FI-ISAC NL publication

MaGMa

Use case framework

MaGMa

•

A joint use case framework from the Dutch financial sector

•

Version 1.0
15.11.2017

Foreword

This use case framework (UCF) was created as a joint effort between several Dutch financial institutions. The focus group operated as part of the Dutch financial institutions information sharing community (FI-ISAC), in particular the SOC/CSIRT working group. The goal of this framework is to stimulate practical sharing of knowledge regarding security monitoring use cases between the Dutch financial institutions. Due to the fact that the framework itself is not specific for the financial sector, but can be applied to any organization, it was decided to release the framework and the accompanying excel tool for framework management into the public domain. The framework is called 'MaGMa UCF' after the most important elements of the framework: Management, Growth and Metrics & assessment. Together, these elements are vital for the creation, maintenance and ongoing added value of the use case framework to the organization.

The MaGMa UCF framework is based on the existing Use Case Framework in use by ABN AMRO and developed by Floris Ladan and Tony Tromp, complemented and modified using insights and experiences from other participating organizations in the focus group. The focus group is thankful for being able to use this excellent framework as a basis for the joint framework.

The MaGMa UCF tool that allows for a practical implementation of the framework described in this document can be downloaded from the site:
https://www.betaalvereniging.nl/wp-content/uploads/Magma-UCF-Tool.xlsx

On behalf of the FI-ISAC focus group on use case management,
Rob van Os

Authors
Rob van Os, de Volksbank, *lead author and UCF developer*
Floris Ladan, ABN AMRO Bank, *UCF lead developer*
Thomas van Casteren, Rabobank, *UCF developer*
Robin Toornstra, Euroclear
Robert Metsemakers, Achmea
Lambrecht Nieuwenhuize, BNG Bank
Holger Grotenhuis, Triodos Bank

Additional contributions by
Kelvin Rorive, Rabobank
Marcus Bakker, ING Bank
Tony Tromp, ABN AMRO Bank, *UCF developer*

Contents

FOREWORD 2

1 INTRODUCTION 4

1.1 SECURITY MONITORING 4
1.2 SECURITY MONITORING USE CASES 5
1.3 INTENDED AUDIENCE 5
1.4 DOCUMENT OUTLINE 5

2 USE CASES 7

2.1 USE CASE ELEMENTS 7
2.2 OVERVIEW 9
2.3 MAGMA UCF TOOL USAGE 10

3 MANAGEMENT OF USE CASES 12

3.1 ONBOARDING (PLAN AND BUILD) 12
3.2 OPERATIONAL PHASE (RUN) 12
3.3 MAINTENANCE (CHANGE) 12
3.4 OFFLOADING (DECOMMISSION) 15
3.5 OVERVIEW 15

4 GROWTH (CAPABILITY AND MATURITY) 17

4.1 CAPABILITY 17
4.2 MATURITY 18
4.3 FINAL REMARKS ON CAPABILITY AND MATURITY 20

5 METRICS & ASSESSMENT 21

5.1 MAGMA UCF EMBEDDED METRICS 21
5.2 MAGMA CONTROL METRICS 22
5.3 MAGMA OUTPUT METRICS 23

6 BEST PRACTICES FOR USE CASE MANAGEMENT 25

7 CONCLUSION 27

1 Introduction

A Security Operations Center (SOC) has a central role in protecting against, and dealing with cyberattacks. In the ever changing landscape of cyber security, there are many threats to protect against. Each of these threats can have unique indicators in different parts of the cyber killchain [[1]]. It is the job of the SOC to recognize the cyber threats facing the organization in an early stage, mount the appropriate response and help to adjust security parameters to avoid breaches. This process of monitoring for manifestations of cyber threats is called security monitoring.

1.1 Security Monitoring

While threats are a core driver for security monitoring, they are not the only driver. Other aspects of the organizational environment must be considered as well. Figure 1 shows the environment for security monitoring.

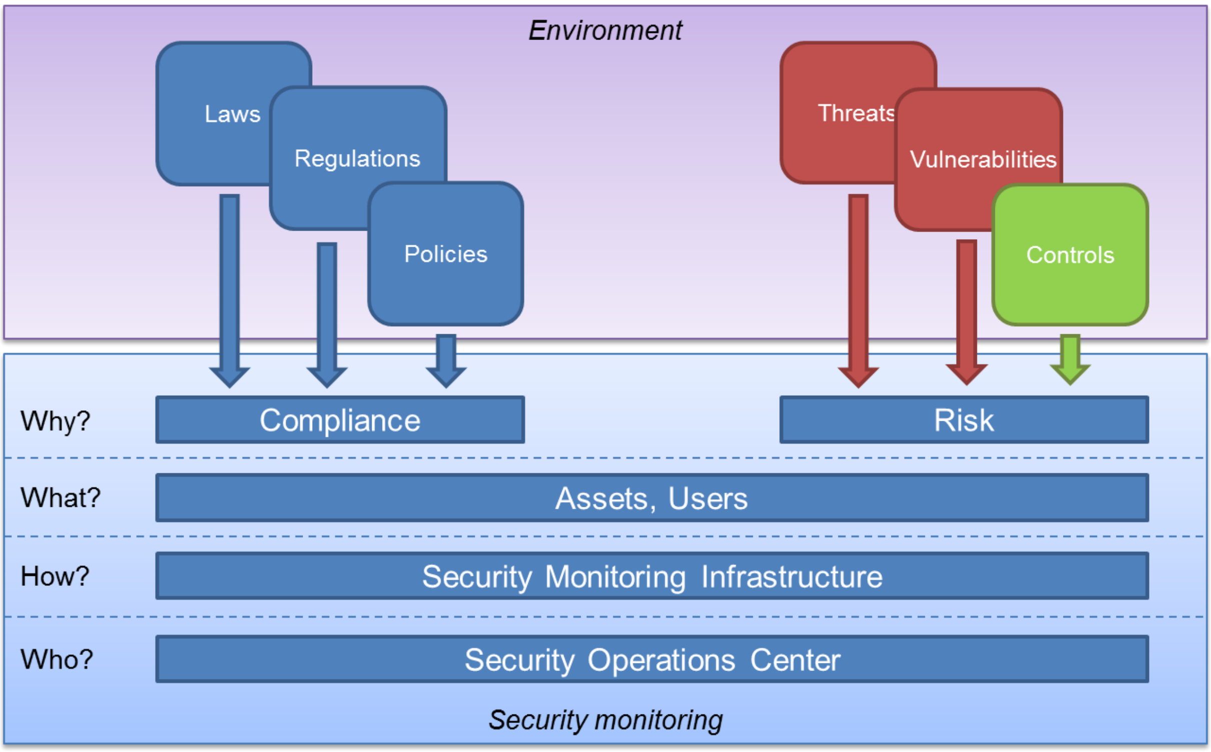

Figure 1: security monitoring environment

To support the security monitoring process, a security monitoring system is used. In particular, a SIEM (Security Information and Event Management) system is an essential part of the security monitoring processes. Additional monitoring systems, such as anomaly detection, intrusion detection (both network based and host-based) and big data platforms for analytics can also be found in more elaborate security monitoring infrastructures. Figure 2 shows an example security monitoring infrastructure.

[1] http://www.lockheedmartin.com/us/what-we-do/aerospace-defense/cyber/cyber-kill-chain.html

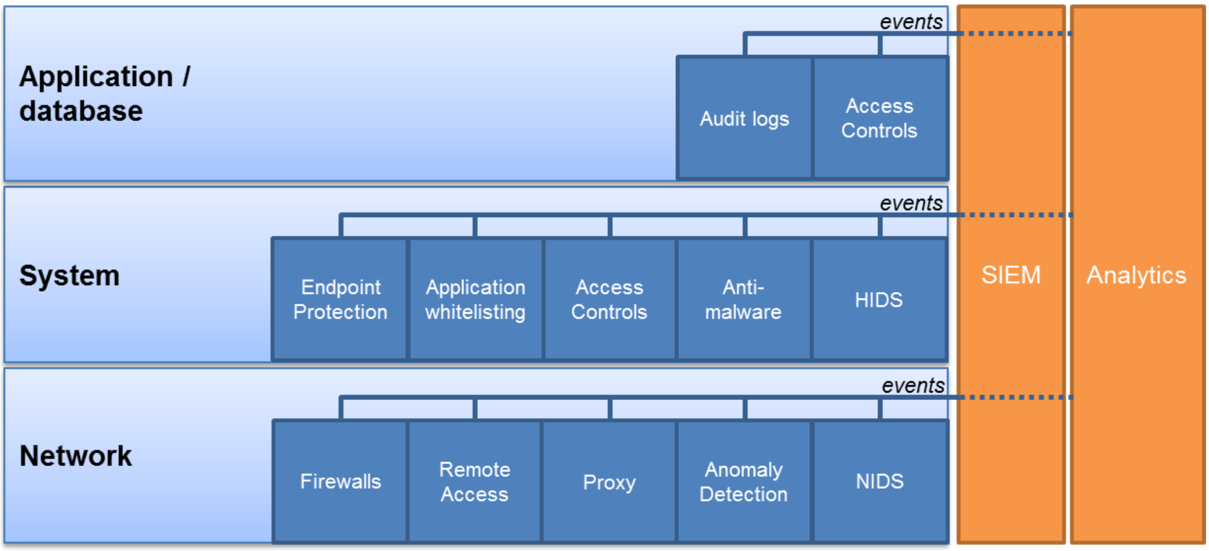

Figure 2: example security monitoring infrastructure

1.2 Security monitoring use cases

To provide a structured approach to security monitoring, use cases are used. Essentially, use cases describe manifestations of threats from a high level (the modus operandi of the cyber criminals) to the lowest level (concrete security events in the infrastructure such as exploits, failed logins, etc.). Use cases also describe follow-up actions (incident response) and are tied with business drivers to show how security monitoring reduces risk in the organization. Within the complexity of the security architecture, use cases can provide structure and overview.

To organize use cases, a use case framework should be used. Such frameworks enable control over use cases and provide insight into identify how well an organization is capable of defending against cyber threats. In this document, The MaGMa Use Case Framework (hereafter called: MaGMa) is presented. This use case framework was created in a collaborative effort of several financial institutions associated with the Dutch Financial Information Sharing and Analysis Community (FI-ISAC) and is based on the existing framework and tool developed and used by ABN AMRO Bank, complemented with views, experiences and best practices from other financial institutions.

To support this document, a UCF tool is provided that can be used for practical implementation of the framework in any organization. For convenience, several high-level use cases have been added to the tool already.

1.3 Intended Audience

This document is intended for security managers, SOC managers, SOC analysts, and professionals within the SOC focusing on design and maturity and the security monitoring process.

1.4 Document outline

The MaGMa use case framework is explained in this document by first going into the elements that compromise a use case. Each of these elements is shortly addressed and the relevance to the framework is explained. Since these elements are also part of the tool, some guidance is provided on practical implementation. Hereafter, the pillars of the framework will be addressed: use case management, growth in terms of capability and maturity and lastly metrics and assessment. Finally, some few best practices will be shared as well.

A note on scope: while use cases may also cover other NIST cybersecurity framework stages such as identify and protect, the focus of this use case framework will be on the 'detect' and 'response' phases, with most emphasis on the 'detect' phase.

2 Use Cases

Before going any further, it is important to provide a definition of use cases. This is necessary, because of the fact that the term is used to describe a variety of elements in different publications. This complicates comparison of publications and frameworks.

In this document, the following definition for uses cases is used: “*a use case is a security monitoring scenario that is aimed at the detection of manifestations of a cyber threat*”. A use case has a strategical, tactical and operational component.

As indicated in the introduction, this definition again emphasizes the fact that the focus of the framework is on the ‘detect’ phase of the NIST cybersecurity framework. Naturally, the events and incidents that flow from the security monitoring architecture can be used to improve protection mechanisms or further refine the threat identification phase.

In discussing use cases, several elements of use cases were identified. These elements will be outlined and discussed. Additionally, the practical application of these elements in the supporting tool is also explained.

2.1 Use Case Elements

The elements that comprise the use case be divided into three layers:

- Business layer. The business layer of the use case describes how the use case is connected to the organization’s business needs.
- Threat layer. The threat layer of the use case describes the threat that the use case is intended for. Several aspects of the threat are important.
- Implementation layer. This is the operational layer, where aspects that are relevant for implementation of the use case in the operational security monitoring architecture are described.

Each of these layers is separately addressed hereafter. Note that in the use case framework itself, two-way traceability is important. Thus, it must be possible to connect elements at the operational layer to elements at the tactical and ultimately strategical layers and vice versa. This allows the SOC to show how business drivers are implemented in operational monitoring (top-down) and which monitoring rules relate to which specific threats and business drivers (bottom-up).

In addition to the previous three layers, MaGMa also differentiates in 3 levels of use case detail. These 3 levels (L1, L2 and L3) describe the threat from high level to low level (actual monitoring rules). These levels are essential to the framework and are prominent in the tool supporting MagMa. The L1 and L2 use cases are part of the tactical layer (threats), L3 is part of the operational layer.

2.1.1 Business layer (strategical level)

The business layer of the use case addresses the elements that are relevant to ensure that the use case supports the business and vice versa. This is the strategic component for the use case. For many use cases, the business layer should be similar if not identical. The following items should be part of the business layer of the use case:

- Purpose. First and foremost, the purpose of the use case should be made clear. Why is this use case relevant for the business? In other words: what business value does an effective implementation of this use case have?
- Drivers: what are the main drivers for this use case? These drivers are usually risk-reduction, avoiding financial loss, avoiding reputational damage or similar. Besides business drivers, compliance drivers (internal policies and policies set by external regulators) may apply as well. Mapping compliance drivers to business drivers and use cases can be an elaborate task, but can assist greatly in the auditability of the framework and the security monitoring process in general.
- Stakeholders: who are the main stakeholders for this use case? The stakeholders are usually tightly coupled with the drivers.
- Output: what output will this use case provide? Output can be generic, such as events and incident response, but can also be more specific. For example, specific reports may be created for this use case. Such reports may even be differentiated for each stakeholder.

2.1.2 Threat layer (tactical level)

The tactical layer of the use case is used to align the use case with the threat management processes. For this layer, several elements are particularly important:

- Threats. This is the single most important element of the use case. What threat is addressed by this use case? As many organizations create a threat landscape as part of their threat management process, it makes a lot of sense to use the threat landscape as the main input for your use cases. In the MaGMa tool, these threats are documented as level 1 (L1) and level 2 (L2) use cases. The L1 use case describes the threat at a higher level (e.g. Exfiltration), while the L2 use case describes the threat more concretely (e.g. exfiltration by malware).

  Note: L1 use cases are comprised of the 7 steps of the cyber kill chain, complemented with additional threats outside the kill chain. Examples of such additional threats are DDoS and policy non-compliance.
- Actors. Threat actors are also relevant for the use cases as different threat actors pose a different threat. This is due to the fact that all actors use a specific set of capabilities, expressed as tactics, tools and procedures (TTPs) and thus have a different threat profile. Potential motivations for actors should also be included. Lists of actors are readily available, for example from the Cyber Security Assessment Netherlands [[2]]. Similarly, lists of motivations are available as well. The list of standard actors and motivations can be extended as needed.
- Security incident response. What actions need to be taken when security monitoring alerts are fired relating to this use case? It is important to determine the appropriate response *before* implementing the use case, because significant added value from the SOC comes from incident response. Within the incident response process, roles and responsibilities for the use case will also need to be described. By outlining the roles and responsibilities beforehand, security incident response can be carried out

[2] https://www.ncsc.nl/english/current-topics/Cyber+Security+Assessment+Netherlands/cyber-security-assessment-netherlands-2017.html

smoothly when incidents do occur. While roles and responsibilities are not part of the use case framework, they are essential for effective security incident response.

- Security analysis. For analysts, it is very helpful to have some guidance on analysis alerts generated by security monitoring rules. Such guidance will aid analysts in correct interpretation of the security monitoring alert and to ultimately decide whether the alert is a genuine threat or a negligible event. Since analysis for many use cases will follow a similar approach, most likely through a Standard Operating Procedure (SOP), this activity belongs at the L2 level.

2.1.3 Implementation layer (operational layer)

The implementation layer of the use cases addresses the operational aspects of the use case in the security monitoring architecture. The following items are part of the implementation layer of the use case:

- Log sources. Which log sources can provide input into this use case? It is important to take the whole architecture into account. For example, threats that act on the application layer have an application component, but likely also a system and network component. Log sources from each of these layers can be included into the use case.
- Scope. What is the scope for this use case in terms of users (accounts), assets, applications, etc.? Consider the criticality of systems as context into the use case. The use case should at least cover the most critical systems first (i.e. the crown jewels).
- Detection mechanisms. Besides scope and log sources, it is also important to indicate which detection technology is used for detection of the L3 use case. This allows for aggregated views in which the added value of different techniques to the security monitoring architecture can be shown. Such views may be used to identify strong and weak areas as well as growth potential.
- Monitoring rules. A vital part of the operational layer is the actual implementation in monitoring rules, as these feed the incident response process. Monitoring rules should relate directly to ‘incident response’ in the operational layer, ‘threats’ from the threat layer and, if required, ‘output’ from the business layer. In the MaGMa UCF tool, these monitoring rules are documented as L3 use cases. Embedded metrics apply at this level. These metrics are described in chapter 5.

2.2 Overview

Figure 3 provides an overview of the use case elements described in this section and also lists the most important input for each of the layers.

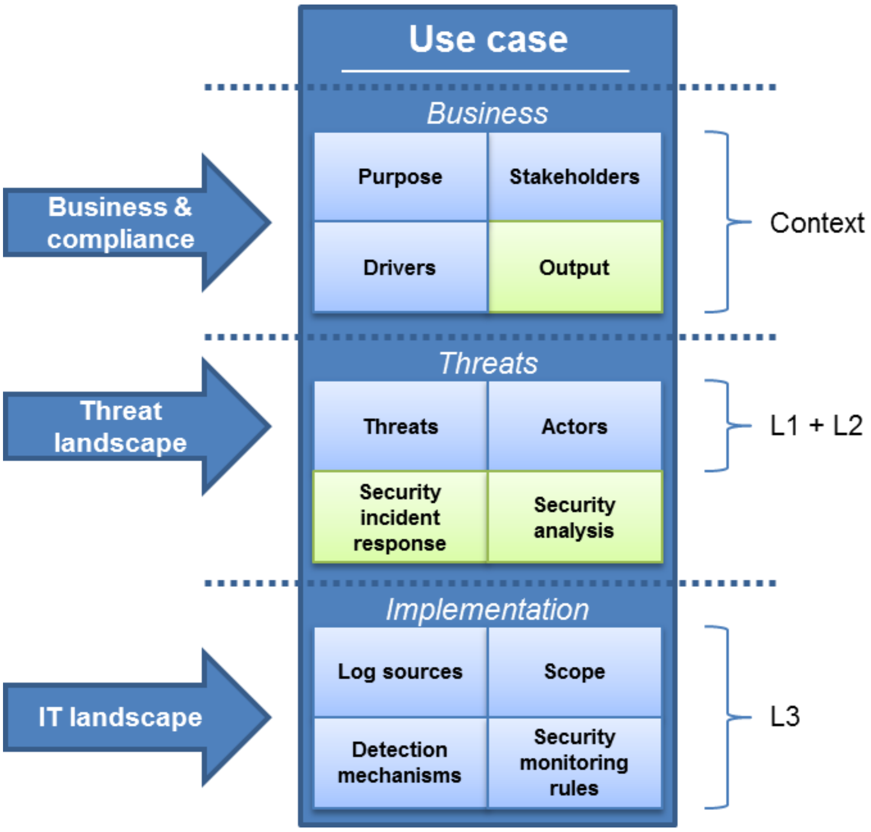

Figure 3: MaGMa UCF use case model

2.3 MaGMa UCF tool usage

The MaGMa UCF tool supports the documentation and management of most of the elements mentioned in this chapter, with the exception of 'output', 'security incident response' and 'security analysis'. The reason for this is that these elements are usually described in much more detail elsewhere. They have been excluded from the tool to avoid overlapping documentation and to limit the size and complexity of the tool.

For 'output', service level agreements are usually the place in which formal agreements on reporting, such as frequencies and content, are documented. Most use cases will not have specific output, but instead are simply connected to the security incident response process. For 'security incident response', documentation regarding incidents and follow up should be part of the security incident response documentation. Lastly, for 'security analysis', guidance on the interpretation of security monitoring alerts should be provided in a security analysis handbook or a similar document, for example an SOP While these 3 elements are not described in detail in the tool itself, references to existing procedures or document can be included if so desired.

All remaining elements can be documented in the MaGMa tool. Usage of the tool is explained in detail in the tool itself, but the basic steps for implementation are the following:

1. First, outline the strategical business layer. This will provide the necessary context for the concrete use cases. Outlining the business layer is done by speaking with

business stakeholders and determining purpose, business drivers and optionally compliance drivers for the framework.

2. Then, create the L1 use cases using the threat landscape. The 7 basic L1 use cases representing the steps in the cyber kill, complemented with additional L1 use cases are already present in the tool. Additional L1 use cases that are specific for your industry should be added. For example: financial fraud (for the financial sector), Industrial control system sabotage (for energy and utilities sector).
3. Once the L1 use cases have been finalized, L2 use cases can be put into place. An extensive list of 62 L2 use cases is already present in the tool. This list was created using input from all participating organizations. The list may contain use cases that are not applicable to your organization, or may not fully cover your security monitoring requirements. Therefore, it is important to carefully select the proper use cases and extend upon this list where required and useful.

   Note that for steps 2 and 3, it is of vital importance to include security management stakeholders to ensure proper alignment with organizational risk & security management processes.
4. Lastly, the L3 use cases should be outlined and operationalized. First, outline all of the actual L3 monitoring rules. Then fill in all other applicable operational elements (log sources, scope and detection mechanisms) for each of the L3 rules as desired. Furthermore, the embedded metrics as described in chapter 5 should be set for each of these rules when they are operationalized.

   A number of L3 rules is already present in the tool. These L3 rules are based on the MITRE ATT&CK Matrix for Enterprise [[3]]. For this purpose, the matrix was first de-duplicated mapped to MaGMa use cases at the L2 level. Not all L1 use cases, and only a fraction of the L2 are covered by the MITRE matrix due to its focus on actions after initial compromise. De-duplication was performed because the MaGMa tool was not designed to map a single L3 use case to multiple L2 use cases. Each L3 use case has a single L2 reference and each L2 use case has a single L1 reference.

2.3.1 Integrating the MaGMa UCF tool into an existing environment

Some organizations have an existing security monitoring infrastructure with many default and custom use cases already present in the security monitoring tooling. These use cases should be added to the tool to provide detailed insight into the L1 and L2 threat categories that are being monitored. Adding use cases should be done by combining a bottom-up and a top-down approach.

The bottom-up approach is done by adding all existent monitoring rules into the L3 level of the MaGMa tool. For each of the rules, an identifier is created that maps the L3 rules to the corresponding L2 rule. If no match can be made, then an additional L2 use case should be added to the tool. Once this process has been completed, a top-down follow-up should be considered. This essentially means that gaps in the L2 layer are identified. The basic question to answer is: are there any threats that are not covered by the current L2 layer? If so, is this

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

due to a missing L1 layer? Any identified gaps should be fixed by adding L2 use cases and, if required, L1 cases to the tool. Corresponding L3 use cases should be created and implemented in the existing security monitoring infrastructure.

3 Management of Use Cases

Naturally, setting up the use case framework is a vital first step in achieving a structured security monitoring process. However, when the use case framework has been created, it also needs to be maintained. This is what use case management is for. Essentially, it is life cycle management for use cases.

3.1 Onboarding (plan and build)

For the onboarding of new use cases, the use case elements as described in chapter 2 should be made concrete. Stakeholders that provide input into the use case must be made part of the process to ensure proper alignment with these stakeholders. Once all relevant information has been identified, the use case can be documented.

Naturally, documentation is only the first step. To operationalize the use case, a planning should be created and the operational aspects of the use case should be implemented. These are activities usually carried out by security engineers. Once the build phase has been completed and the use case description is updated, the newly built use case can be handed over to the operational SOC manager. From this point on, analysts will be analyzing output from the monitoring rules and provide the engineers with the required input for tuning the monitoring rules.

3.2 Operational phase (run)

In the run phase of a use case, all operational elements are implemented and running as part of daily security operations. Concretely, this means that:

- Log sources have been added to the security monitoring systems and supplying the required information.
- Scope has been determined and implemented for this use case.
- Security incident response is known and documented.
- Roles and responsibilities for this use case have been formally documented and accepted. This includes roles and responsibilities for security incident response as well.
- Security monitoring rules have been implemented, tested and documented in the build phase and are now triggering security alerts

3.3 Maintenance (change)

There are several types of input that lead to changes within the use case. Most likely, these changes will be carried out at the implementation level, although changes at the threat and business levels will occasionally be required. This section identifies potential sources for input into change management for use cases. These can be divided into 2 main drivers:

1 Environmental drivers. These are changes to use cases resulting from changes in the organization.

2 Operational drivers. Additionally, operational drivers can lead to change as well. Red team testing as incidental input for improvement and lessons learned from incident response as a continuous input for improvement are important to consider. Threat hunting is discussed separately, as this is an extensive process.

3.3.1 Environmental drivers

Environmental drivers are the influences outside the SOC. Changes in the input previously described in the chapter on use cases and outlined in figure 1 are important factors to use case life cycle management.

Changes to the threat landscape

Threat actors may move to new modus operandi, or new threats may arise that need to be added to the use case framework. The threat management process should be connected to the use case management process to ensure changes in threat level are properly reflected in security monitoring.

Changes to the business

Businesses evolve continuously. Thus, it is possible that changes to the business layer of the use case may be required. These changes may also need to be reflected in operational implementation. For example, changes in demands regarding use case output may lead to new assets being monitored or new monitoring rules being implemented. Additionally, risk assessments conducted by the business may lead to new insights regarding risk. Such insights can also lead to changes within the use case framework.

Changes in rules and regulations

New or updated rules and regulations, mostly through introduced audit & compliance departments, may require changes to SOC use cases. For example, new output may be required and thus, new monitoring rules must be implemented to deliver such output.

Changes in the IT infrastructure

Changes in the IT infrastructure must be reflected in the operational layer of the use case. For example, new systems may need to be added and systems that were removed from the infrastructure must be removed from security monitoring systems as well. Ideally, this process should be automated. However, sometimes this is not possible, especially when new types of systems are added to the infrastructure. In such cases, additional parsers may need to be created to help the security monitoring systems understand the output from the new system or application. Additionally, use case updates will be required when major architecture changes are implemented. For example, new network routes may render intrusion detection rules useless.

3.3.2 Operational drivers

These are change drivers resulting from operational security activities. Such activities may be conducted within the SOC or by separate teams or departments.

Red team testing

Red team testing is a great way to simulate cyberattacks in a controlled fashion. This allows the SOC to test and validate the monitoring rules and the effectiveness of the incident response without actual risk to the organization. The findings from the red team test will lead to changes in the use cases, most often in the operational layer. Additionally, red team tests may lead to the development of new use cases.

Lessons learned from incident response

Besides red teaming, incident response activities originating from security alerts are also input into the use case management process. This loopback mechanism ensures that monitoring is effective and efficient. After all, breaches must be detected and false-positives should be avoided as much as possible. As with red team testing, incident response outcome can be used for enhancing existing use cases or creating new use cases. The latter is true when the initial trigger for incident response did not originate from the security monitoring process.

Threat hunting

Security monitoring, in general, is reactive in nature: events are generated in the infrastructure, monitoring rules or anomaly detection tools fire alerts and an incident response is mounted by the SOC or incident response team. Threat hunting is part of a more proactive approach to security monitoring. Instead of waiting for events to trigger, the hunting teams proactively search for signs of threats or potential threats in the infrastructure. For example, hunting teams may use event information available in the SIEM or analytics systems to search for aberrant patterns. Alternatively, hunting teams may use information regarding new 0-day attacks to determine which potential threats are facing the organization.

Hunting activities should always start out with a hypothesis. In the example of the 0-day, the hypothesis could be: "*the recently uncovered 0-day exploit can be used by cybercriminals to compromise client systems and gain foothold in our organization's infrastructure*". Testing that hypothesis could be done by attempting to reproduce parts of the kill-chain: what delivery methods are available for this exploit? Will this exploit work out-of-the-box on our client security architecture or are modifications required? Is it possible to create a persistent foothold? Figure 4 shows these actions as part of the killchain. Red blocks indicate activities that are not carried out within the threat hunting process, while green blocks indicate activities that are part of the process. Weaponization is presented in orange, as 0-day exploit code may be readily available.

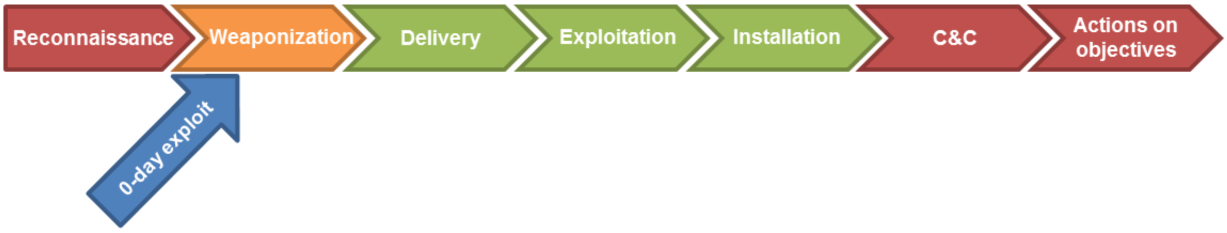

Figure 4: cyber killchain in hunting activities – 0-day hypothesis

Another hypothesis could be: “*cyber criminals use existing tools on our systems for lateral movement*”. These dual-use tools include powershell, net, rundll, wmic and scheduled tasks. Figure 5 shows the killchain relevance for this hypothesis.

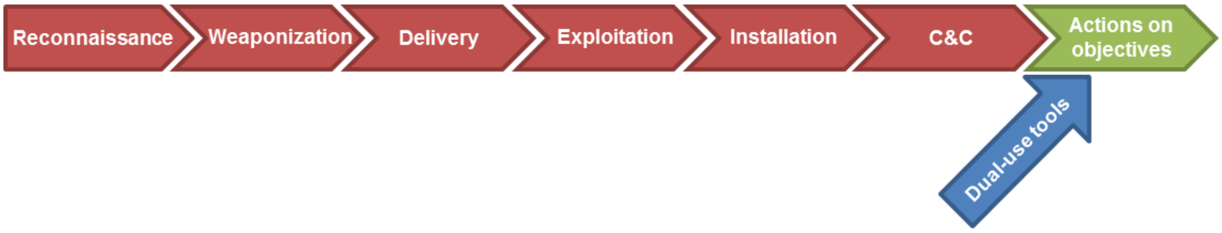

Figure 5: cyber killchain in hunting activities – dual-use hypothesis

This type of hunting activity overlaps with red teaming activity. When this type of hunting is performed continuously, this is called purple teaming.

Hunting activities need not be as sophisticated as described in the 0-day example. Lab research regarding new malware to observe behavior and derive indicators can also be part of hunting. Using indicators from threat intelligence feeds and determining whether or not these indicators are present in the infrastructure is another example of hunting activities. Note that more elaborate hunting will focus on higher levels in the pyramid of pain [[4]], in particular on tactics, techniques and procedures (TTP). Hunting for Indicators of Compromise (IoCs) should be avoided as much as possible, as this can be automated using threat feeds directly into security monitoring systems.

In any case, hunting activities may lead to new insights about threats and security monitoring and are therefore input in the lifecycle management for use cases. Similar to red team testing and security incident response, hunting activities can be input for changing existing use cases or creation of new use cases.

3.4 Offloading (decommission)

When use cases are no longer required, an offloading process should be followed to remove the use case from the framework at each of the layers. The same inputs that feed the change management of the use case may trigger the decommissioning of the use case. The offloading process should focus on:

- Removing operational elements from the security monitoring system
- Decommissioning any specific procedures associated with the use case
- Informing the owner of the threat landscape that the use case regarding a particular threat is being removed from the framework
- Informing business stakeholders and recipients of reports generated specifically for that use case that of the removal of the use case

3.5 Overview

Figure 6 provides an overview of the life cycle management process and the input received.

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

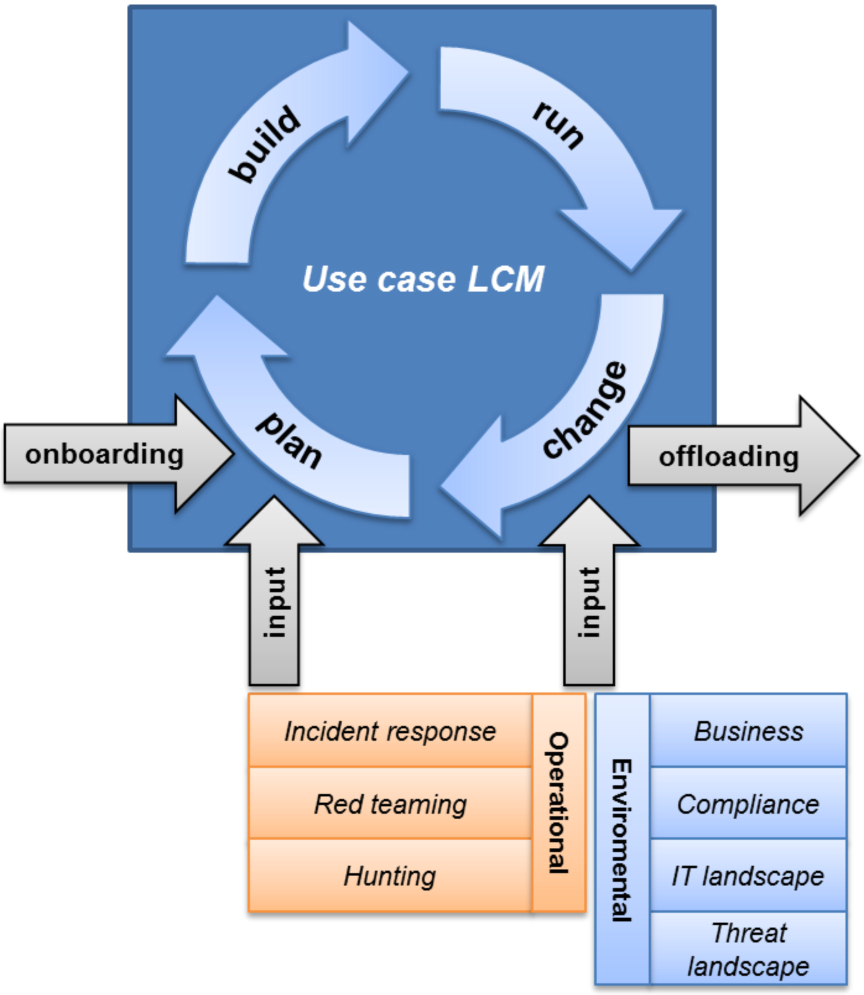

Figure 6: Use case management overview

4 Growth (Capability and Maturity)

With a use case framework in place and several use cases implemented, an important question arises: how to move towards more mature and more effective security monitoring? To answer this question, 2 aspects must be evaluated. These aspects are capability and maturity. These aspects may seem to be closely related, but need to be evaluated separately. A highly capable SOC may have a very ad-hoc approach with little documentation and standardization. A highly mature SOC may have very limited capabilities. It is recommended to ensure SOCs grow both in capability and maturity. This applies to use cases as well as other parts of security operations. This section deals with capability and maturity for use cases and what activities should be employed at each level.

4.1 Capability

Use case capability deals with how well a SOC is capable of detecting security threats. Highly capable SOCs will have multiple layers of detection (detection-in-depth) and advanced security tooling for security monitoring. Of course, these SOCs also have highly professional personnel to ensure that their advanced technical capability is backed with equally advanced analyst knowledge.

Growing in use case capability can be done by extending in two directions:

- Framework completeness. This type of growth is aimed at the growth in number of use cases. By implementing more use cases, the total capability level will increase. However, the number of use cases implemented is not the only direction, as the completeness of each use case itself is also an important factor.
- Use case completeness. This type of growth is aimed at growth in operational implementation level of a single use case. This can be done in multiple ways. One of those ways is increasing capability by increasing the number of assets (log sources), provided that existing rules will automatically process events from those sources. Another way is by adding new types of assets to the security monitoring systems. Lastly, growth can be achieved by increasing the number of security monitoring rules. In many cases, all such activities will need to be carried out simultaneously in order to increase the use case implementation level.

Figure 7 outlines these directions in capability growths and potential growth paths that can be followed. One approach is to first focus on completing the framework; another approach focusses on completing the operational implementation for a limited number of use cases. A third approach is a linear approach where both directions are balanced out. In any case, decisions on which use cases to implement first should be risk-driven. This means that high-risk use cases should be implemented first. The implementation level should be matched to a satisfying level of risk reduction. Thus, the most effective risk-based approach to capability growth is by first using risk levels to identify high-risk use cases and then using residual risk to determine the required implementation level per use case before moving to the next use case.

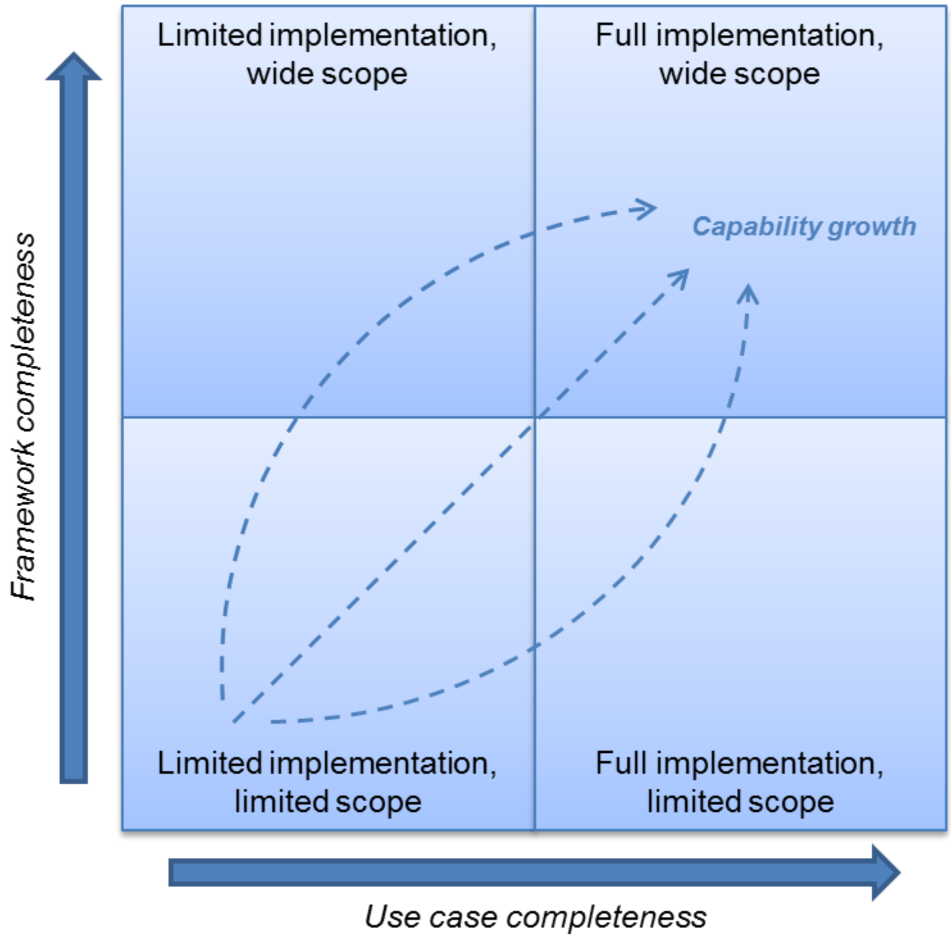
Limited implementation, wide scope
Full implementation, wide scope
Capability growth
Limited implementation, limited scope
Full implementation, limited scope
Framework completeness
Use case completeness

Figure 7: alternative routes to higher use case capability

4.2 Maturity

Maturity deals with how well a SOC is able to provide effective security monitoring continuously and how well it can adequately deal with the output of security monitoring. A highly mature SOC will provide consistent and repeatable output in terms of high-quality reports and effective incident response. Additionally, a highly mature SOC will improve continuously and consistently through a standardized process. Therefore, standardization and documentation are significant aspects of maturity.

According to the CMMI [[5]], the following maturity levels can be identified:

1. Initial
2. Managed
3. Defined
4. Quantitatively Managed
5. Optimizing

Each of these levels will be discussed within the context of SOC use cases.

[5] https://www.sei.cmu.edu/cmmi/

4.2.1 CMMI Level 1: initial

The first maturity level in the CMMI is characterized by ad-hoc and chaotic operations. In terms of use cases and security monitoring, a level 1 mature SOC is a SOC where no standardization is applied. Security alerts are not dealt with consistently, but depend heavily on the individual analyst processing the alert. There is no formal reporting, nor clear assignment of roles and responsibilities. Use cases will be documented scarcely, if at all, and measurement of implementation levels is not performed. A structured use case framework, such as the one outlined in this paper, is not in place. A SOC performing at maturity level 1 is unlikely to adequately defend its organization against cyber threats.

4.2.2 CMMI level 2: managed

A level 2 (managed) SOC moves away from ad-hoc security monitoring to structured security monitoring. Output and service delivery are formalized and service level agreements may be set up (depends on the organization). A process is in place to measure whether the security monitoring service delivery is in line with expectations and to deal with any shortcomings. A use case framework is likely to be in place, although it is unlikely that such a framework is adequately aligned with all relevant stakeholders.

4.2.3 CMMI level 3: defined

In maturity level 3, the SOC moves to a fully structured use case framework. All relevant documentation regarding use cases is formalized and aligned with relevant stakeholders. Service delivery regarding security monitoring is consistent, measured and reported on. The output from these measurements is used as input for improvement, but this 'feedback loop' is not continuous. The full use case template, as outlined in chapter 2, is used for every use case implemented in security monitoring. This provides a means for measuring the implementation level for each use case. This information can be used by the SOC to inform stakeholders. Since security monitoring can help to mitigate risk associated with threats (and thus quantify residual risk), the measurement of use case implementation becomes part of the threat management process.

4.2.4 CMMI level 4: quantitatively managed

Maturity level 4 is characterized by the fact that goals are set for quality in the use case management process. Concretely, this means quality of output delivered by the SOC, in reports, incident management, documentation and communication. Also, the performance of the use case management process and completeness of the framework should be measured. These measurements are used to improve the process and framework in a structured manner.

4.2.5 CMMI level 5: optimizing

Finally, maturity level 5 is characterized by continuous improvement. This improvement does not only come from operational findings, but also by continuously aligning with relevant stakeholders. This ensures that the security monitoring services that are delivered by the SOC match expectations from those stakeholders. At this maturity level, the SOC, through effective security monitoring and structured use case management, becomes a strategic partner for threat management and risk reduction.

4.3 Final remarks on capability and maturity

There is no such thing as an ‘optimal maturity level’. This depends entirely on the organization and the ambition level of the SOC. While higher maturity levels provide for better SOC output, more effort and budget is required as well. Such effort may not be possible for smaller SOCs with fewer personnel. Thus, an ambition level for the SOC should be decided on. This ambition level is the maturity goal for SOC use case management. While a generic optimal maturity level may not exist, it can be stated that acting at CMMI level 1 provides little added value to the organization. Thus, level 2 should be the minimal maturity level for SOC use case management.

Similarly, there is no optimal capability level either. Growth in completeness requires effort and time. It must be noted that maintaining a capability level requires effort as well. Tuning security monitoring rules and adapting to changes in the environment is a continuous activity that will be more cumbersome in a highly capable security monitoring system. More rules and more assets simply require more maintenance. Lack of maintenance will result in increasingly high false-positive rates that are time consuming and degrade the quality of the use case and security monitoring process as a whole. More on false-positives and other metrics can be read in the next chapter.

5 Metrics & Assessment

To determine the effectiveness of the use case framework in delivering optimal security monitoring, metrics are required. The MaGMa framework should be regularly assessed using these metrics. The outcome of such assessments can be used to identify strong and weak spots within the framework. Note that the focus here is on quantifiable metrics. Three types of metrics have been identified: embedded metrics, control metrics and output metrics. Each of these types will be examined.

5.1 MaGMa UCF embedded metrics

There are several metrics embedded into the UCF. These metrics are used to provide steering information regarding the effectiveness of the use case framework. The following embedded metrics can be found in the UCF:

- Effectiveness. This value is set manually in the framework at the L3 level. Use this metric to indicate the effectiveness of the detection mechanism. For example, a proxy inspecting traffic is much less effective if it is not able to inspect HTTPS traffic.
- Implementation. This value is set manually as well at the L3 level. Use this metric to indicate is how well a detection mechanism has been implemented. For example, the implementation level of an IDS is much lower if the ruleset is incomplete or has not been tuned.
- Coverage. Just like weight, this value is set manually in the framework at the L3 level. Use this metric to indicate the level in which this detection mechanism covers the use case. For example, a use case focused on firewall events has less coverage if not all traffic is routed through the connected firewall.
- Weight. This metric is a calculated overall score of effectiveness, implementation and coverage.
- Potential. This metric is a calculated value that indicates how much improvement can be gained by investing in coverage and implementation. Thus, high potential for growth is awarded to monitoring rules that are potentially very effective, but have insufficient coverage and implementation.

The MaGMa UCF tool also provides some high-level statistics. The following overviews are found in the tool:

- Average values for all embedded metrics.
- The number of operational use cases per threat. The MaGMa UCF tool provides for measurement of the number of L2 and L3 use cases per L1 use case. Thereby providing insight into the number of use cases assigned to each threat. Since these threats are connected to business and possibly compliance drivers, this relation can be used to show that the SOC is in control.
- The number of identified business and compliance drivers. This information can be used to show alignment of the use cases to business and compliance.
- The number of unique detection mechanisms as well as the number of unique log sources.
- Effectiveness, implementation and coverage per L1 use case.
- Weight and potential per L1 use case.

Naturally, since all information is in Excel format, other overviews can be easily created. Possible additional overviews to consider are:

- Number of use cases per detection mechanism to determine added value of detection mechanisms to security monitoring.
- Number of use cases per log source to determine the relevance of certain log sources to security monitoring.
- All L1 overviews above created for the L2 level.

A final note: if an organization finds it difficult to quantify the embedded metrics, simplifying the metrics to qualitative metrics (e.g. none, low, medium, high, full) can be considered. Note that this will make reporting on MaGMa control metrics (see next paragraph) more difficult. Qualitative metrics are not supported by default in the MaGMa UCF tool, so quantitative values corresponding to qualitative indicators could be used. For example, 'none' could be set to 0%, 'low' to 25%, medium to 50%, etc.

5.2 MaGMa control metrics

Control metrics provide governance information on the framework itself: how is the framework changing over time? What maturity level are use cases currently on? This paragraph outlines the control metrics that can be used in the MaGMa framework.

5.2.1 Changes to the framework

This metrics provides information on the number of changes to the framework in a set period of time (for example: 1 quarter). This provides information on the stability of the framework.

Framework changes should be reported as a quantity of changes per use case. It makes sense to report this as the L2 level (as the L1 level is fairly static), although underlying numbers on the L3 level can be used to determine which rules have undergone most changes.

Note that this metric requires tracking changes to the framework and thus some form of change management procedure. If no such procedure is in place, an alternative metric can be used that reports on the number of rules that have been modified in the last quarter. The 'last modification date' should be updated as an absolute minimum.

5.2.2 Growth in number of use cases

Growth in number of use cases is another change indicator, in this case a change in monitoring scope. This indicator can be used to quantify efforts in extending security monitoring capabilities to the SOC manager.

The growth in number of use cases should be expressed as an absolute value. Similar to the previous metric, it makes most sense to report this on the L2 level.

5.2.3 Growth in weight

In the framework, weight is a calculated value, using the product of effectiveness, implementation and coverage as input (see embedded metrics). A higher weight value means that security monitoring is becoming more effective. Growth in weight can be reported high-level (change to weight average), or in more detail on the L2 and L3 levels.

Similar to the previous metrics, a full snapshot is required for L3 detailed reporting.

5.2.4 Changes to potential

This signifies growth potential of the use case relative to its effectiveness. Similar to weight, potential can be reported high-level (change to weight average), or in more detail on the L2 and L3 levels. Use cases with high potential are candidates for improvement.

If detailed reporting is required to determine which L3 use cases have shown the most growth or which use cases have been impacted adversely, a full snapshot of the previous assessment is required.

5.3 MaGMa output metrics

The last category of metrics is metrics on the output of the MaGMa framework.

5.3.1 Number of alerts

For each use case, the number of alerts resulting from the security monitoring systems should be reported on. These numbers are indicators of which use cases are triggered often and which are triggered rarely. Combined with the information on false-positives, this information is input into the threat management process.

The number of alerts should be expressed as an absolute value per use case. This should either be reported on L1 or L2 use cases, depending on the level of detail of the report (i.e. are you reporting to management, or to the SOC manager?). L1 and L2 use cases with a high level of alerts, and a high false-positive ratio are candidate for tuning at the L3 level.

5.3.2 Number of incidents

Not every alert will also lead to an incident response. Therefore, the number of incidents should also be reported on.

Just like with alerts, the number of incidents should be expressed as an absolute value per use case. This should either be reported on L1 or L2 use cases, again depending on the level of detail of the report.

5.3.3 False-positive ratio

The false-positive provides an indication of the quality of the security monitoring system. False-positives occur when a security alert is triggered, while there's no actual security incident. While false-positives are fact-of-life in any SOC, high numbers of false-positives should be avoided at any cost.

False positives can be expressed as the ratio between the total number of alerts and the number of alerts relating to incidents:

$$FP\ percentage\ = 100 - 100\ \times \left(\frac{number\ of\ alerts\ relating\ to\ incident}{total\ number\ of\ alerts}\right)$$

5.3.4 Number of false-negatives

Information on false-negative, similar to the false-positive ratio, provides insight into the quality of the operational security monitoring rules. False-negatives occur when an actual incident has taken place that is within scope of one of the use cases, but was not detected by any L3 operational monitoring detection mechanisms. This could be due to improper tuning of correlation rules, incorrect configuration of the correlation rules or simply because no detection mechanisms exists as the attack vector was either unforeseen or not implemented yet.

False negatives should be reported as a quantity of missed security incidents. Note that false-negatives may occur continuously without being noticed. Usually, triggers that false-negatives have occurred come from the security incident response process.

6 Best practices for use case management

During the creation of MaGMa, the focus group has identified some best practices that can be applied to use case management. These best practices are described hereafter.

Challenge the customer

Here, the customer is the SOC customer; the department head requesting the onboarding of a new use case. Always challenge the customer so that they are able to accurately define the risk the use case is meant to mitigate. Also challenge the customer to think about remediation and response. Thinking ahead will save precious time in case of actual incidents.

Make sure incident response is known beforehand

Never start the onboarding procedure if follow-up for a specific use case it unknown. If it is not possible for the security incident response team to mount an effective response process, a use case will have limited or no added value.

Obtain required mandate for incident response

In addition to the previous best practice, it is important that the security incident response team has the required mandate beforehand to ensure effective incident response. Obtaining mandate outside of business hours during an incident may prove to be difficult and can slow down the incident response process.

Use playbooks for security incident response

To support the SOC in analysis and incident response, a playbook can be used. This allows for structuring of required information about follow-up for use cases. More on playbooks can be found here:
http://blogs.cisco.com/security/using-a-playbook-model-to-organize-your-information-security-monitoring-strategy

Always evaluate incident response

Incident response evaluation should always be carried out after security incidents have been resolved. These evaluations should focus on actions performed in incident response process, and the effectiveness and efficiency of those actions. This information can be used to improve the security monitoring process as well as the security incident management process. Such evaluations can also be used to show added value of the SOC or security incident response team (and the security monitoring process indirectly) in limiting impact of security incidents.

Provide sufficient context for security monitoring

Security monitoring can be greatly improved by adding context regarding users, assets, business processes and threat intelligence. Consider such enrichments when designing and deploying L3 use cases for optimal security monitoring efficiency.

Use a grace period

Before the newly implemented use case is made fully operational, it makes sense to apply a grace period. In this grace period, the security alerts resulting from newly implemented security monitoring rules are aggregated. Periodically (for example, daily or weekly), the aggregated results are analyzed and the security monitoring rules are tuned accordingly. After a predefined period of time, the security monitoring rule is operationalized and the number of false-positives should be limited.

Be vigilant about false-positives

False-positives can lead to security incidents, because analysts may not be able to differentiate actual threats from events that can be ignored. Overwhelming numbers of false-positives reduce analyst capability to identify actual threats greatly.

Go beyond detection and response where possible

Showing added value of the SOC is imperative for retaining SOC budget. Thus, ensure that activities from the SOC are seen by relevant stakeholders. Go beyond detection and response and provide input for protection against threats to business, IT and development teams. Inform stakeholders of new threats possibly facing the organization. Help the organization in implementing security measures where possible. Simply put: be proactive about security and do not sit around waiting for incidents to happen.

Support risk management process

The SOC can support the risk management process by providing information on the threats facing the organization and the impact of such threats. This way, the SOC plays a role in the strategical process of managing operational risk. Threat reporting is also important to achieve SOC visibility.

Have the UCF audited or reviewed

Regularly let another department (for example: audit or security management) review the current state of the UCF. By having someone outside of the SOC provide a ‘fresh’ perspective on the UCF, additional points for improvement can be identified.

7 Conclusion

Managing security monitoring in enterprise environments can be a difficult task due to the sheer complexity and diversity of the infrastructure. The environment changes as well, with new threats emerging and new or updated laws and regulations demanding a versatile approach to security monitoring. To remain in control with such a changing landscape, use cases must be created and managed in an orderly fashion. This is where the MaGMa UCF comes in.

The MaGMa UCF provides a framework for management of use cases. It supports insight into the existent use cases, categorizing these use cases and additionally helps to identify potential gaps in security monitoring. Also, the framework is used to measure use case effectivity and thereby supports growth and maturity of the SOC as a whole.

In this document, the MaGMa framework and all of its features have been outlined. This document provides guidance for the implementation of the supporting MaGMa UCF tool into the organization. The tool itself should be used in a flexible manner and tailored to your organization. Add use cases, additional metrics, overviews and graphs where needed to optimize added value to your SOC. Filling the tool with all of your existing use cases may seem a daunting task at first. Therefore, it is advisable to start out small, by simply adding the use cases and mapping them to existing L2 use cases. In a next step, all of the log sources and detection technology can be added. Following this step, the coverage metric can be set and so forth. This way, added value can be obtained quickly. The details can be filled in once the framework outline is up and running. These details provide additional steering and quality information that is essential in a more mature security monitoring process.

Finally, the most important thing the MaGMa UCF provides is the capability to be in control over your security monitoring process and the alignment of security monitoring to business and compliance needs. The framework provides the ability to prove to your stakeholders that the SOC is in control and adequately managing and decreasing risk in the enterprise.

MaGMa

Use case framework

FI-ISAC NL Publication

©2017