic IPs shows connection to a multihoster system, where the actual URL accessed was not compromised.

2.3.2 Description/Significance. This category reflects false alerts that were raised due to configuration errors in the central log management system, often a SIEM, rule. These errors in the detection algorithm should be corrected as fast as possible to prevent further false alarms. Following false alarms by this rule can lead to threat alert fatigue, which can occur if alarms are not taken seriously anymore by analysts because they already expect it to be a false alarm. Threat alert fatigue should be avoided at all cost to keep up analyst motivation and considering that appropriate use of an analyst's resources is the simplest form of respect that can also prevent burnout.

2.3.3 Benefits. An analyst's motivation can be kept on a higher level as the trust in the ruleset can be sustained. This category can reflect the quality of the rules, as especially rules that were properly vetted in the baselining phase and moved to a productive state should not raise false positives by the algorithm anymore. If false positives by SIEM rule configuration error persist, skills of the SIEM engineers might need to be improved or the rule logic might need to be adjusted. As SIEM rules are usually maintained by the SOC, SOC analysts can directly influence this category. This category can be a valuable key performance indicator (KPI), as too many SIEM rule errors reflect basic SOC skill or understanding of infrastructure problems.

2.4 Detection Device/Rule Configuration Error

2.4.1 Example.

- The IDS sends an alert to the SIEM for a suspicious pattern detected within application traffic in a subnet, where this application is actually not located.
- The IDS rule pattern has created false positives because the rule is not specific enough.

2.4.2 Description/Significance. This category reflects rules on detection devices, which are usually passive or active components of network security. In bigger organizations, these tools are often maintained by, for example, the network team. By the defined nature and scope of the teams, needed adjustments of these rulesets can often not be visible to the responsible maintaining teams, as they rarely are focused on processing security alerts. Therefore, a close and trusted interaction with the SOC is essential. As both teams are met by different types of KPIs (usually network and firewall teams are measured by operation-related KPIs like uptime, and SOC's are measured by the response time and completion/quality of analysis reports), adjustments to the rulesets and configuration can conflict with SLAs or might only be performed with a lower priority. To create KPIs for the SOC on this category would therefore not be fair.

This value can also be important in an SOC as a Service/MSSP relationship, as these devices are often not controlled by the SOC company, and therefore bringing visibility into customer configuration ruleset quality might be needed to support changes.

2.4.3 Benefits. Direct visibility into device configurations and rulesets to support the need for configuration adjustments. If no improvements to the ruleset can be performed by the responsible team, this can indicate the need for a different product or structural adjustments to the responsibilities of the involved teams or the company processes.

2.5 Bad IOC/Rule Pattern Value

2.5.1 Example.

- An alert for an IP address categorized as Command and Control connection can upon analysis be classified as an obsolete indicator, which no longer hosts malicious services.

2.5.2 Description/Significance. Products often require external indicator information or security feeds to be applied on active or passive infrastructure components to create alerts. The quality of these indicators should be measured in a separate category as the used tool might actually fulfill SOC requirements, but the additional security feed might not uphold to the requested standards. This category can directly support needed changes in the SOC supply chain.

2.5.3 Benefits. The source of the false positive can be clearly identified and proves that the SOC performance is better than superficial results might suggest. With this information, adjustments to the SOC supply chain can be initiated.

2.6 Test Alert

2.6.1 Example.

- The alert was created for testing purposes by the SOC team.

2.6.2 Description/Significance. This alert reflects alerts created for testing purposes, which can be caused by regular unit tests, if such processes are in place, or single tests performed when baselining or fine tuning a rule. It should be counted in a separate category to not confuse it with other measures. No direct improvement suggestions can be derived from this category, but they might be a seal of quality because it supports a working SIEM rule configuration process. These incidents should also always be excluded when the number of attacks is calculated for managers or customers.

2.6.3 Benefits. Direct visibility into business relevant amount of security incidents.

2.7 Confirmed Attack with IR Actions

2.7.1 Example.

- An alert for an outgoing connection to a URL provided by an IOC reveals an infection. Further host analysis is performed.

2.7.2 Description/Significance. This alert represents the classic true positives, where all security controls in place were circumvented, a security control was lacking, or a misconfiguration of a security element occurred. In this case, a full Incident Response has to be started (according to NIST 800-61) and an in-depth analysis needs to be performed to identify prospective needed prevention measures during the lesson learned phase.

2.7.3 Benefits. This value represents the classic true positives, which is often the number that managers or MSSP customers are most interested in. They represent a direct risk for the confidentiality and integrity of company data and should therefore be processed with cautious steps, possibly involving external specialists where needed.

2.8 Confirmed Attack Attempt without IR Actions

2.8.1 Example.

- An antivirus alert is raised on a client device for detection of a malicious software. Infection was prevented.

Table 2. False Positive/True Positive Comparison by Perspective

Case	C-Level Perspective	SOC Perspective	MSSP Account Manager Perspective	Follow-Up Action
<i>Key Driver</i>	<i>Does this alert inform me about an actual threat to the company?</i>	<i>Are our SIEM rules/detection capabilities working correctly?</i>	<i>Were the MSSP service systems configured correctly to detect a threat to my company?</i>	<i>What lesson can be learned from this event?</i>
<i>Announced administrative/user action</i>	No—false positive	Yes—true positive	No—false positive	Update suppressions for announced actions
<i>Unannounced administrative/user action</i>	No—false positive	Yes—true positive	Yes—true positive	Update information process
<i>Log management rule configuration error</i>	No—false positive	No—false positive	No—false positive	SIEM rule correction needed
<i>Detection device/rule configuration error</i>	No—false positive	No—false positive	No—false positive	Detection device/rule configuration correction needed
<i>Bad IOC/rule pattern value</i>	No—false positive	No—false positive	No—false positive	IOC provider should be accredited
<i>Test alert</i>	No—false positive	Yes—true positive	Yes—true positive	Should be excluded from reporting
<i>Confirmed attack with IR actions</i>	Yes—true positive	Yes—true positive	Yes—true positive	Lesson learned should point out needed infrastructure improvement
<i>Confirmed attack attempt without IR actions</i>	No—false positive	Yes—true positive	No—false positive	To be included in SOC report to reflect well spent budget

2.8.2 Description/Significance. This category reflects an attempt by a threat actor, which in the end could have been prevented by in-place security measures but passed security controls associated with the delivery phase of the Cyber Kill Chain [4]. It is the desired state of a security incident, as it proves current measures to be sufficient. This category can also reflect alerts, where the risk was accepted by company security policy and no further investigation officially is intended. Confirmation for prevented infection should always come from active infrastructure components. Mapping this resolution category to passive infrastructure components only detecting and not preventing threats can indicate handling errors as the confirmed deny action was not recorded and the analyst might have come to wrong assumptions. If the number of these category incidents is low, infrastructure improvements earlier in the Cyber Kill Chain (Phase: Delivery, Installation, and Exploitation) should be prioritized because it may indicate lack of successful detection capabilities [5].

2.8.3 Benefits. This value management and MSSP customers are most interested in the general security monitoring process, as it represents actual attack attempts and can be used to illustrate and support security spending.

3 RESULTS

These categories allow a more granular definition of false positives, as a false positive is actually subjective to the decision maker. For example, from the perspective of a CSO, all alerts except those belonging to the category

“Confirmed attack with IR actions” should be seen as false positives, as they represent actions performed by the SOC, where no harm to the infrastructure was done and thus all current security measures were sufficient. If these numbers are evaluated well, they can already today present a great image of where the budget was spent well. From an SOC perspective, although an alert categorized as “Unannounced administrative actions” and “Confirmed attack attempt without IR actions” represents a well-working ruleset, it actually detected suspicious behavior and is therefore rather a true positive. Even “Announced administrative actions” can represent a well-working ruleset, although the process for suppressing such alerts might want to be improved in case the rules have been extensively proofed to be working well in the past. Of course, the CSO/CRO would not want this value to be counted in the number of attacks represented to the board of managers to indicate additional budget needs.

An SOC manager or architect, however, can use these values to see on what the team is most spending its time on. The values become indispensable for basing strategic decisions for further needed infrastructure components, partners, or resources. Table 2 highlights the main differences in perspective and what follow-up action should be taken after classification of the event.

4 DISCUSSION

Other methods to categorize incidents usually relate to the type of event that was analyzed, such as Denial of Service, Forensics, and Compromised Asset [6], or related to other activity or criticality. The MISP project [7] created a great overview of known taxonomies for such cases. This categorization method, however, rather focuses on what lesson can be learned from each monitoring event by identifying the real root cause for the rule to trigger, so it actually reflects considerations usually done at the post-incident activity phase by NIST 800-61, which might not be already applied by the standard continuous monitoring process. It should therefore also be done at the end of a security incident.

This method of categorizing security incidents created by continuous security monitoring rulesets has been proven valuable in the business context of an MSSP, but also within bigger companies succumbing to standards by NIST, PCI DSS, and Swift, among others.

ACKNOWLEDGMENTS

I originally came up with the dislike for false positives when maintaining security device rulesets of a former employer. Thank you to Christoph Weber and Michael Kurth for supporting and vetting this idea and successfully implementing it at our shared former employer. Thank you to Raphaël Vinot, Corsin Camichel, Eireann Leverett, Florian Roth, Ian Amit, Meline Sieber, Frank Boldewin, Jochen Raymaekers, Francesco Picasso, and Amanda Berlin for further vetting and correcting this article. I additionally want to thank FIRST for sponsoring this publication.

REFERENCES

- [1] NIST. 2011. Information Security Continuous Monitoring (ISCM) for Federal Information Systems and Organizations. Retrieved January 31, 2020 from <https://csrc.nist.gov/publications/detail/sp/800-137/final>.
- [2] NIST. (n.d.). Security and Privacy Controls. Retrieved January 31, 2020 from <https://nvd.nist.gov/800-53>.
- [3] NIST. 2012. Computer Security Incident Handling Guide (800-61). Retrieved January 31, 2020 from <https://csrc.nist.gov/publications/detail/sp/800-61/rev-2/final>.
- [4] Lockheed Martin. 2020. The Cyber Kill Chain®. Retrieved January 31, 2020 from <https://www.lockheedmartin.com/en-us/capabilities/cyber/cyber-kill-chain.html>.
- [5] Eric M. Hutchins, Michael J. Cloppert, and Rohan M. Amin. 2010. Intelligence Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains. Retrieved January 31, 2020 from <https://www.lockheedmartin.com/content/dam/lockheed-martin/rms/documents/cyber/LM-White-Paper-Intel-Driven-Defense.pdf>.
- [6] GitHub. (n.d.). Incident Categories by FIRST CSIRT. Retrieved January 31, 2020 from https://github.com/MISP/misp-taxonomies/tree/master/csirt_case_classification.
- [7] GitHub. (n.d.). MISP Taxonomy Overview. Retrieved January 31, 2020 from <https://github.com/MISP/misp-taxonomies>.

Received July 2019; revised October 2019; accepted October 2019