Visualization for information system security monitoring

Damien Crémilleux

UNIVERSITE BRETAGNE LOIRE / MATHSTIC

CentraleSupélec

Thèse de Doctorat

Visualization for information system security monitoring

Visualisation pour la supervision de sécurité des systèmes d'information

Par :
Damien CRÉMILLEUX

CentraleSupélec
COMUE Université Bretagne Loire

École doctorale N° 601
Mathématiques et Sciences et Technologies
de l'Information et de la Communication
Spécialité : informatique

Thèse présentée et soutenue à Rennes, le 15 février 2019
Unité de recherche : IRISA
Thèse n° : 2019-02-TH

Rapporteur avant soutenance :

Hervé DEBAR	Professeur, Télécom SudParis
Guiseppe SANTUCCI	Professeur, Sapienza Università di Roma

Composition du jury :

Présidente :	Isabelle CHRISMENT	Professeur, Télécom Nancy
Examinateurs :	Hervé DEBAR	Professeur, Télécom SudParis
	Benjamin MORIN	Responsable de la division détection, ANSSI
	Giuseppe SANTUCCI	Professeur, Sapienza Università di Roma
Dir. de thèse :	Christophe BIDAN	Professeur, CentraleSupélec
Co-dir. de thèse :	Frédéric MAJORCZYK	Expert technique, DGA-MI et CentraleSupélec
	Nicolas PRIGENT	Expert technique, LSTI

Abstract

Security operations centers (SOCs) are the central place for the security of information systems. Two distinct teams of security analysts work together in a SOC to detect, analyze, and respond to security incidents. Tier 1 analysts are the first to receive security alerts and dispatch them to Tier 2 analysts for further investigations if needed. Despite the evolution of techniques and procedures over time, there remain some significant difficulties impacting the efficiency of SOCs. This situation results in failure to timely react to real attacks.

First of all, this thesis highlights and classifies the limitations of SOCs into two categories, technology and process. Tier 1 analysts are particularly affected by these limitations, resulting in a high burnout rate and strenuous working conditions. Visualization tools have been proposed to fulfill the tasks accomplished in SOCs. We identify that these tools are often not in adequation with the needs of security analysts, showing a lack of comprehension of their work and constraints.

We address the problem of the high number of IDSes alerts, the repetition of the same tasks, and the lack of creativity with VEGAS (Visualizing, Exploring and Grouping AlertS), a visualization and classification tool. VEGAS is based on a combination of data summarization and visualization. We compare the methods for dimension reduction and our proposition uses principal component analysis as a first step to produce a two-dimensional scatterplot that visually correlates the IDSes alerts. Our tool allows Tier 1 analysts to explore the various fields of similar alerts, to analyze them quickly, and to generate meaningful rules to group IDSes alerts. Our evaluations with a case study and experts have demonstrated that VEGAS is useful to quickly detect similar IDSes alerts and group them efficiently.

After addressing the problem of IDSes alerts triaging, we tackle the remaining process limitation of the lack of feedback alongside the limitations of threat escalation and rhythm of networks. We propose TheStrip, a process with a visual tool, to enhance the collaboration between Tier 1 and Tier 2 analysts. Our process creates a feedback loop between Tier 1 and Tier 2 analysts and improves rules to define security meta-events. Our concept of security meta-events adds time and collaborative features to enable better groups and the creation of attacks scenarios We propose a visualization tool to support this new process. This tool is organized around a timeline view, providing a quick perception of the context and easy reconstruction of attack scenarios.

Résumé

Notre monde est de plus en plus dépendant des systèmes d'information connectés. En 2018 il y avait plus de 3,66 milliards d'internautes et ce chiffre devrait dépasser les 4,13 milliards en 2021 [1]. Il est indispensable de veiller à la sécurité de ces systèmes d'information, face à des attaquants de plus en plus organisés et disposant de budget et de moyens de plus en plus importants. Les mesures de sécurité préventives s'avèrent insuffisantes pour assurer cette sécurité, les attaquants finissent par trouver des failles pour s'introduire au sein des systèmes ou les endommager. La plupart des systèmes d'information sont aujourd'hui assistés par un SOC (*Security Operations Center*, centre opérationnel de sécurité), afind'assurer une sécurité réactive avec la supervision et la gestion des incidents de sécurité.

L'activité principale d'un SOC est donc le triage et l'analyse des évènements de sécurité du système d'information, et en particulier des alertes provenant des systèmes de détection d'intrusion (IDS). En moyenne, un SOC collecte plusieurs milliers et même plusieurs millions d'évènements par jour, et a pour objectif de trouver ceux étant symptomatiques d'une véritable intrusion. En effet, la grande majorité des alertes reçues par un SOC sont des faux positifs et ne témoignent pas d'une véritable menace pour le système d'information. Ce nombre élevé d'alertes, associé à leur taux de non pertinence, est un problème car les vraies attaques ne sont pas détectées ou alors très tardivement. En outre, ce problème engendre une pression sur les analystes travaillant dans les SOCs. Ces analystes ont un temps très limité pour décider de la sévérité et de la véracité des alertes. Ils souffrent d'un taux de burn-out important : la durée moyenne de travail d'un analyste est de trois ans au sein d'un SOC.

Cette thèse a pour objectif d'améliorer la supervision au sein des SOCs. Les principales contributions que nous apportons sont :

- Une revue des SOCs et de leur organisation, aboutissant à l'expression de leurs limites. Ces limites sont de deux catégories différentes, technologie et processus.

- VEGAS [2, 3], un outil de visualisation et de classification qui permet aux analystes de première ligne dans les SOCs de regrouper les alertes grâce à leur représentation en analyse en composantes principales.

- TheStrip [4], un nouveau processus associé à un outil pour améliorer la collaboration au sein des SOCs.

En plus de ces contributions, cette thèse propose une revue des solutions de visualisation pour la sécurité selon la connaissance situationnelle (*situational awareness*). Nous avons ajouté la collaboration entre les analystes de sécurité au concept de connaissance situationnelle afin de retranscrire complètement les différents cas d'utilisation rencontrés dans un SOC.

Centre opérationnel de sécurité

Les SOCs sont un élément central pour la sécurité des systèmes d'information. Ils doivent remplir des nombreuses missions, dont la plus importante est le triage et l'analyse en temps réel des évènements de sécurité. Les SOCs sont organisés autour de deux équipes d'analystes, les analystes Tier 1 et les analystes Tier 2. Les analystes Tier 1 sont situés en première ligne et disposent de quelques secondes ou minutes pour décider de la gravité d'une alerte. Ils utilisent comme source de données les alertes IDS. En cas de suspicion, l'alerte est envoyée à des analystes Tier 2 pour une analyse plus approfondie et l'élaboration de la réponse face à cette menace. Les analystes Tier 2 disposent de plus de temps pour effectuer leur analyse et ils utilisent toutes les sources de données à leur disposition tels que les logs des serveurs, des clients ou encore des applications.

Malgré l'évolution de leur mission et de leurs outils, les analystes au sein des SOCs souffrent d'inconvénients liés au processus et aux technologies utilisés. Cette liste de limites est une des contributions de cette thèse. Les limites de type technologique sont:

- De nombreuses données et sources de données qui ne sont pas nécessairement liées. Même avec les seules alertes IDS comme source de données principale, les analystes Tier 1 doivent faire face à un volume d'évènements de sécurité important. Ce problème existe aussi pour les analystes Tier 2, le nombre de données qui leur est transmis étant encore plus important. De plus, les sources de données sont variées telles que les antivirus, les alertes IDS, le trafic réseau. Comme ces sources de données ne sont pas nécessairement liées entre elles, il est nécessaire d'avoir une expertise pour chaque type. Enfin, cette caractérisque rend la corrélation et l'exploration difficiles.

- La progression de la menace. Il est particulièrement important de savoir si un évènement est isolé ou bien s'il fait parti d'un scénario d'attaque plus important. La connaissance du contexte actuel, des menaces et des incidents, est nécessaire pour que les analystes de sécurité puissent élaborer une réponse efficace.

- Le rythme du réseau. Les analystes de sécurité sont familiers du rythme du réseau qu'ils surveillent, ils connaissent les évènements fréquents et les conséquences qu'ils engendrent. La compréhension de ces évènements et du nombre classique d'erreurs au sein du système est insuffisament exploité.

Nous nous sommes aussi intéressés au processus du travail mené dans les SOCs et nous en avons indentifié les limites suivantes :

- La répétitivité des tâches. Les analystes Tier 1 accomplissent des tâches répétitives en suivant des procédures pré-établies. Cet aspect est également valable pour les analystes Tier 2 car ils doivent traiter les mêmes types d'évènements envoyés par les analystes Tier 1. Cela résulte en une perte de temps et une appréciation diminuée pour le travail accompli par l'analyste Tier 1.
- Le manque de retour. Une fois leur décision prise, les analystes Tier 1 perdent la trace de leurs actions. Ils ne sont pas notifiés du résultat de l'analyse faite par les analystes Tier 2 et ainsi ne savent pas s'ils ont pris la bonne décision.
- Le manque de créativité. Les analystes Tier 1 suivent des procédures qui limitent sevèrement leur créativité et n'en dévient pas, avec pour conséquences des difficultés pour réagir correctement à de nouveaux types d'attaques.

Les analystes Tier 1 sont particulièrement touchés par ces difficultés, engendrant des conditions de travail difficiles. Nous pensons que la visualisation de sécurités, associée à une meilleure collaboration au sein des SOCs, est une réponse à ces problèmes.

Visualisation de sécurité

La visualisation de sécurité au sein des SOCs est utilisée avec plusieurs objectifs. Dans cette thèse nous avons passé en revue les solutions de visualisation de sécurité selon le concept de connaissance situationnelle. La connaissance situationnelle se définit selon trois étapes qui sont la perception, la compréhension, et la projection, ainsi que cinq cas d'utilisation : la supervision, l'inspection, l'exploration, la prévision, et la communication. Nous proposons un sixième cas d'utilisation qui est la collaboration car collaborer est une nécessité pour les analystes travaillant au sein d'un SOC. La collaboration entre les analystes de sécurité, tout comme la communication, est un cas d'utilisation transverse aux trois phases de la connaissance situationnelle.

Cette revue des solutions de visualisation de sécurité a mis en avant les différentes techniques avec leurs avantages et leurs inconvénients. Les outils de supervision dédiés aux analystes Tier 1 utilisent des représentations visuelles simples pour que les analystes perçoivent l'état actuel du système d'information. Cependant, le passage à l'échelle de ces solutions est souvent limité, et ces solutions ne sont pas toujours capables de supporter la charge d'une utilisation réelle. Les outils d'inspection et d'exploration proposent plus d'interactions et sont capables d'exploiter plus de sources de données. Ils sont utilisés par les analystes Tier 2, et sont de notre point de vue encore insuffisants par rapport aux limites que nous avons exposées concernant la progression de la menace et le rythme du réseau. Les solutions de prévision sont moins présentes dans la littérature scientifique, et exploitées généralement dans les SOCs ayant une certaine maturité sous le nom de *threat intelligence*.

Alors que les analystes ont besoin de collaborer et de communiquer au sein d'un SOC, il semble y avoir un manque de solution visuelle pour effectuer ces tâches. Nous sommes persuadés qu'une manière plus efficace de gérer la large quantité de données est de rendre

la collaboration entre les analystes de sécurité plus simple par une meilleure organisation du processus de travail à travers la visualisation.

VEGAS

Pour faciliter la tâche de triage des analystes Tier 1, nous proposons un outil appelé VEGAS (« Visualizing, Exploring, and Grouping Alerts ») [2, 3]. Cet outil permet de visualiser les alertes IDS dans un espace en deux dimensions, afin de pouvoir ensuite facilement les regrouper. En effet, les alertes de sécurité sont composées de plusieurs dimensions telles que le port source, le port destination, ou encore les adresses IP, et leurs représentations graphiques directes sont difficiles à appréhender pour un analyste n'ayant qu'un court instant pour effectuer sa tâche. C'est pourquoi un algorithme permettant de transposer les alertes suivant une représentation en deux dimensions a été choisi. Cet algorithme doit être efficace d'un point de vue rétention d'information (des alertes similaires doivent être proche sur l'espace deux dimensions), doit passer facilement à l'échelle sur des milliers de données dans un temps contraint, et doit être non supervisé. Ces pré-requis ont orienté notre choix sûr l'analyse en composantes principales.

VEGAS propose ainsi de visualiser les alertes en deux dimensions grâce à l'analyse en composantes principales. Cette représentation permet à l'analyste de rapidement trouver des ensembles d'alertes ayant des caractéristiques similaires. Les ensembles d'alertes identifiés sont ensuite diagnostiqués à l'aide de représentations visuelles des caractéristiques de ces alertes. L'analyste Tier 1 peut sélectionner les caractéristiques qui sont pertinentes pour ce groupe. Une règle est générée afin de pouvoir rediriger directement les alertes similaires, passées comme futures, diminuant ainsi le flux de nouvelles alertes non diagnostiquées pour l'analyste travaillant en première ligne. Cette combinaison de l'analyse de données et de la visualisation est le cœur de notre proposition.

VEGAS permet de répondre à la limite technologique du nombre élevé d'alertes pour les analystes Tier 1. En outre, VEGAS est aussi une réponse concernant la répétitivité des tâches et le manque de créativité touchant ces analystes. En effet, de part la création de règles redirigeant les alertes similaires, les analystes Tier 1 n'ont plus à traiter plusieurs fois les mêmes alertes et peuvent se concentrer sur les alertes jusque là non rencontrées. De plus, la création de règles suite à une analyse rapide des alertes permet de répondre à la monotonie de la tâche due au manque de créativité.

Nous avons tout d'abord évalué VEGAS à l'aide d'un cas d'utilisation (challenge VAST 2012 [5]) puis nous avons interrogé douze experts-analystes de sécurité, de sexe masculin avec une à dix années d'expérience dans ce milieu. Les retours de ces tests sont positifs. Tout d'abord ils confirment la pertinence du problème de triage des alertes que nous cherchons à résoudre. Ensuite, les experts ont positivement évalué les visualisations et les interactions proposées par notre outil. Enfin, ils ont déclaré que VEGAS permet une amélioration de la productivité des analystes Tier 1 au sein des SOCs.

TheStrip

Les entretiens avec les experts nous ont permis de mettre en avant le manque de collaboration au sein des SOCs, et les limites qui en dérivent. Pour remédier à celles-ci, nous proposons TheStrip [4], un nouveau processus de collaboration entre les différents analystes et un outil associé pour mettre en œuvre ce processus.

Notre processus est basé sur le concept de méta-évènements, avec l'introduction d'une boucle de retour entre les analystes Tier 1 et les analystes Tier 2. Nous étendons les règles définies avec VEGAS pour regrouper les évènements en méta-évènements, permettant une division dans le temps des évènements capturés par une règle et une meilleure collaboration entre les analystes. Une fois la règle créée par un analyste Tier 1, les analystes Tier 2 peuvent ajouter des caratéristiques liées au temps, modifier la règle si besoin, et renseigner les différentes personnes coopérant dans la résolution de ces évènements de sécurité. Ce processus permet aux évènements identifiés par les règles d'être redirigés directement vers les analystes dédiés à ce type d'évènement. En outre, les analystes Tier 2 peuvent relier les meta-évènements afin de retracer les scénarios d'attaques.

Nous avons développé un outil pour mettre en oeuvre ce processus. Cet outil permet une perception rapide de la situation actuelle système à l'aide d'une vue sous forme d'une *timeline*. Les interactions de cette vue permettent aux analystes de visuellement corréler les méta-évènements et de reconstruire facilement les scénarios d'attaques. Des vues dédiées proposent la visualisation et la modification des règles pour les méta-évènements et les scénarios.

Ainsi TheStrip permet de répondre au manque constaté de retour envers les analystes Tier 1. Ceux-ci sont avertis des changements effectués sur les règles et peuvent comprendre les raisons de ce changement. Les fonctionnalités offertes par l'outil sont également une réponse aux limites de la progression de la menace et du rythme du réseau grâce aux représentations proposées.

Conclusion

Au cours de cette thèse, nous avons cherché à améliorer la supervision de sécurité. Nous avons explicité les limites des SOCs et montré que les solutions actuelles de visualisation de sécurité ne permettent pas de répondre complètement à ces limites. Nous avons proposé des solutions pour y remédier avec VEGAS et TheStrip. VEGAS assiste les analystes Tier 1 pour le triage des alertes avec des représentations adaptées et la création de règles. TheStrip renforce la collaboration entre les différents analystes au sein des SOC grâce à un nouveau processus associé à un outil de visualisation.

Acknowledgements

First of all I would like to express my gratitude to my advisors Christophe Bidan, Frédéric Majorczyk, and Nicolas Prigent for their guidance throughout this thesis. I had a great freedom in the choice of my research subjects and I appreciate all your contributions of time and ideas to make my PhD experience stimulating.

Completing this work would have been all the more difficult were it not for the support provided by the other members of the CIDre team. You have been companions in countless informal discussions, billiard games and pints, and a source of friendships as well as good advice. I thank you for making this research journey a pleasant experience.

During this thesis, I encountered some technical difficulties and I am especially grateful to Christopher Humphries for his guidance in these moments. You have been so helpful.

I am also grateful to the members of the DGA who took the time to test my work. Your feedback and comments were invaluable in making this thesis more relevant to the needs of security experts.

I want to express my deepest gratitude to my parents and my family for their endless support and constant encouragement I have gotten over the years. During this thesis, I have been amazed by your willingness to proof read countless pages of my work. Finally, I would like to thank Lauriane for her continued support and encouragement, especially during the final stages of this PhD. You lovingly tolerated my long hours of work, thank you.

Contents

List of Figures

List of Tables

1 Introduction

Contents

Our world relies on information and computing systems. With mobile devices, computers, servers, the Internet, and IoT[1] products, our addiction to a connected world is increasing. The adoption of IT systems has created a new playground for attackers. Data breaches and intrusions occur daily aiming at a wide range of targets going from end users to governments including large and small companies. Attackers have been so successful that in 2012 at the RSA Cyber Security Conference, the director of the FBI, Robert S. Mueller, stated [6]:

> I am convinced that there are only two types of companies: those that have been hacked and those that will be. And even they are converging into one category: companies that have been hacked and will be hacked again.

His assessment still stands nowadays. Attackers are now well organized and can be sponsored by nation states meaning they have more resources and better technical skills, resulting in even more severe damages. The WannaCry ransomware [7] hit the headlines in 2017 with total damages ranging from hundreds of millions to billions of dollars. The United States, the United Kingdom, and Australia attributed this attack to North Korea.

Successful attacks do not just result in a loss of money: they also reveal personal information. The consumer credit report agency Equifax[2] was featured on the headlines of news articles after losing the personal data of more than 145 million US citizens [8].

Firewalls, antivirus, secure programming, configuration hardening or cryptography are used to enforce the security policy and avoid security incidents[3]. However, proactive

[1] Internet of Things

[2] `https://www.equifax.com/`

[3] A security incident is defined as a single or a series of unwanted or unexpected information events that have a significant probability of threatening information security [9].

security i.e., measures taken to prevent perceived threats, are insufficient to protect information systems. Security incidents cannot be entirely prevented. Breaches will be found on information systems and attacks will be successful. The attackers only have to discover one way in while defenders have to defend the whole perimeter of the information systems.

This asymmetry brings an intense focus on incident detection, investigation, and reaction capabilities. While automated tools exist that try to detect and stop attacks reactively, their efficiency is still partial and intrusions are still occurring on information systems. Given the complexity and the evolving techniques of current and future attacks, these automated tools often fail to detect attacks and to adapt to the new threats. Humans with their contextual knowledge and intuition have a crucial part in incident detection, and we are convinced that this situation will continue. Even more, we advocate that humans have to increasingly collaborate together to understand the organization's network, do the security data triaging, investigate the problems and remediate them.

This context and the motivations are presented more deeply in the first part of this introduction, with the description of security monitoring for information systems. Then we sum up the contributions of this thesis and present the structure of this document.

1.1 Information systems and security monitoring

Reactive capabilities, detection and reaction to threats, are often concentrated within a cybersecurity operations center (SOC) dedicated to handling security events and security incidents. The first SOCs were designed around 1990 [10] and nowadays most large information systems are monitored by one. SOCs are either internal or outsourced and managed by another company.

The most prominent and time-consuming activity of a SOC is monitoring, meaning triage and analysis of large numbers of security events. A typical SOC collects from thousands to hundreds of millions of security events every day [10] with the objective of finding which of them require priority attention. Among the numerous data feeds that are available, the primary source of security events are intrusion detection systems (IDSes).

While IDSes have proven to be useful, they are well-known to raise large quantities of alerts, with more than 90% of them unrelated to relevant security issues [11, 12]. The high volume of irrelevant security events and the way they are currently handled lead to the fact that real attacks are often missed and ignored. Consequently, there may be a significant delay (up to several months) between an intrusion and its discovery [13, 14], resulting in severe damages to the company owning the targeted information system. An extreme example is the Yahoo![4] data breaches [15]. A first data breach occurred in August 2013, followed by a second in late 2014. Both breaches were discovered in July 2016. Three billion Yahoo! accounts were impacted with names, email addresses, telephone numbers, dates of birth, and other personal information of users revealed.

[4] `https://www.yahoo.com`

The high number of security events to be reviewed by security analysts also put them under pressure. They have to answer immediately and correctly to the alerts raised by IDSes. This pressure leads to poor judgments when looking at security events. Security analysts sit all day in front of a computer screen, looking through thousands of raw alerts and security events coming from IDSes and multiple sensors to ignore them or escalate them, with the anxiety of missing real alerts tied to an actual attack. As a consequence security analysts suffer from a high burnout rate [16] and the turn-over is high.

1.2 Research objectives and contributions

This thesis contributes to improving system security monitoring and helping security analysts in collaborating and discovering intrusions inside information systems. SOCs are complex organizations which require specific knowledge of their internal mechanisms in order to improve their efficiency and solve these problems. The initial step towards this objective is a better understanding of its operations, and then the expression of the root reasons for the current situation.

First of all, fully automatic systems are not the silver bullet to maintain the security of information systems. We advocate that an efficient way to handle a large number of alerts and events is to make collaboration among security operators easier by better organizing the workflow through visualization [17, 18]. Visualization enables human analysts to stay in the loop, help them interact with the data and easily understand the context. Visualization for information systems security is a growing field, yet few pieces of work are specifically targeting security analysts working in SOCs.

The main contributions of this thesis are as follows:

- the review of SOCs and their organizations resulting in the expression of their limitations, classified in two different categories: technology and process. This review gives a reliable starting point for the two next contributions of this thesis and the development of adequate solutions for the security analysts working in SOCs.
- VEGAS [2, 3], an alerts visualization and classification tool that allows first line security operators to group alerts visually based on their principal component analysis (PCA) representation. VEGAS is included in a workflow in such a way that once a set of similar alerts has been collected and diagnosed, a filter is generated that redirects similar alerts to other security analysts that are specifically in charge of this set of alerts, in effect reducing the flow of raw undiagnosed alerts.
- TheStrip [4], a new collaboration process and a design prototype to enhance the cooperation between security analysts inside SOCs in order to quickly process security events and define a better workflow that enables them to efficiently exchange feedback.

Alongside these contributions, we also review visualization tools and their techniques according to situational awareness, with our addition of collaboration between security

analysts as a use case for situational awareness. This review shows us the lack of collaboration features of current solutions, and the different strengths and weaknesses of visualization techniques.

1.3 Thesis structure

Chapter 2 presents a literature review on security operations centers, based on three criteria: missions, architecture, and human organization. This review exposes the current limitations of SOCs with respect to two aspects, process and technology.

Chapter 3 discusses applications of visualization in cybersecurity inside SOCs. It describes the different use cases and how current visualization solutions try to answer the identified limitations are described.

Chapter 4 tackles the issue of triaging large quantities of alerts with our contribution, VEGAS [2, 3], an intuitive visualization tool that allows grouping similar alerts easily and dispatching these groups of alerts among security operators for further analysis.

Chapter 5 is dedicated to TheStrip [4], our new process and tool to enhance collaboration between security analysts.

Finally, Chapter 6 summarizes our contributions and outlines perspectives.

2 Security operations centers

Contents

Security operations centers are the central point for monitoring, analyzing and acting on threats so as to prepare for, detect and respond to security incidents. SOCs are receiving large amounts of data and security analysts are collaborating to discover the intrusions and attacks inside them. Most security analysts are struggling to efficiently accomplish these objectives, and real attacks are often missed and ignored [19], lost among the data stored by SOCs. Consequently, there is often a delay measured in months between an intrusion and its discovery, resulting in harmful repercussions for the victim organization.

Another important aspect is that the pressure put on security analysts in SOCs results in poor judgments when looking at security events and in a high burnout rate [16]. Due to this fact, the period of work of a security analyst generally spans between one to three years [20].

In this chapter, we provide a brief history of SOCs to understand their evolution over time. We then explain in details the functions they provide. We then explore two essential aspects of SOCs: system architecture and human organization. We advocate that taking these two aspects into account is necessary to cover the whole perimeter of a SOC on different scales and to draw the global picture. For each aspect, we differentiate the technology from the process. Based on these aspects, we finally highlight the current limitations inside SOCs into two categories: technology and process.

2.1 A brief history of security operations centers

The origins of computer security date back to 1975 with the creation of antivirus and firewall software, and their usage in government and military organizations [21]. Indeed with the creation of computer networks, first intrusions and abuses appeared. For instance as soon as 1979, Kevin Mitnick broke into the computer network of the American corporation DEC (Digital Equipment Corporation) and copied their software [22]. In this period of time, due to the low bandwidth, the computer security of an organization was handled by a single person having skills in network technologies.

The first generation of SOCs started around 1990. At this time, computer security was no more reserved for governments and military organizations and large organizations started to design and build their SOCs [10]. Attackers of that era created worms and bots to amplify their actions, like the Happy99 computer worm for Microsoft Windows [23]. SOCs reacted to these new attacks by focusing on intrusion detection and companies began to sell IDSes (Martin Roesch created the Snort[1] IDS in 1998) and firewalls (the first commercial firewall was DEC SEAL in 1991). The concept of security information event monitoring (SIEM) was introduced at the end of this generation with aggregators and correlators [21]. We can see that in this first generation of SOCs the response to intrusions was a technical one with the development of products.

The numbers of attacks increased rapidly after 2000. States constituted national CERT (Computer Emergency Response Team) to handle computer security incidents. This is an evolution from the first CERT founded in 1988 by the CERT Coordination Center at Carnegie Mellon University. In 2000, in France, the Prime Minister set up the CERTA (now CERT-FR[2]). The United States Congress created the US-CERT[3] in 2003, with the responsibility to analyze and reduce cyber threats and vulnerabilities, disseminating cyber threat warning information and coordinating incident response activities to cyber defense, incident response, and operational integration center [24]. The same year, California state law SB 1386 [25] regulated the privacy of personal information, becoming the first US breach notification law. Along with adhesion to security and data protection standards and legal requirements, SOCs formalized their procedures and focused on early detection capabilities and prevention rather than strictly detection. This second generation was a turning point: the answer to digital threats was no longer only technical but involved processes and legal requirements.

The end of that decade was marked by the beginning of sophisticated and state-sponsored attacks, like the first publicly known cyberwar consisting in the Russian aggression on Estonia in 2007 [26] or the Stuxnet Trojan targeting Iranian SCADA systems in 2010 [27]. Organizations became aware that intrusions happen regardless of the deployed security measures and SOCs worked on improving exfiltration detections and containment capabilities [21]. This third generation of SOCs also induced the development of information sharing since no single SOC had all the data necessary to detect all threats

[1] `https://www.snort.org/`
[2] `https://www.cert.ssi.gouv.fr/`
[3] `https://www.us-cert.gov/`

and SOCs understood that they needed to collaborate to meet their objectives. SOC exchanged indicators of compromise (IOCs), such as virus signatures or hash values, to collectively improve their detection capabilities. Even if this exchange was already happening between the CERT Coordination Center and its partners for instance, this generation is marked by the increasingly active information sharing activity [21].

Nowadays the rate of attacks is still increasing, and the amount of data handled by SOCs has never been so high. The missions of the SOCs have extended over time. They now have to detect threats, prevent them and even forecast them with threat intelligence. Constant collaboration among organizations is now fundamental to improve security by maintaining up-to-date information and awareness.

2.2 Missions

We showed that SOCs have adapted continuously to face evolving threats. We argue that this evolution was following a reactive process and not a proactive one: when new threats emerged, modifications were made. We should underline that SOCs was not a central topic in the scientific community of computer security and the literature was mainly written by SOCs vendors or detection services providers. This section details the different missions that a SOC should or can provide according to the literature.

McAfee [28], Fortinet [29], Hewlett Packard Enterprise [20, 30], Splunk [31], or IBM [32] have all released white papers. Joseph Muniz, architect at CISCO, wrote a full book about operating SOCs [33]. Consulting groups like EY [34], Deloitte [35] or Tata Consultancy Services [36] also have edited guides on how to manage SOCs. The security consultant David Nathans wrote in 2015 a guide [37] on this subject. Other documents describing SOCs and best practices were written by institutes, agencies or magazines: The Information Security Journal [38], MISC [39], the SANS Institute [40], CLUSIF [41], MITRE Corporation [10], ANSSI [9] or the NIST [42]. Each of these documents gives a specific view of SOCs. For instance, CLUSIF [41] focuses on how to start a SOC from nothing, and Saâd Kadhi [39] shows tools dedicated to a SOC. All these documents agree on the fact that SOCs must enable business continuity and efficient recovery as well as prevent threats from impacting the business. SOCs also have to provide insightful risk and compliance reporting, and ensure that groups managing critical infrastructure components from a regulatory perspective are aware of potential threats to enable quick remediation of risks.

In 2013, Jacobs et al. [43] observed that there was no model to measure the effectiveness of SOCs. They proposed to evaluate a SOC through the functions it provides (named *capabilitites*) and the maturity of these functions. The capabilities cited by Jacobs et al. are divided into two groups, primary and secondary. Primary capabilities are essential to a SOC while secondary capabilities are functions offered in addition to the primary SOC capabilities. The list of capabilities offered by Jacobs et al. is technical. For instance, Jacobs et al. emphasizes log collection, retention, archival, and correlation. However, this paper lacks some capabilities like triaging or collaboration with other SOCs, and the adequacy of the architecture and the relevance of human processes are not evaluated.

A more complete list of capabilities a SOC is given in the report of Carlos Zimmerman for the MITRE Corporation [10]. This report updates the work done by West-Brow et al. [44] and aims at giving an as complete as possible list of capabilities. We separate the capabilities proposed by Zimmerman between primary and secondary, following Jacobs et al., in order to highlight the essential ones. The primary capabilities are:

- *Real-Time Analysis.* A SOC is the place to receive reports about security events and to do a real-time triage of data feeds to detect potential intrusions. This capability is described in the literature as the most significant offered by a SOC.
- *Intel and Trending.* A SOC collects and analyzes cyber intelligence. This includes the creation of new signatures, long-term analysis of event feeds, and threat assessment.
- *Incident Analysis and Response.* Security analysts working in the SOC perform an in-depth analysis of potential intrusions and provide recommendations on how to respond. This includes tradecraft analysis, countermeasures implementation and sometimes the actual response.
- *Artifact Analysis.* Gathering, storing and analyzing artifacts such as malware, network traffic or data from mobile devices are part of the capabilities of a SOC.
- *Audit and Insider Threat.* A SOCs collects data for long-term retention to enable further audits or analysis. It may be the support for insider threat analysis and investigation.
- *Scanning and Assessment.* A SOC maps its constituency networks to understand them better. The scan includes vulnerability scanning to find weaknesses in the network. Analysts working in SOCs can also perform penetration testing and red teaming to simulate attacks and test the security of the system.

Secondary capabilities are:

- *SOC Tool Life-Cycle Support.* The SOC contributes to its own IT, including ensuring the security of devices like the firewalls or proxies. This capability involves tuning the diverse sensors as well as developing original tools and signatures if needed. Staying up to date with the threat model implies to monitor the evolution of the commercial tools and the research in this domain.
- *Outreach.* A SOC can extend its core functions with product assessment, security consulting, training and awareness building, situational awareness, redistribution of adversary's tactics, techniques and procedures as well as media relations.

2.3 Architecture

White papers from SIEM or SOCs providers focus on how to implement their solutions in an information system. They emphasize the central place of their solutions in the

security of the information system. They describe what are the different components of their solutions and how to link the different data feeds like the firewalls or the IDSes to their tools. These specific frameworks do not show the global architecture of a SOC. For instance, the architecture given by HP [30] only shows how the connectors, the data-centers and the correlation tools work together, but not how they interact with the capabilities of a SOC.

Joseph Muniz et al. [33] try to regroup all the SOC technologies under a cohesive architecture shown in Figure 2.1. This architecture depicts a possible organization of components and their relationships. On top, the data sources, divided between internal data and external data, constitute the input of the SOC. Then the SOC's technologies (event collection, storage, correlation and so forth) are communicating together to efficiently use these inputs. At the bottom, alerts and actions constitute the output of the SOC. The arrows in this figure represent the relations between the diverse components. For instance, the alert dashboard uses the results from the anomaly detection. This architecture illustrates the different components of a SOC. However, it stays still close to the technology.

We believe that this technical approach to the architecture does not suffice to fully understand a SOC. The architecture of a SOC can be approached with a less technical and higher level point a view. For instance the French Cybersecurity Agency[4] (ANSSI, Agence nationale de la sécurité des systèmes d'information) has edited a document that describes all the requirements a security incident detection service must comply with [9]. In this document, and in accordance with the capabilities of a SOC described before, three distinct activities are mandatory: event management i.e. collection and storage of security events, incident management i.e. identifying, qualifying and managing security incidents, and reporting i.e. to communicate with the targeted information system stakeholders. If a SOC does not comply with the requirements, it can not be endorsed by the ANSSI for certain missions.

The agency takes into account the three activities cited previously to design the architecture of a SOC. The reference architecture published by the ANSSI is shown in Figure 2.2. This architecture is independent of software solutions or technologies.

First, the SOC is separated from the information system it monitors, and the SOC is located in a trust zone. Even if the SOC is internal, it should be separated from the monitored system to prevent compromise of the SOC from the monitored network. The events gathered from probes are collected on the monitored information system and then sent through a safe connection to a collection zone inside the SOC, as presented on the left part of Figure 2.2. This first part constitutes the event management activity.

Events are then used in the analysis zone. Using dedicated tools, security analysts group events into meaningful incidents and are able to understand the situation of the information system.

The third step is the transfer of the incidents to a zone dedicated to report to the stakeholders of the monitored systems. Operators inside the SOC must transmit their findings to them and help in providing an appropriate response with the collaboration

[4] `https://www.ssi.gouv.fr/`

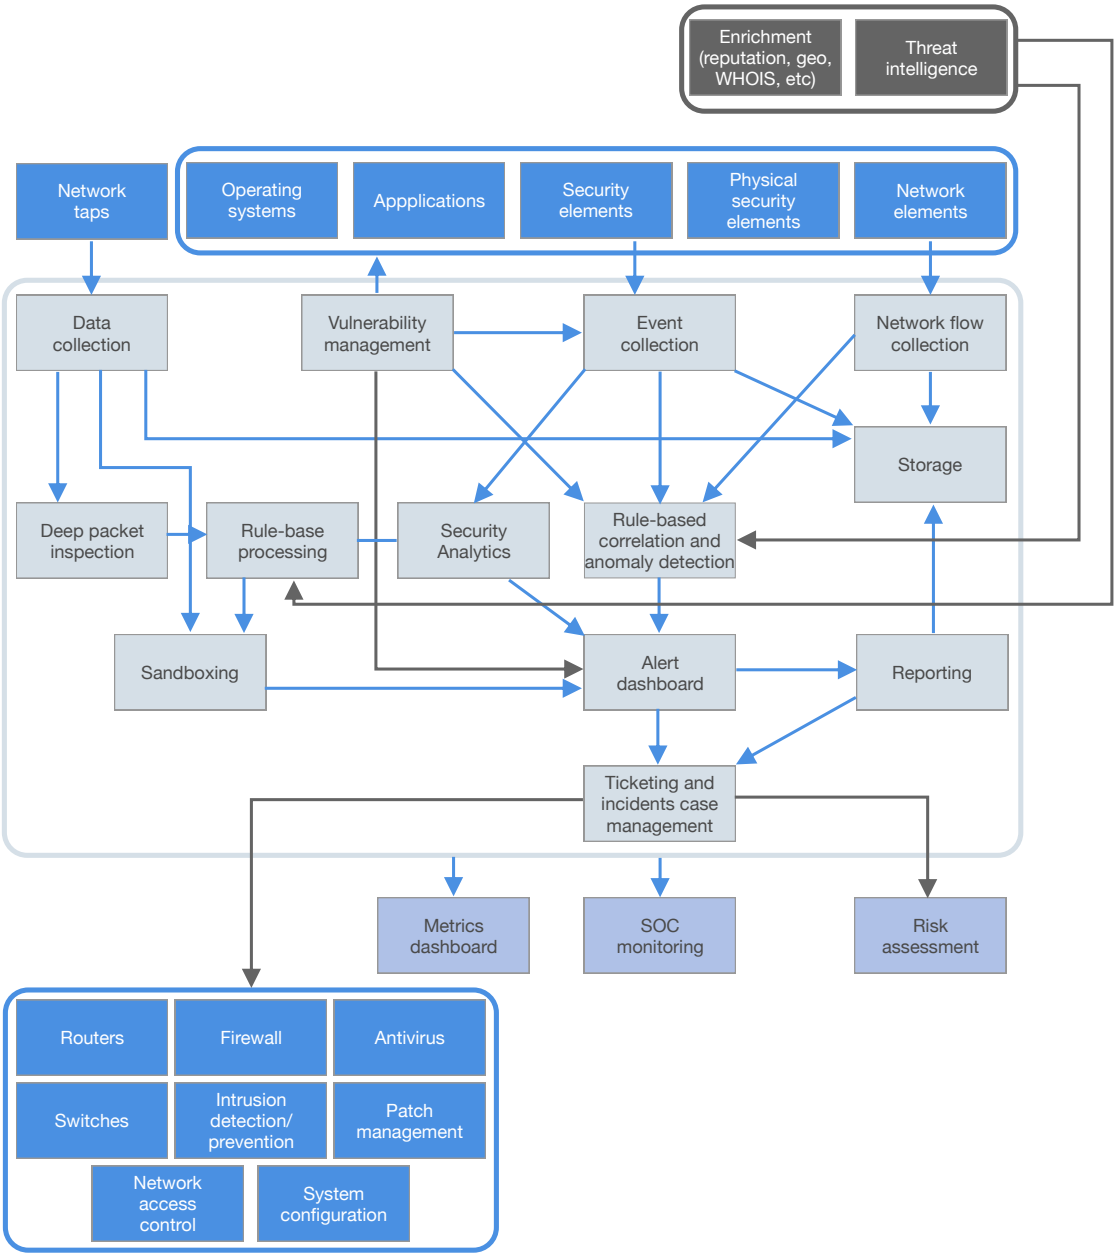

Figure 2.1: Conceptual technical architecture of a SOCaccording to Joseph Muniz et al. [33] (adapted).

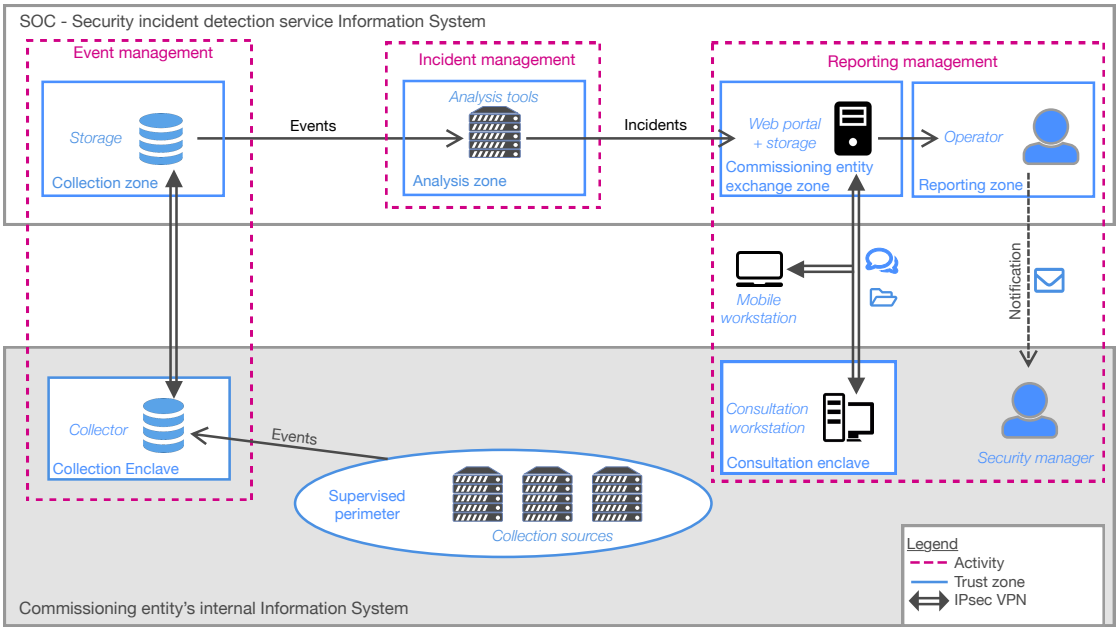

Figure 2.2: Illustrative diagram of the architecture of a security incident detection service according to the ANSSI [9], version 2.0.

of the security manager. This appears on the right part of the figure. This reference architecture enables a clear separation between the different activities. Separations of activities must be taken into account when designing the network architecture which helps in ensuring defense in depth.

The classification of the capabilities of a SOC according to the ANSSI activities is in Table 2.1. This table shows that the capabilities not linked to an ANSSI activity are those which are secondary to assist the primary activities of a SOC.

After this presentation of the global picture of a SOC, the next phase is to focus on the incident management zone. This zone is arguably the most important of a SOC, where security analysts accomplish their jobs. Security data is analyzed and response to threat created in this zone. Section 2.4 presents its organization.

2.4 Organizational model of the incident management zone

SOC members should be able to analyze large volumes of data and know when to carry out further investigations. They also must have the appropriate training to deal with the evolving nature of the job of security analyst. Therefore, having a successful SOC requires security analysts with a broad range of skills and a variety of experiences.

Capabilities	ANSSI activities
Real-Time Analysis	
Call center	Event management
Real-Time Monitoring and Triage	Incident management
Intel and Trending	
Cyber Intel Distribution Cyber Intel Creation Cyber Intel Fusion Trending Threat Assessment	Incident management
Incident Analysis and Response	
Incident Analysis Tradecraft Analysis	Incident management
Incident Response Coordination Countermeasure Implementation On-site Incident Response Remote Incident Response	Report management
Artifact Analysis	
Forensic Artifact Handling Malware and Implant Analysis Forensic Artifact Analysis	Incident management
SOC Tool Life-Cycle Support	
Sensor Tuning and Maintenance	Incident management
Border Protection Device O&M[5] SOC Infrastructure O&M Custom Signature Creation Tool Engineering and Deployment Tool Research and Development	None
Audit and Insider Threat	
Audit Data Collection and Distribution Audit Content Creation and Management Insider Threat Case Support Insider Threat Case Investigation	Incident management
Scanning and Assessment	
Network Mapping Vulnerability Scanning Vulnerability Assessment Penetration Testing	Incident management
Outreach	
Redistribution of TTPs[6]	Incident management
Product Assessment Security Consulting Training and Awareness Building Situational Awareness Media Relations	None

Table 2.1: Classification of a SOC capabilities according to the ANSSI zones

In a report dedicated to building a SOC [41] from scratch, the CLUSIF[7] provides a workflow and specify how security analysts collaborate inside the incident management zone.

The security analysts of a SOC are separated in several categories, each with different duties. Tier 1 analysts are, in number, the biggest category. They are responsible for continuously monitoring the alert queue, triaging security alerts, monitoring the health of security sensors and endpoints as well as collecting data and context information necessary to initiate Tier 2 tasks. The tasks performed by Tier 2 analysts consist in fulfilling deep-dive incident analysis by correlating data from diverse sources, determining if a critical system or dataset has been compromised, advising on remediation and reporting periodically to the SOC manager.

According to the CLUSIF Tier 1 and Tier 2 analysts cooperate in two processes shown in Figure 2.3:

- *The detection process.* This process is shown in green on top of Figure 2.3. The Tier 1 analyst is responsible for the quick triage of the security alerts. If the security alert is associated with a written procedure, the Tier 1 analyst follows it; otherwise, he or she calls a Tier 2 analyst, who will execute the qualification process.

- *The qualification process.* This process is shown in blue on the Figure 2.3. Tier 2 analysts study the given set of security alerts and try to better understand them. Depending on the result of their analysis, Tier 2 analysts follow a procedure related to the incident if it exists or design new procedures.

Dealing with education and experience, Tier 1 analysts hold a Bachelor and a few years of experience while Tier 2 analysts usually hold a Master degree and several years of experience.

Tier 1 and Tier 2 analysts are under the supervision of the SOC manager who is responsible for prioritizing work and organizing resources to ensure that the SOC is running efficiently. He or she manages resources (personnel, budget, shift scheduling and technology strategy) to meet the SLAs[8]. He or she communicates with management and other authorities, serves as the primary contact for business-critical incidents, and provides overall direction for the SOC and input to the global security strategy. The SOC manager also handles internal and external communications.

When the size of a SOC increases, other positions exist to fulfill specific capabilities [10]. Regarding the security analysts, there are sometimes Tier 3 analysts who take over Tier 2 analysts when specific analysis techniques are needed. Some SOCs may have trending analysts, scanning analysts, or product assessment analysts, having specific skills in these fields. Moreover, the SOC tool life-cycle support capabilities can be accomplished by dedicated engineers. The MITRE Corporation proposes a representation for a SOC with these positions (see Figure 2.4). Systems administrators and owners collaborate with

[7] `https://clusif.fr/`

[8] Service-Level Agreement. A SLA is a commitment between two or more parties, usually a service provider and a client. Particular aspects of the service (quality, availability, responsibilities) are agreed between the parties.

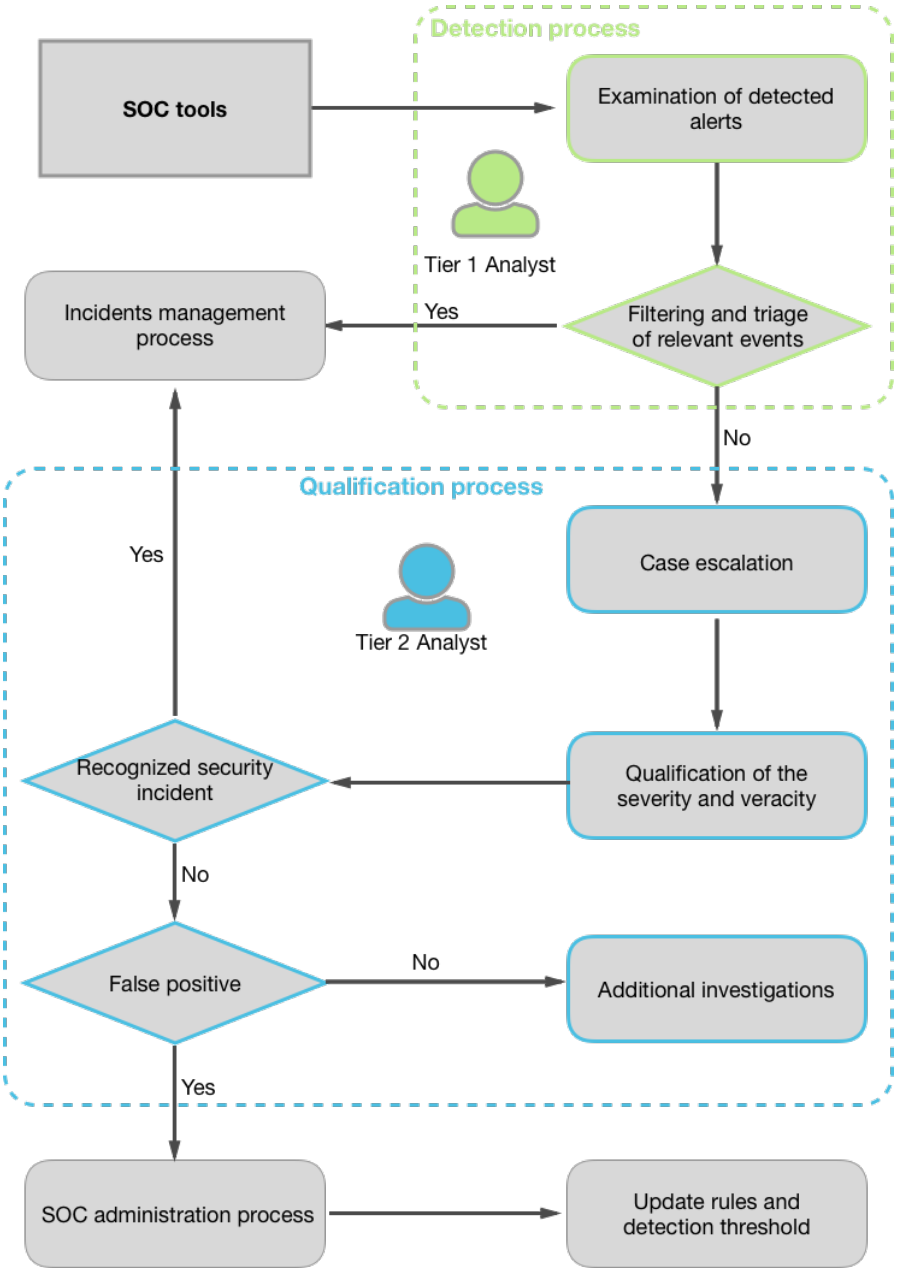

Figure 2.3: Organizational model of the SOC incident management zone according to the CLUSIF.

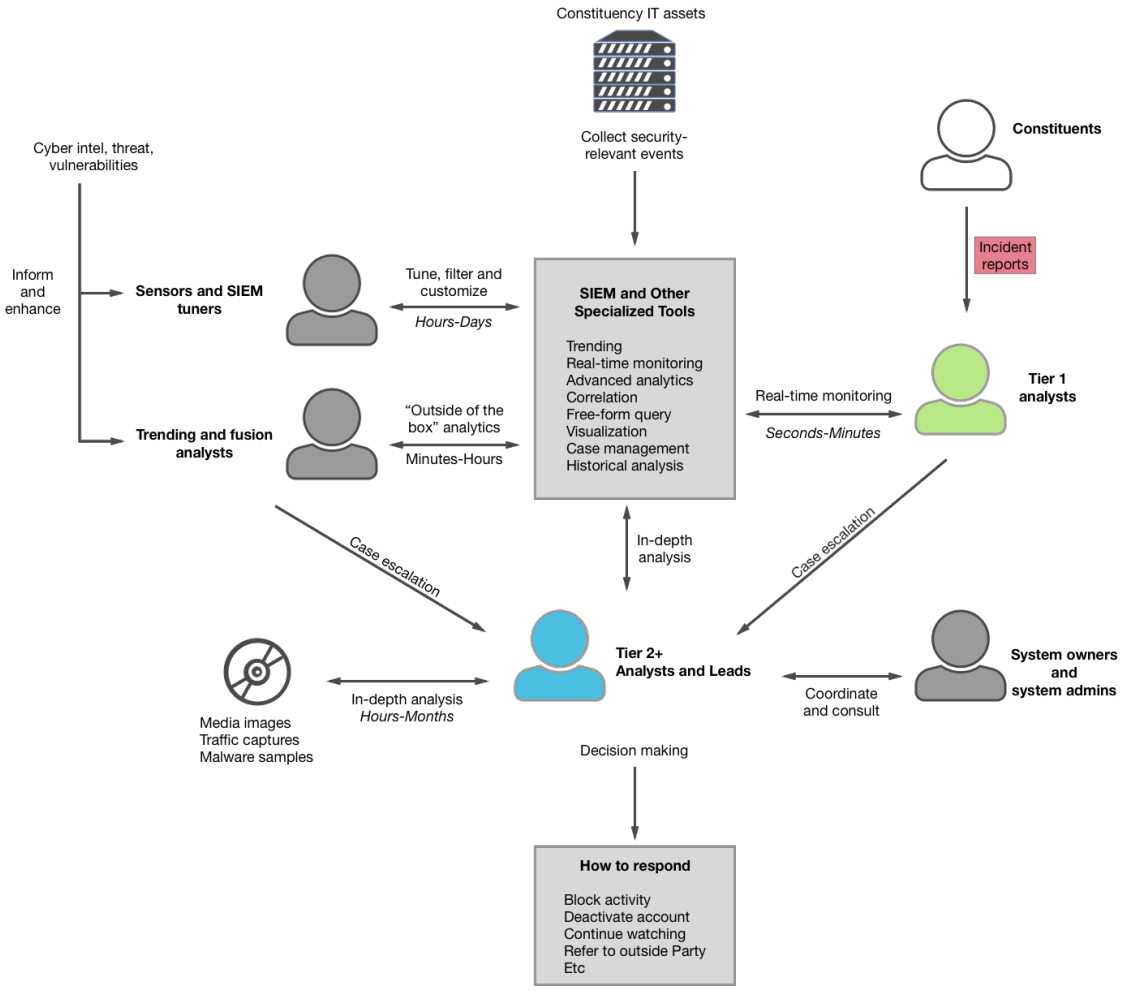

Figure 2.4: Organizational model of a SOC according to the MITRE.

Tier 2 analysts to find an appropriate response. Another source of case escalation, in addition to the Tier 1 analysts, are the security analysts doing sensors tuning and "out of the box" investigation. This representation details the timeframe given to each analyst. Tier 1 analysts only get a few seconds or minutes to accomplish their task, whereas Tier 2 analysts have more time, up to several months, to perform in-depth investigations.

2.5 Limitations of Security Operations Centers

A report of MWR InfoSecurity [45] focusing on the reasons why SOCs are failing stated that around 5 percent of the 100 attacks they simulate to evaluate SOCs are actually detected. So SOCs often missed attacks, the sign of these attacks being hidden among all the data stored. In its threats report from December 2016 [46], McAfee states that 93 percent of SOC managers are overwhelmed by alerts and unable to triage all potential threats. This alerts' overload leads to organizations being only able to investigate 25 percent of their security alerts. We assert that this is due to intrinsic limitations. In this section, we present the current limitations of SOCs, divided into two categories: technology challenges, and process challenges.

2.5.1 Technology challenges

SOCs face technology challenges related to data, the progression of threat, and knowledge of the rhythm of the network.

Lots of unlinked data and data sources

For Tier 1 analysts, the main source of data is the security alerts from the IDSes. IDSes are categorized into two complementary types, knowledge-based and behavior-based [47]. Knowledge-based IDSes detect intrusions through established patterns of well-known attacks such as signatures. When a known pattern is detected, an alert is raised. This approach requires a complete and up-to-date knowledge base of attacks patterns, otherwise attacks will not be detected. This means that zero-day attacks do not raise an alert. This type of IDSes generates many false positives, i.e. alerts that are not related to an effective attack due to erroneous detection rules or incorrect tuning of the IDS.

Behavior-based IDSes detect a deviation from the normal or expected behavior of the system or the users. This means that the correct behavior must be learned or defined before. The IDSes then compare the reference model with the current activity and raise an alert when it observes a deviation. If the training phase is too specific, normal behavior marginally different from what has been learned will raise alerts which are false positives. On the contrary, if the model learned is too broad, some attacks will not be captured. Moreover, behavior of users evolves over time. Therefore, the model must be kept up-to-date which is a challenging task.

Even with a single data source, Tier 1 analysts have to face a huge volume of security events and only have seconds or minutes to accomplish their task in order to cope with

the flow of alerts. This challenge also exists for Tier 2 analysts. The amount of data given to them is prodigious, in the order of millions of security events to explore. Even when Tier 1 analysts help them so that they only look at interesting events, there is a lot of data to explore in order to understand threats.

Beyond IDSes, the data sources are various: antivirus, system events, network traffic, various logs. This diversity is a supplementary challenge to Tier 2 analysts, expertise in each of these data sources being required. In addition to the multiple sources of data, the sources used by Tier 2 analysts are not necessarily linked between them. Each data source e.g., IDSes alerts, netflow traces, packet data, web server logs, etc., has its own structure and semantic. Due to this fact, correlation and pivoting in the data is a hard task. Security analysts often have to understand which events are mirrored between the sets of data, how a value is translated from a set to another, etc. An example is the network address translation performed automatically when packets go through some routing devices that make it hard to link data packets and events before and after routing devices.

Progression of threat

Tier 2 analysts evaluate the level of threat of a given alert and investigate the events given the elements of context of the information system. It is particularly important to evaluate if the event is isolated or if it is a part of a bigger scheme. It is also required to discover the technicity of the attack,i.e., it comes from a script kiddy or if the information system if targeted by an APT[9]. The knowledge of the current context, threats, and incidents currently happening help the Tier 2 analysts to make a decision.

Rhythm of networks

Security analysts responsible for the security monitoring of an information system gain a specific knowledge of the rhythm of their network. They know the particular events happening periodically and what will follow such events. Security analysts frequently do not have the right to modify the sensors of the network, so they learn to deal with such events. In the context of visual security tools, Daniel Best et al [48] named this situation "Cadence of Network". The understanding of such cadence and the typical amount of errors in the system is currently insufficiently exploited. We should mention that it is a part of the collection strategy required by the ANSSI [9].

2.5.2 Process problems

A SOC is a place of secrecy due to the sensitivity of the information handled and the high workload, making it challenging for researchers to easily investigate this subject. Two anthropological studies of SOCs were performed by Sundaramurthy et al. in 2014 [49] and

[9]Advanced Persistent Threat. An advanced persistent threat is a broad term used to describe a set of continuous computer hacking processes, often orchestrated by a intruder or team of intruders targeting a specific entity.

2015 [16] to observe the security analyst burnout happening in SOCs. The collaboration inside a security team is also addressed by Rajivan et al. [50] who employed a hybrid methodology, a mix of field observations and simulations, to focus on the team situational awareness, meaning the global comprehension at a team level of the security events happening in the network. Even if the different teams observed in their study do not have the same objective as a SOC, some observations are relevant to our subject. Their findings, related to process problems, are divided into three categories.

Repetition of the same task

Tier 1 analysts perform repetitive tasks and are spending most of their time following known procedures. When the same type of events keeps coming, they have to repeat the same procedure over and over again.

This aspect is also true for Tier 2 analysts. Because Tier 1 analysts keep sending the same type of events, Tier 2 analysts have to deal with these same events. Even if they do not need to perform an analysis again, the consequence is a loss of time and a diminished perception of the work accomplished by Tier 1 analysts.

Lack of feedback

Once their decision is made, Tier 1 analysts lose track of their actions. They do not have the result of the analysis performed by Tier 2 analysts and therefore will not know if they performed correctly. Furthermore, if a procedure seems inefficient from their point of view, they often do not have the permission to change it. Therefore, Tier 1 analysts feel that they are not correctly empowered by the management, resulting in a decreasing motivation at work over time. Sundaramurthy et al. [16] pointed out that security analysts feel enthusiastic when they see the impact of their effort and perform better.

Lack of creativity

In addition to the repetition of the same task, the creativity of Tier 1 analysts is severely constrained. Creativity refers to the ability of analysts to handle an operational scenario that differs significantly from those they have encountered so far [16]. Tier 1 analysts simply execute the written procedure and stay with what they know; they are not empowered to deviate from the norm if the situation requires it. This results in the failure to react appropriately to a novel operational scenario.

In organizations where procedures are strict and rigid, security analysts burn out more quickly. The lack of rotations between analysts and assignments is also a decreasing factor for the creativity of security analysts.

The limitations we exposed in this chapter are impacting the efficiency of SOCs. Their division in technology and process is the first step to address them. Thanks to the better understanding of the cause of the problems of SOCs, we can propose appropriate solutions to both aspects.

2.6 Conclusion

SOCs are a central place for the security of information systems. Despite the evolution of their missions and tools, they suffer for process and technology limitations. Tier 1 analysts are particularly affected by these limitations, resulting in high burnout rate and strenuous working conditions. The situation of these security analysts is so precarious that organizations are struggling to find candidates for these positions and to keep their employees [51].

We strongly believe that visualization can be a response to the challenges met in cybersecurity, as Daniel Best et al. explained in their paper [48]. The current state of security visualization inside SOCs is the topic of Chapter 3 and the conclusions are then used to tackle the problem of overwhelming data given to Tier 1 analysts. Our proposition for the triaging of security alerts is detailed in Chapter 4.

Regarding the process limitations, we believe that collaboration is yet insufficient in SOCs. Rajivan et al. [50] show that collaboration and information sharing, when done correctly, have a positive effect for the triage and the investigation of complex alerts. For known and simple security events, there is no gain to put a team effort. Sundaramurthy et al. [52] highlights the necessity to empower Tier 1 analysts to reduce the process limitations. Chapter 5 presents our proposition for a better process in SOCs, taking into consideration these insights. We also address the technology problems related to the progression of threats and the rhythm of the networks with visualization by designing a prototype corresponding to this new process.

3 Security visualization inside security operations centers

Contents

Analyzing, processing and communicating about an ever-growing amount of data are daily tasks for the analysts working in a SOC. In cybersecurity, data is primarily stored in textual form called logs. Logs come from various sources such as IDSes alerts, servers, clients, applications, Processing such a quantity of text is a difficult task for the human brain. On the other hand, the human visual system is a powerful pattern seeker and processes images and pictures with more ease. Humans acquire more information through vision than through all other senses combined, and visual displays have the highest transmission capacity from the computer to the human [53]. In the context of cybersecurity, analysts show a preference for visualization as it helps them acquire a higher number of accurate insights [54].

The first part of this chapter is dedicated to the purpose of visualization for information systems security. Current visualizations suitable for tasks carried out by SOCs can be separated into different categories (monitoring, inspection, exploration, forecasting, communication and we add collaboration) and stages (perception, comprehension, and perception) according to situational awarenesss. We then review the state of the art of security visualization according to these categories, these stages, the type of analysis performed and the involved analysts in the remainder of this chapter, The scope of this thesis being the security monitoring and the work of analysts in SOCs, other areas of security visualization like malware or binary code visualization [55] are not discussed. The lessons taken from this study of visualization for security are used for our propositions called VEGAS and TheStrip, detailled in Chapter 4 and Chapter 5.

3.1 Situational awareness and purpose of visualization for security

Visualization is the process to generate visual representations of data. The concept of using visualization to understand data has been around for a long time. Some of the first representations were related to geographic information. One of the earliest known map, found on the walls of the Lascaux caves, dates back to 14 500 BC. Statistical graphics were conceived in the 17th century [56], to show quantity, time-series, scatterplots, and multivariate data. One of the most cited examples of statistical graphics and described as the possible "best statistical graphic ever draw" by Edward R. Tufte is the combination of data map and time-series drawn by Charles Joseph Minard in 1869 in Figure 3.1. It portrays the losses suffered by Napoléon during its Russian campaign. The thick tan flow-line shows the size of the army, its direction (gold when going into Russia, black for the retreat) as well as its location during the campaign, and the visualization links this sets of data to temperature and time for a more in-depth understanding of the event. Minard illustrates the explanatory power of visualization with this effective multivariate graphic.

Computers and technology enable the process of large amounts of data and empower the development of data visualization. They give us new ways to explore, interact with and communicate about large security datasets. The interest in visualization for cybersecurity is increasing with more research paper being published. As Lane Harrison pointed out [57] the need for better visualization tools is now widely recognized and supported.

Visualization relevant to SOCs' missions can be reviewed from the scope of situational awareness. Situational awarenesss is presented in the next section and other possible classifications are then discussed.

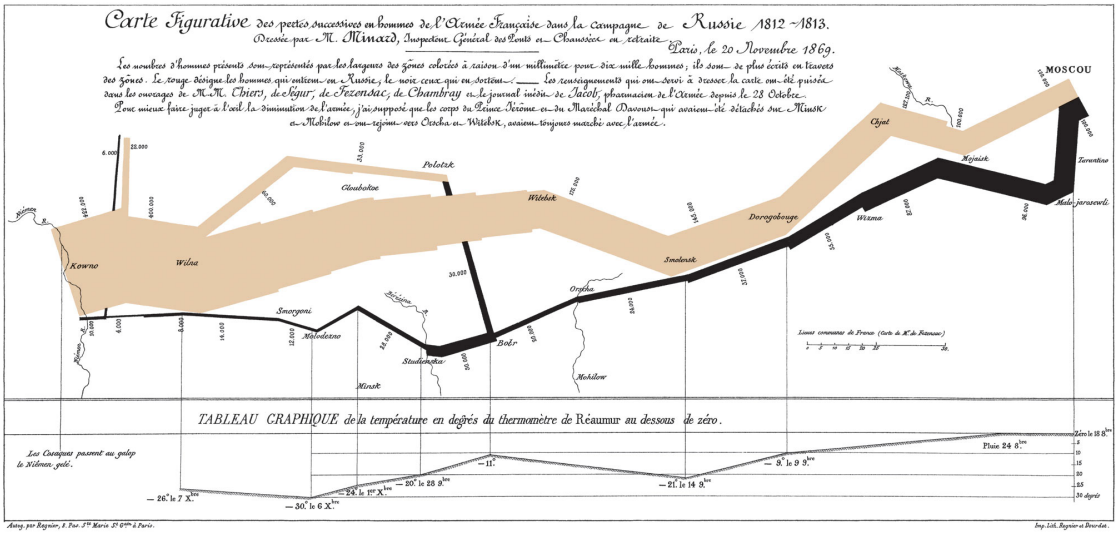

Figure 3.1: Map of the successive losses in men of the French Army in the Russian campaign 1812–1813 by M. Minard.

3.1.1 Situational awareness

Situational awareness is traditionally defined following the seminal work of Mica Endlsley who describes situational awareness as "the perception of the elements in the environment within a volume of time and space, the comprehension of their meaning, and the projection of their status in the near future" [58]. In other words, situational awareness consists in being aware of one's surroundings and understanding what information is relevant, not necessarily knowing everything but just what needed for the time being. Historically, situational awareness comes from the discipline of aviation, and so is a concept of the physical world. However, the definition fits the requirement needed by a SOC: prepare for, detect, and respond to security incidents. As noted in [59], situational awareness is now acknowledged as being provided by SOCs.

Situational awareness is made of three distinct stages of increasing complexity:

- *Perception* of data and the elements of the environment. For instance, in a SOC it means knowing the overall network activity, current IDSes alerts, etc.

- *Comprehension* of the meaning and significance of the situation. This stage involves an analysis of the elements perceived in the previous stage, their combination and integration to understand their meaning and how the situation impacts the system's objectives.

- *Projection* of future states and events. This refers to the capability to project and anticipate the future actions of the elements in the environment. Thanks to the knowledge gained from the previous stages, future states can be extrapolated. In a

SOC, the understanding of the modus operandi of an attacker can lead the analyst to predict likely future attacks.

These three stages denote a progressively increasing awareness level, starting from basic perception of important data to interpretation and combination of data into knowledge and finally to the ability to predict future events and their implications.

We can see that the three stages refer to different types of tasks accomplished inside a SOC. The perception stage is linked to the task of real-time analysis done by Tier 1 analysts, whereas the comprehension stage is the responsibility of Tier-2 analysts. The last stage, projection, is performed by specific analysts in mature SOCs and is related to threat intelligence.

Visualization is a response to provide situational awareness. In 2005, Anita D'Amico and Michael Kocka proposed a diagram showing the relationship between the three stages of situational awareness and the categories of uses of visualization in a security context [60]. They found five major uses for visual data presentation: monitoring, inspecting, exploring, forecasting, and communicating. Monitoring is the understanding of the current state of applications and systems, which is continually changing. It is part of the perception stage of situational awareness. Inspection is the search for specific details to find answers to the situations perceived during monitoring. By contrast, exploration is the examination of data without any hypothesis or questions. The objective of exploration is to generate questions or find data of interest. Both inspection and exploration start at the end of the perception stage and are part of the comprehension stage. forecasting consists in predicting some future events or states based on current data. It is related to the projection stage of situational awareness. Finally, visualization tools are used as a mean to communicate with other people. Depending on the particular goal of communication, this use is relevant on all stages for situational awareness.

Extending D'Amico and Kocka's work, we propose a sixth use to visualization: collaboration. Since security analysts in SOCs are not working on their own, they need to cooperate to achieve SOCs missions efficiently. Like communication, collaboration is relevant on all stages for situational awareness. As a consequence, we present a modified version of the diagram of Anita D'Amico and Michael Kocka in [60]. We also add the type of analysts accomplishing the work in the diagram. Figure 3.2 depicts the relations between the stages of situational awareness and the uses of visualization, the types of analysis performed, and the analysts using those types of visualization. This diagram will be our reading guide to review the state of the art of visualization related to SOCs. By keeping this diagram in mind, we explain why some applications of visualization are used in the wrong way and are not appropriate or present limitation regarding their users and objectives.

3.1.2 Other classifications

Other classifications for visualization solutions have been proposed in the literature. Raffael Marty [61] and then Christopher Humphries [62] divided the visualization for security events in three categories according to their objectives. The first category is

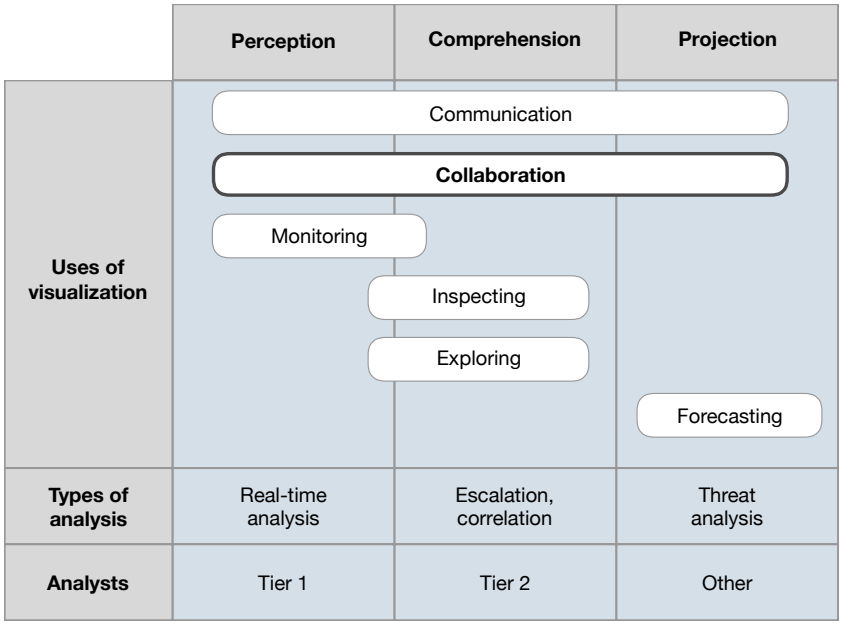

Figure 3.2: Relationship between the stages of situational awareness and the uses of visualization, the types of analysis performed, the analysts using those types of visualization. Modified version based on [60], with our addition of collaboration and the analysts.

real-time monitoring, having the objective to understand the current state of applications and systems. The second category is the analysis.This category is a merging of inspecting and exploring uses. The third category encompasses the reporting, communicating and displaying data. There are some strong similarities between this classification and ours. Forecasting is missing, and communicating is named reporting. We can notice that according to Humphries, collaboration is just related to reporting. We advocate that collaboration, like communication, is not only a part of reporting but encompasses all three stages of situational awareness.

Another taxonomy based on use cases have been proposed by Hadi Shiravi et al. [63]. They take a data-driven approach and classify network security visualization into five use-case classes: host/server monitoring, internal/external monitoring, port activity, attack patterns and routing behavior. Fabian Fischer [64] enriches this classification with two use cases, malware behavior and attack attribution, which were not in the focus of Shiravi's work.

Fischer also emphasizes that the most important data sources found are network traces, security events, user/asset context, network events, host events, application logs, and malicious data. He also classifies visualization solutions according to their visualization types using the categorization proposed in [65].

We believe that these other classifications are useful but do not prevent visualization tools from falling short of expectations and being irrelevant because the visualization is misused.By associating the classification according to the purpose of visualization and the objectives of the analysts, the review shows the usefulness and limitations of the current solutions. The remainder of this chapter investigates the six different uses of visualization we propose. When the proposed solutions are associated with several stages of situational awareness, we classify them in the most relevant stage. A greater focus is done on monitoring because this task is accomplished by Tier 1 analysts who are the main interest of our study.

3.2 Monitoring

Visualizations that target monitoring give an understanding of the current state of systems and applications, by showing events of interest. Real-time aspects are essential since Tier 1 analysts have a very short amount of time to understand what is happening and to deal with the data they receive. They need to comprehend the status of the system at a glance and to perceive state changes; therefore, visualizations are usually already configured to help analysts detect unusual patterns. Tools dedicated to monitoring are generally made of dashboards, as seen in the examples cited in this section. Stephen Few, the author of the reference guide about dashboards [66], gives the following working definition:

> A dashboard is a visual display of the most important information needed to achieve one or more objectives that has been consolidated in a single computer screen so it can be monitored at a glance.

Interaction within a dashboard is intentionally limited so analysts can react as fast as possible. A dashboard has the capability to see trends and changes over time that should attract the attention of the analysts.

IDSes alerts are the prominent source of data for monitoring and the security visualization community has proposed several solutions to represent these alerts [59]. Monitoring is intimately linked to the time, so time has an essential place in these representations. In this section, we describe the most frequently used visualization techniques.

3.2.1 Scatterplots

One of the most popular visualization techniques for monitoring is the scatterplot. A scatterplot is a diagram using Cartesian coordinates to display values for typically two variables, and is used to see how the two variables are related. Scatterplots are great for monitoring because they can quickly show relations on a two-dimensional plane and Tier 1 analysts may see trends or detect clusters. However, the difficulty lies in the choice of the variables to represent.

SnortView [67] uses a scatterplot to manage the flow of alerts created by the Snort IDS by showing the relation between the time and the source IP addresses. Alerts are displayed using the time each alert was raised on the abscissa and the source IP on the ordinate. Different colors and icons are used to represent alert, according to their classification (e.g., attacks on an email server), as displayed in Figure 3.3. A destination matrix is displayed on the right side, a red circle representing communication between a particular source and a particular destination. By using SnortView, analysts rely on familiar visual patterns and watch for outliers indicating a change in events and potential attacks. However, SnortView is limited to displaying a maximum of 40 different alerts that happened over four hours. This is insufficient when the volume of alerts is significant.

The use of scatterplots to correlate two dimensions is also adopted in IDS Rainstorm [68]. Alerts are displayed on a scatterplot using time on the abscissa and destination IP addresses on the ordinate. This representation allows identifying the main targets and threats on the network. However, the classification of the alert is not taken into account. Furthermore, due to the time-sliding window, alerts can be missed and like SnortView the scalability of this solution is limited. The same technique is used in NVisionIP [69] and in IP Matrix [70]. They propose a scatterplot to represent communication between hosts: source and destination IP addresses being defined as the axes of the scatterplot, so that particular clusters of communication can be discovered.

3.2.2 Link graphs

Link graphs are efficient to represent communications and interaction inside the information system. They can quickly show patterns of communications, which are useful for Tier 1 analysts.

NIMBLE [71] employs a design based on node-link techniques. This graph-based visualization represents hosts as cards and IDSes alerts as links between the cards as

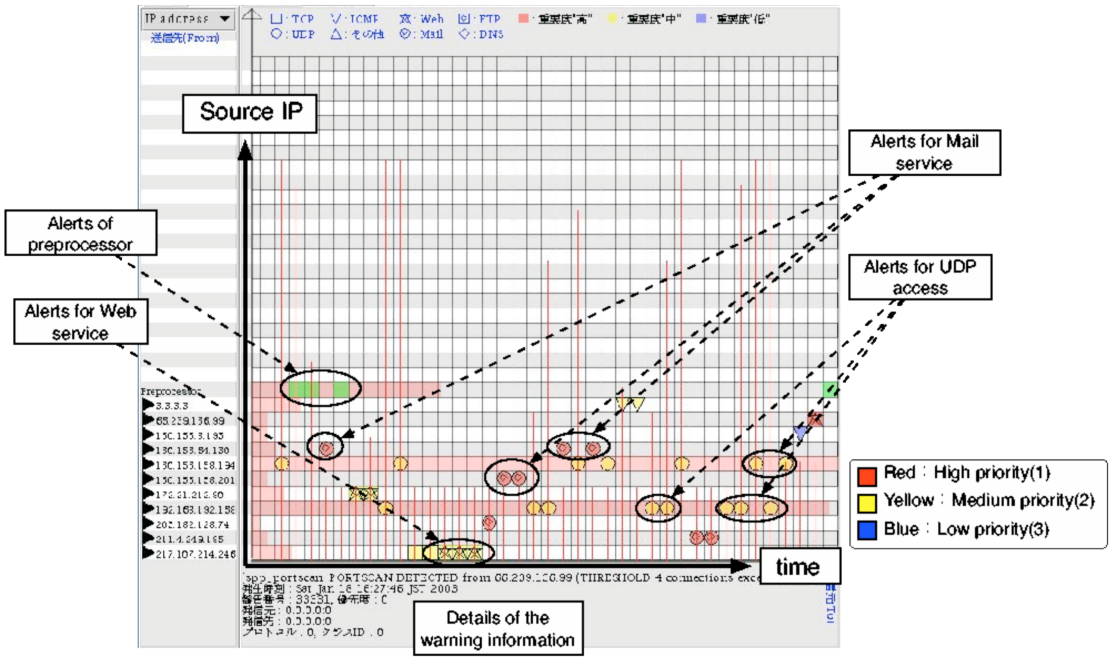

Figure 3.3: Snapshot of SnortView [67]

illustrated in Figure 3.4. Links have descriptions for the type of alert and NIMBLE suggests incident categorizations or "explanations" to help analysts.

Instead of a classic node-link graph, Avisa [72] uses a radial display to visualize the relation between alert types and hosts. The result is similar to a chord diagram. Alerts triggered by an IDS are drawn as arcs starting from the alert type panel on the top left of the ring and ending at the host affiliated with the alert on the other side of the ring. This representation enables analysts to detect strange patterns of alert such as multiple failed login attempts or scanning. A similar design is proposed in [73].

VisAlert [74] proposes another radial display for monitoring and visual correlation of network alerts. The visualization (Figure 3.5) is explicitly designed to understand the nature, time and location of events. The developed system displays the local network topology graph (the "Where") in the center with the various alert types on a surrounding ring (the "What"). The ring's width represents the time (the "When") and moves outwards as it ages. An arc is drawn from a specific attack type on the outer ring to a particular host on the topology graph to represent a triggered alarm. The objective is that changes in visual patterns lead Tier 1 analysts to detect signs of potential anomalies.

A limitation of link graphs is the size of the graph, meaning the number of nodes and edges. If there are too many elements to show, the visualization will become unreadable or confusing. Moreover adding additional dimensions, like the type of alerts or its severity, is hard with this technique. This is detrimental to Tier 1 analysts in their triage.

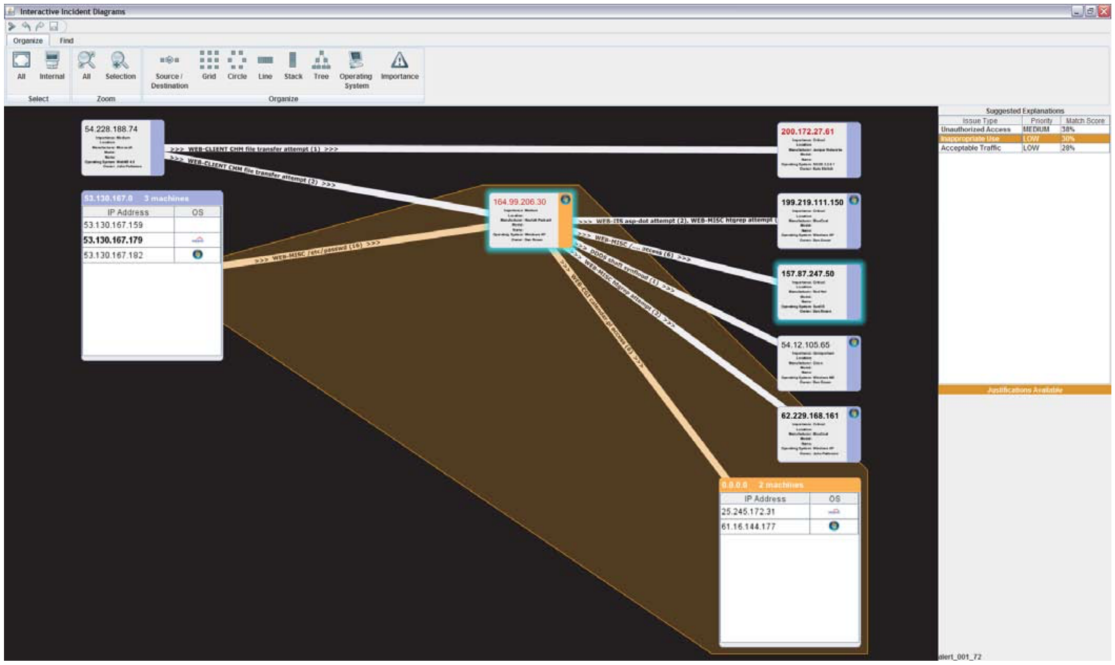

Figure 3.4: NIMBLE interface [71]

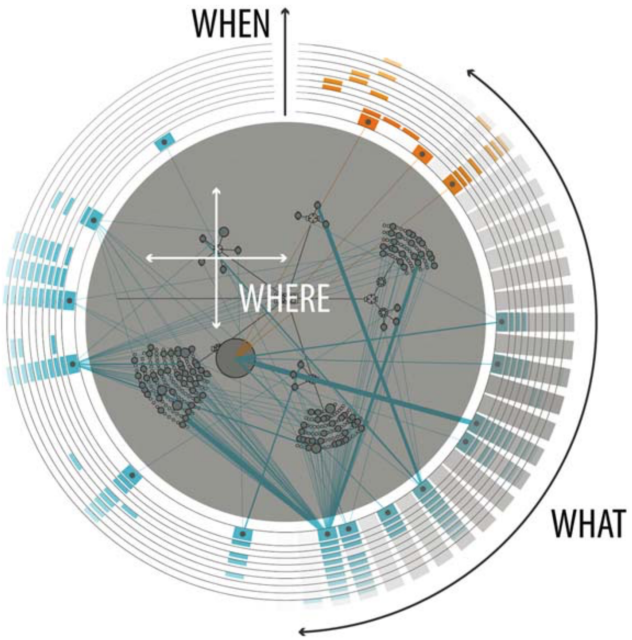

Figure 3.5: VisAlert [74]

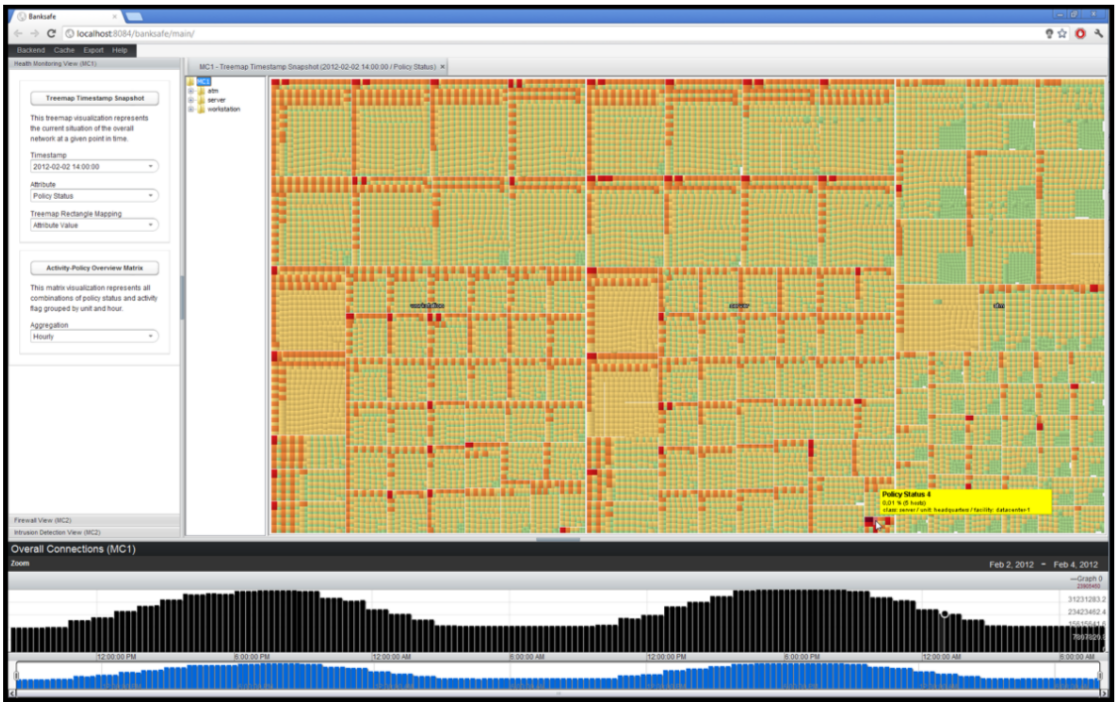

Figure 3.6: Treemap visualization for BANKSAFE [77]

3.2.3 Treemaps

Treemaps are a method to visualize hierarchical structure [75] using nested figures. In our context, treemaps are relevant to represent computer network, by dividing them into subnets. Due to its space-filling properties, treemaps are useful to display large networks and can show more hosts than a link graph; however, the links between components are lost. This is an essential property for Tier 1 analysts monitoring a large information system. Usually, treemaps are drawn with rectangle areas, to have a space-filling layout. There has been a tentative to use circular treemaps [76] with specific glyphs to better support comparative tasks.

A typical treemap example is BANKSAFE [77], shown in Figure 3.6. The main visualization of this dashboard provides an overview of health based on the IDSes alerts received for millions of computers. The hierarchy of the treemap is done according to organizational levels of the networks, with the sizes of the rectangles indicating the number of underlying hosts. The color of the rectangles encodes the health of the hosts based on IDSes alerts. Green indicates a healthy host, whereas red means that the host is attacked. This representation is well-adapted to instantly view the globals status of a network.

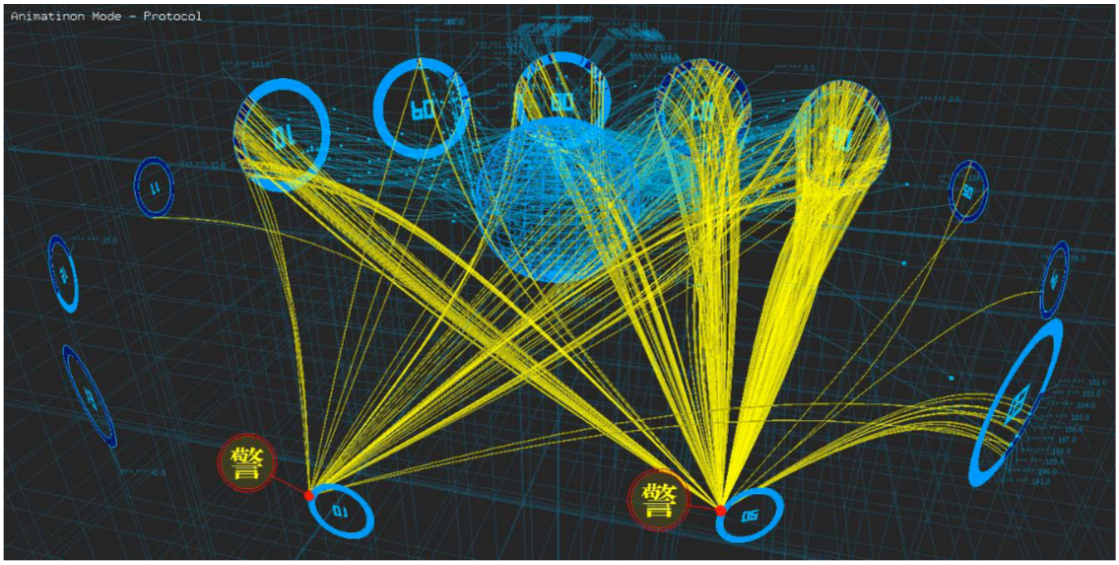

Figure 3.7: Overview of DDoS attack by Anonymous in DAEDALUS [79]

3.2.4 Three-dimensional techniques

Three-dimensional views have been proposed to improve the scalability of visualization tools. Using three dimensions extends the limits of screen space and enables the display of more dimensions to the analysts. Nevertheless, three-dimensional visualization suffers from several issues: occlusion, meaning that an information may be missed because it is hidden, and perspective, where a pattern can be seen from only a particular point of view.

Toa [78] simultaneously maps source IP address, destination IP address and destination port within a cube to allow analysts to directly see correlations between all three dimensions at once. This visualization is particularly efficient to detect port and network scans. In DAEDALUS (Direct Alert Environment for Darknet And Livenet Unified Security) [79], three-dimensional views have the objective to monitor the darknet on a large-scale. DAEDALUS has two components: a central sphere and several rings around it. The sphere represents the Internet, and the rings represent the organizations livenets and darknets that are monitored. In between the sphere and rings, alarms triggered are made visible with lines. Figure 3.7 shows DAEDALUS during a DDoS attack by the hacktivist group Anonymous.

3.2.5 Interaction

Interaction enables Tier 1 analysts to explore the data instead of just looking at static displays, it brings back human in the loop. Regarding the tasks to be performed, interaction is limited for Tier 1 analysts and is usually designed to fit the need of quick triaging. For some visualization techniques, interaction is a necessity. This is the case

for three-dimensional techniques, where interaction is used to navigate through the data. In two-dimensional techniques for monitoring, interaction is mostly limited to simple triaging and sorting.

For instance, Curtis et al. [80] present a tool to assist in the rapid browsing of alerts that displays them in a list with a specific color encoding for quick visual analysis. It provides sorting functions to be applied to variables, and specific IP can be tagged as malicious for later processing. A graph displaying the alerts frequency is also available.

3.2.6 A priori processing

Additional computations can be done a priori to assist the analysts. Instead of having representations of raw data, operations are performed to show a more interesting part of the data or a more meaningful view or a variable computed from the raw data.

TVi [81] applies PCA-based computation to detect anomalies in network flows based on the variation of the entropy. A quick change in the entropy means a different traffic on the network and can indicate a network problem such as an intrusion or a virus. TVi user interface displays the timeline of the entropy values for every subnetwork. The user can select which features (e.g., TCP traffic, DROP packet) to show in the timeline graph and select the dimensionality of the vectors of the reduced base used in the anomaly detection step. The objective is to detect outliers thanks to the entropy computations.

Other representations try to have predefined patterns for a given type of attacks. Ngoc Anh Huynh et al. propose a dashboard to detect malware that leaves periodic traces in network traffic [82]. Indeed some malware, like Zeus[1], are acting at recurring time intervals and leave specific traces in the network traffic. The dashboard shown in Figure 3.8 has been proposed to quickly detect this type of malware. The left panel (panel A) shows the alerts returned by the periodicity detector, based on Fourier transforms. A circular graph is used to visualize the timeseries associated with the alerts to allow quick temporal comparison. Panel B shows the time series associated with a particular alert and enables to zoom a given timeframe with filtering. The panel on the right (panel C) shows the textual content of a particular netflow. This dashboard allows quick recognition of periodic malware; however, netflows still needs to be manually analyzed which can be a daunting task in big networks.

3.3 Inspecting

Analysts need to inspect the data in more details when suspicious activities have been found. They have hypotheses about the causes of such activities and visualization is used to verify them. Visualization tools to support Tier 2 analysts in the inspection task are the most numerous in the literature and we believe that they provide an appropriate response regarding this task.

[1] `https://github.com/Visgean/Zeus`

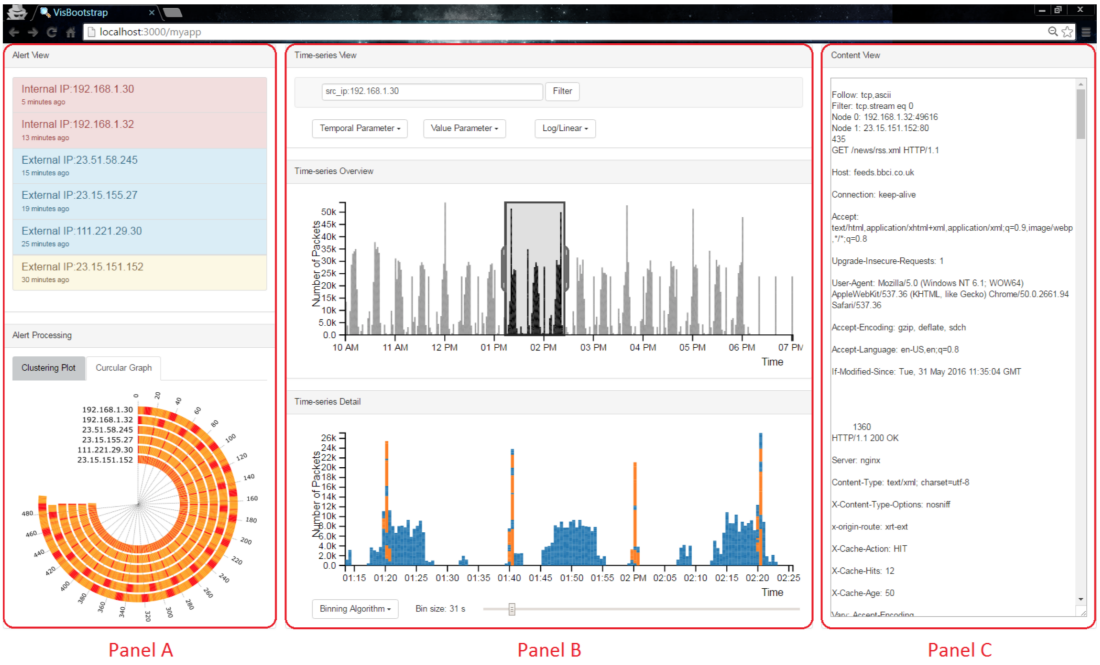

Figure 3.8: Visualization Dashboard of [82]

Visualizations dedicated to inspection are highly interactive to enhance the process. CORGI [83] is designed for in-depth analysis of multiple data sources and is a representative example of this category of tools. CORGI helps experts collect points of interest in logs and pivot around multiple sources of data. An overview of CORGI is shown in Figure 3.9 with the timeview of the different logs files on the left and the field summary view which represents the distribution of the values for the variables of the logs files. On the right, the full-sized chart view is the main panel where analysts can filter data and add values of interest. The values of interest collected are displayed in the header panel and can be applied to interact with the data.

The visualizations for inspection are more complex than these dedicated to monitoring and require the analysts to have more time to dive deep into the data. Walton et al. propose the use of a parallel coordinates visualization to helps analysts with queries with conditional attributes (QCATs) [84]. A parallel coordinates graph is used to simultaneously visualize several dimensions, with the dimensions plotted along the vertical axes. The user interface shown in Figure 3.10 consists of multiples panels. The analyst creates a QCAT with panel C to formulate his hypotheses (for instance login tentative from a given IP address between midnight and 3 am). Then he views the results on panel A with a parallel coordinates visualization to detect specific patterns in the data. Panels to share and manipulate QCATs or select specific logs files are also available. When efficient interation is in place, parallel coordinates graphs are efficient to examine how dimensions relate to each other and therefore are often dedicated to Tier 2 analysts.

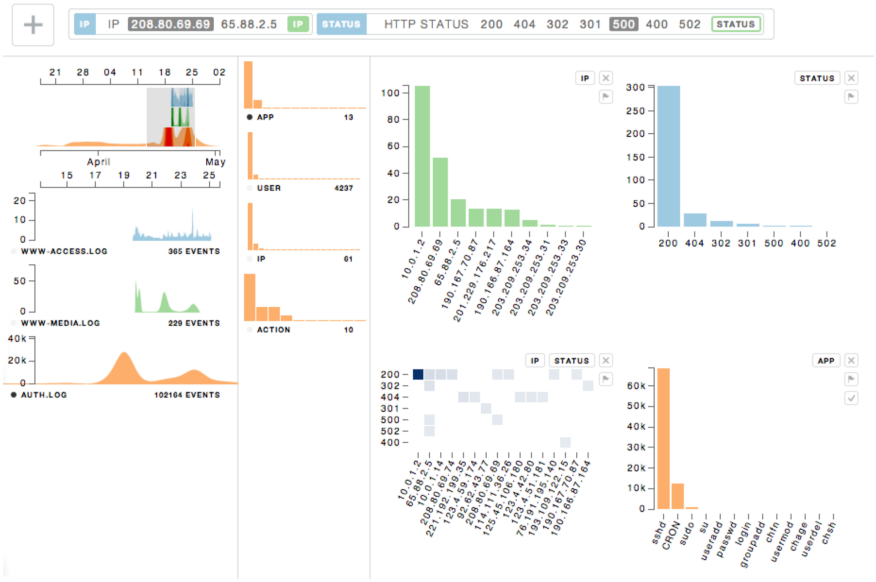

Figure 3.9: An overview of CORGI [83].

Figure 3.10: User interface to visualize queries with conditional attributes [84].

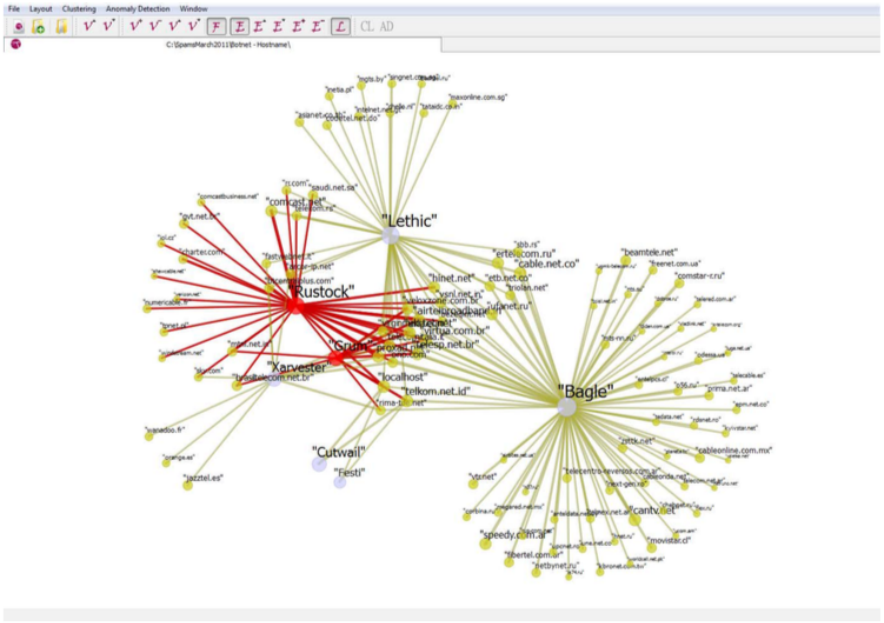

Figure 3.11: Visualization of a spam campaign with [85].

3.4 Exploring

Along with the inspection of specific data, Tier 2 analysts can be interested in exploration, i.e., discovery without a known plan in mind. They can try data combinations, experimentations with data views, to find interesting patterns and search through a long period into the data looking for unusual trends.

For instance, Orestis Tsigkas et al. design a visual analytic tool to explore months of spam data and detect associated botnets [85]. The interactive graph-based visualization, shown in Figure 3.11, represents a spam campaign from March 2011 after the selection of features by the analyst.

As stated earlier, Tier 2 analysts have many data sources available to explore. Therefore, the visualization techniques for the exploring stage are very diverse so that the analysts can freely explore the data looking for different types of patterns. ChangeLink [86], for instance, displays an overview of the directory change to understand how a particular computer was used. Hyungseok Kim et al. propose a three-dimensional representation to facilitate exploration of the current control conditions of firewalls [87].

3.5 Forecasting

Forecasting is the ability to predict future states of the systems. This objective can be accomplished with the visualization of the current state and its progression, to help analysts know what would be the next potential courses of actions. This stage of

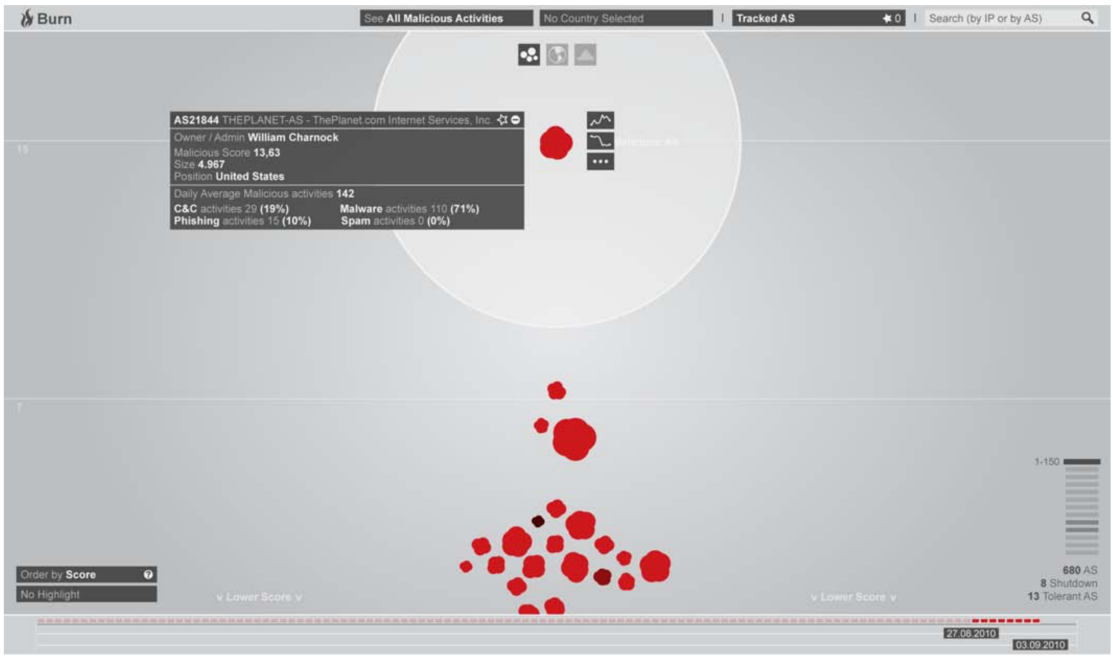

Figure 3.12: Main view of BURN [89]

situational awareness is less popular than monitoring and analysis, so fewer solutions have been proposed in this area.

EMBER (Extreme Malicious Behavior viewER) [88] displays malicious activity at the city level, with a normalized metric called the standardized incidence rate to avoid the emphasize on population-dense metropolitan areas in developed countries (where Internet connectivity is the highest). Using EMBER, Tamara Yu et al. observe that malware preferentially spread to regions with already high levels of malicious activity according to their metrics. By using this tool, analysts can view the current dispersion of malware and try to predict the next targeted region.

Similarly to EMBER, BURN (Baring Unknown Rogue Networks) [89] helps analysts visualize and then predict the malicious activities, but at the autonomous systems level. The goal is to exhibit rogue activities of AS[2] and to find misbehaving networks through visual and interactive exploration. In the main view of this tool shown in Figure 3.12, AS are represented by bubbles. Bubbles are animated so that malicious activity stands out. BURN allows analysts to see data by rank of maliciousness, by geographic location and by historical activity. The understanding of the current dynamic may lead the analysts to understand the future states for the different AS.

[2]Autonomous System. An autonomous system is a network or a collection of network controlled and supervised by a single entity.

3.6 Communication

Communication is rarely an objective in itself; communication is an added feature to the primary objective, it comes in addition to the perception, comprehension, or projection purpose.

Visualization is a great tool to communicate and display information. In a SOC visualization is often used for reporting to colleagues or manager. Visualization for communication focuses on past information, the critical point of communication being that the audience should be able to understand quickly what information is displayed and what is the main takeaway. As a consequence, the visualization employed are pretty simple to highlight this main point.

In our context, communication is about displaying data in a way to transmit specific information. According to our research, no paper on visualization for cybersecurity explicitly targets communication.

Sometimes, the visualization created for another use is self-sufficient to pass the information. This is especially the case when the recipient of the information is in the same team. For instance, the tool presented in [90] proposes a hierarchical sunburst visualization for firewall rules and can show rules based on keywords. For instance, if an analyst wants to communicate about a vulnerability due to firewall rules, he or she can just hand over the visualization resulting from this tool.

VIAssist [91] has an integrated report designer component to allow users to construct these reports easily without even leaving the application. These reports can contain screenshots of the workspace and its components, as well as text and simple drawing annotations. The result in Figure 3.13 shows a report template, automatically populated with two of the visualizations and manually annotated to highlight the data.

3.7 Collaboration

Collaboration is the next step after communication. We propose the addition of collaboration alongside the other uses described previously. People do not simply transmit information, they work together on it. Analysts can work in shifts or at the same time, and they need to share their work to achieve their goal faster or more efficiently. A few pieces of work address collaborative visual analytic between security analysts.

The first pieces of work to add collaborative features in our context were based on tagging. Security analysts add tags (metadata) to characterize and describe the security data they analyzed. By adding these tags, they can share their conclusions and avoid unnecessary jobs for the colleagues. This is the case with FlowTag [92]. VIAssist [91] also offers annotations and tagging to allow users to communicate hypotheses, findings, and status across shifts and locations.

OCEANS [93] is a web-based collaborative interface that allows collaboration between experts of the information system and goes further in the collaboration. The data sources are netflows, IDSes logs and host status logs, with views to detect anomalies

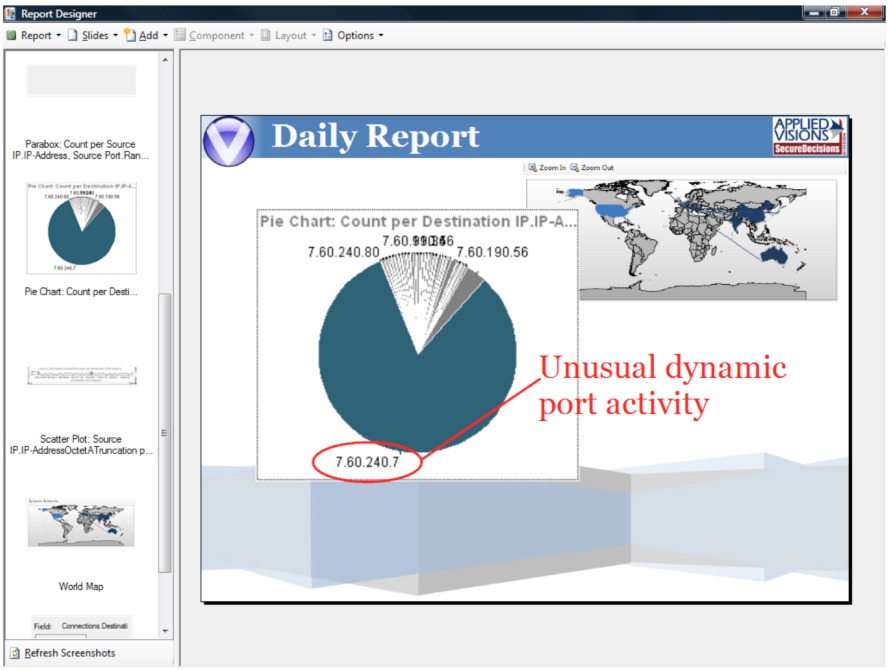

Figure 3.13: VIAssist report [91]

inside these flows. The visualization proposed are advanced enough to be used by Tier 2 analysts. Analysts can submit events and comment them thanks to collaborative features as illustrated in Figure 3.14. OCEANS is dedicated to Tier 2 analysts, with advanced interaction to inspect and interact with the data.

Cyber Analyst Real-Time Integrated Notebook Application (CARINA) [94] is a collaborative tool with the objective to help analysts investigate systems and take decisions. CARINA has dedicated roles (analyst, supervisor, manager, and director) with actions implemented for them. This tool is not designed to deal with primary data sources like log files but to exchange and annotate investigations and reports.

3.8 Conclusion

Visualization is used for multiple purposes in information systems security. We have reviewed in this chapter many visualization tools according to their objective regarding situational awareness with our addition of collaboration between security analysts. Monitoring tools dedicated to Tier 1 analysts use simple visual representations to help analysts perceive the current state of the information system. However, the scalability of these visualizations is usually limited. While many tools are working with a small volume of data, they are often not capable of copping with for real-world scenarios and loads. Inspection and exploration tools offer more diversity and interaction, and are capable of working with diverse data sources. They are used by Tier 2 analysts, and

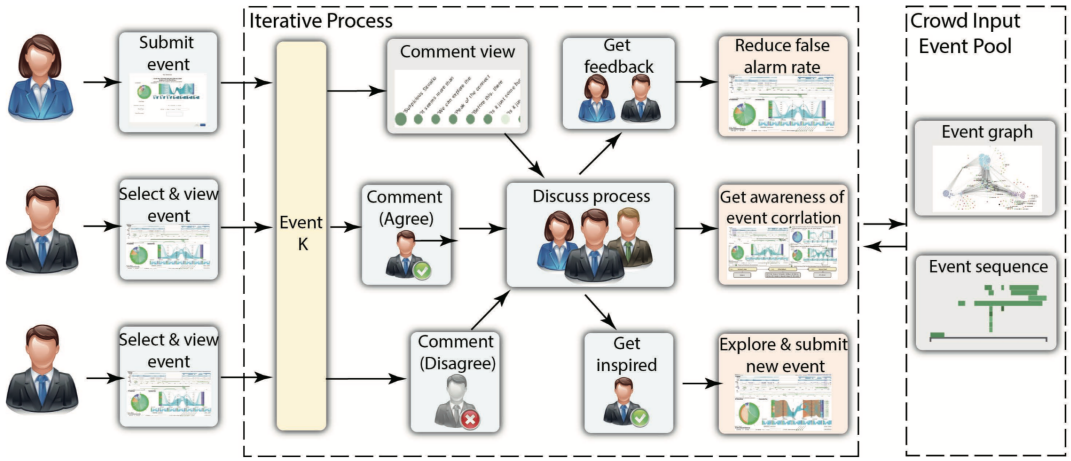

Figure 3.14: OCEANS [93] collaboration diagram.

we advocate that they are a correct response to the technical limitation regarding data presented in Chapter 2. Threat escalation and rhythm of networks are not yet fully answered. Forecasting tools are a category less prevalent in the scientific community, and only present for threat intelligence in mature SOCs. While analysts need communication and collaboration tools to work together, there seems to be a lack of tool offering visual communication and collaboration that could be used in SOCs.

We strongly believe that an efficient way to handle the large quantity of data is to make collaboration among security operators easier by better organizing the workflow through visualization. Our proposal for this are detailed in Chapter 4 and Chapter 5.

4 Visualization for quick triaging

Contents

The number of security alerts arriving in a SOC being high, Tier 1 analysts are drowning into the data and are struggling to efficiently accomplish their tasks. To answer these limitations, we have designed and developed VEGAS (Visualizing, Exploring and Grouping AlertS) [2, 3], an intuitive visualization tool that allows grouping similar security alerts easily and dispatching these groups of alerts among Tier 2 analysts for further analysis. Once a Tier 1 analyst has identified a group of alerts, any forthcoming security alert that belongs to that group will be forwarded automatically to the Tier 2 analyst in charge of the group for further analysis. Therefore, VEGAS reduces the number of alerts received by Tier 1 analysts and makes the flow of security alerts more manageable.

Specifically, we propose the following contributions through VEGAS:

- *Visual exploration of alerts based on principal component analysis (PCA).* Our system uses this technique to convert incoming alerts into two dimensions and represent them on a scatterplot, making it easier for the Tier 1 analyst to identify and group similar attacks.

- *Assisted generation of rules to dispatch alerts.* Once a group of similar alerts has been identified, the Tier 1 analyst can easily interact with VEGAS to bring his skills and generate a rule that describes alerts belonging to this group.

- *Filtering of incoming alerts based on previous rules.* Alerts are dispatched according to the rules that have been generated. Therefore, VEGAS only displays new alerts (i.e., alerts that do not belong to an identified group of alerts that has already been taken care of) and alerts that belong to a group are automatically dispatched in this group for further analysis or persistent storage.

The first part of this chapter addresses the workflow of VEGAS and the technique used to display the alerts in two dimensions. We then describe the interface of this tool and the different interaction mechanisms we developed, before explaining the implementation. VEGAS has been tested both with a use case and by security analysts, the results of the evaluation being detailed in the last section of this chapter.

4.1 Working with security alerts

In this section, we present how VEGAS organizes the workflow of security analysts to make alerts flows manageable. First, we describe the data source in more details and how to represent security alerts adequately thanks to PCA. Then, we explain the global workflow of VEGAS.

4.1.1 IDSes alerts as a data source

IDSes alerts are one of the main sources of data for Tier 1 analysts. Many IDSes are available, Snort[1], Suricata[2] and Bro[3] being common examples. In most cases, and despite the fact that a common alert description format [95] is available, each IDS produces alerts in a format of its own. Whatever format is used, an alert is generally made of a set of fields and in most cases, all the alerts produced by a given IDS are made of the same fields.

Figure 4.1 shows an alert generated by Snort. Several fields describe the network packet that was identified as malicious and the involved rule:

- the *text of the specific rule violated* (between [**]) specifying the triggered rule,

[1] `https://www.snort.org/`
[2] `http://suricata-ids.org/`
[3] `https://www.bro.org/`

```
[**] [1:2100538:17] GPL NETBIOS SMB IPC$
     unicode share access [**]
[Classification: Generic Protocol Command Decode]
[Priority: 3]
04/05-17:55:00.933206
172.23.1.101:1101 -> 172.23.0.10:139
TCP TTL:128 TOS:0x0 ID:1643 IpLen:20
DgmLen:122 DF ***AP*** Seq: 0xCEF93F32  Ack: 0xC40C0BB
Win: 0xFC9C  TcpLen: 20
```

Figure 4.1: A Snort alert.

- the *classification* indicating the type of alerts,
- the *priority* describing the level of criticity of the alert (1 being the most severe and 4 the least severe),
- the *timestamp* defining when the event occurred,
- the *source IP*,
- the *source port*,
- the *destination IP*,
- the *destination port*,
- and other arguably less important fields of the packets. These fields can be useful for specific exploration or inspection by Tier 2 analysts, but we assert that given the little amount of time given to Tier 1 analysts, the other fields should be discarded for the triage.

4.1.2 Displaying alerts

As stated earlier, the main objective of VEGAS is to propose an efficient way to make the flow of security alerts manageable. It does so by allowing Tier 1 analysts to create groups of alerts that are then transmitted to Tier 2 analysts. Groups of alerts should be made logically so that a group has a unique meaning and can be apprehended as a whole.

Several approaches are available to group security alerts. The difficulty comes from the fact that IDSes alerts present many features that are as many dimensions that need to be taken into account to group them. A first possible approach is based on clustering algorithms. Clustering algorithms are part of machine learning and data analysis techniques and are dedicated to the grouping of data points. This requires minimal intervention for Tier 1 analysts, groups being created by the chosen algorithm. Another approach is the visualization. Visualization can show security alerts in specific

representations to enable visual correlation, and Tier 1 analysts will have to use the given representations to manually group the alerts.

We reject clustering methods for the creation of groups for several reasons. First, clustering techniques are not purely automatic tasks. It is often necessary to tune parameters, such as the distance functions or the number of clusters, which requires an expertise in the domain of data analysis that Tier 1 analysts may not have. Furthermore, the issue of the interpretation of the clusters comes up. Tier 1 analysts do not have the time to understand the meaning of the clusters and the possible outliers they will see, which are also linked to the choice of techniques and parameters.

As opposed to clustering algorithms, visualization gives more flexibility to Tier 1 analysts and does not require a specific expertise. Selecting and grouping IDSes alerts to create meaningful ensembles and rules is the task of Tier 1 analysts, given their knowledge about the monitored network. Moreover, it makes sure to bring back human intelligence in the loop and avoid grouping unrelated alerts, the cost of misclassified alerts possibly leading to dangerous consequences regarding the security of the information system.

Visual correlation is effective if humans can detect patterns in the graphical representations, and many dimensions make it difficult to detect relevant patterns. As seen in Chapter 3, several visualization techniques can be used to represent multidimensional data. In our context, we believe that complex representations which try to show all the characteristics of the data are not adapted to the highly intense rhythm of the Tier 1 analysts' work. Due to that fact, we perform a dimension reduction to offer Tier 1 analysts a simple enough representation to detect similar alerts. We choose to display alerts in only two dimensions, on a scatterplot in order to be easily manipulated and grouped. The purpose of our approach is the data summarization as a useful starting point for Tier 1 analysts.

Dimension reduction can be achieved through two different types of techniques: feature selection and feature construction. VEGAS employs a combination of these two techniques. Feature selection is the process of selecting or discarding attributes based on their usefulness for analysis. In VEGAS, we choose not to display the fields we defined as "less important" in the previous section, especially since experiments show that these fields do not allow to create useful groups of alerts.

By contrast, feature construction creates new features using functions applied to the original features, these new features being informative and non-redundant. Several techniques are available to generate two-dimensional layouts from high-dimensional data. The following points were important to choose the technique which best suits needs:

- Unsupervised procedure. Unsupervised procedures mean that only input data is given, in comparison to supervised procedures where training data is needed with data labeled with the appropriate classification. In our opinion, supervised techniques are not to be taken into consideration because training on security data is a very challenging task. There is not necessarily training data available or adapted to the information system, and attacks continuously evolve so the training may become irrelevant.

Moreover, groups of security alerts are subjective and there may be no ground truth to assemble alerts. Some analysts and SOCs may have different policies regarding how to compose a group. In case of a scan of different services, like a DNS and a web server, some analysts want a single group whereas others will prefer a group for each service. This is a crucial difference compared to pattern recognition in computer vision, for example, where the classification output is fixed.

- Interpretation of the result. The visualization resulting should be understood by Tier 1 analysts without requiring too much additional work.

- Information retention. With only two dimensions remaining, the loss of information must be minimal.

- Scalability. The number of alerts being high, the chosen technique must have a low complexity to stay efficient in the long term.

Principal component analysis (PCA) [96], multidimensional scaling (MDS) [97], Isomap [98], Locally-linear embedding [99], and Laplacian eigenmaps [100] were considered. They are all unsupervised and validate the first criterion.

PCA is a transformation that computes linear combinations of the original data into a set of values of linearly uncorrelated variables called principal components. The principal components are ordered by the largest possible variance (meaning accounted for as much of the variability in the data as possible) with the constraint that each principal component is orthogonal to its predecessors. Regarding the interpretation of the result, PCA allows to understand the key variables in the data and to spot outliers. PCA has a complexity of $O\left(p^2 n\right)$ with p being the number of dimensions and n the number of alerts. So this technique has a linear complexity regarding the number of alerts. Moreover, to limit memory requirement, iterative PCA computation mechanisms are available.

Isomap, Locally-linear embedding, and Laplacian eigenmaps are non-linear methods to reduce dimensions, contrary to PCA which is a linear composition method. These type of algorithms assumes that the data of interest lies on an embedded non-linear manifold within the higher-dimensional space in order to produce relevant results, so they are popular in computer vision. We assert that security data does not have this characteristic and the interpretation of the result criterion is not fulfilled.

Multidimensional Scaling is also a non-linear method but reduces data dimensions so that distances between points in the data are preserved. We believe that for an output on a scatterplot, in two dimensions, the result from a PCA is more easily understood by Tier 1 analysts because the PCA is linear so that Tier 1 analysts may be able to view patterns. Moreover, the scalability of PCA is better than MDS, MDS having a complexity of $O\left(n^3\right)$ with n the number of alerts.

This comparison shows that principal component analysis is the candidate of choice.

4.1.3 Computing PCA

We present how we compute PCA using the covariance method. The input data is the set of security alerts. We put them in a single matrix X of $p * n$ dimensions such that it holds n alerts of p dimensions. The dimensions contained in an alert consist of numerical values and categorical values [101]. Therefore, categorical variables are transformed into numeric ones for the computation of the PCA, using the dummy variable creation technique [102]. For each category, a new variable is created, and elements belonging to this category take the value 1 for the new dummy variable, else 0. Thanks to this technique, alerts are now only composed of numeric values and can be used as input for the principal component analysis. The data is also normalized.

The next step is to calculate the covariance matrix C of the alerts. This can be done using the following equation with X^T being the transpose of X:

$$C = \frac{1}{n-1} X^T X \tag{4.1}$$

The covariance matrix C is symmetric and contains the variance of dimensions as the main diagonal elements and the covariance of dimensions as the off-diagonal elements. We now calculate the eigenvectors and eigenvalues for the covariance matrix C. This is done by computing the matrix which diagonalizes C. Eigenvalues and eigenvectors are ordered and paired, and they contain useful information about the data. The eigenvalues represent the distribution of the source data's information among each of the eigenvectors. The eigenvectors are perpendicular to each other and form a basis for the data.

We then sort the eigenvectors according to their eigenvalues in decreasing order. We choose to keep the first two principal components to display our data, with the risk of rejecting other potentially relevant components. We point out that the loss of variance resulting from the selection of the two first principal components can be quantified according to the value of the remaining eigenvalues.

Finally, we transform the original alerts into two dimensions with the next equation. These two dimensions will be used to plot the alerts on a scatterplot. We indicate the dimensions of the different matrices in this equation, and the matrix N is the outcome of the PCA.

$$N_{2*n} = [top\ 2\ eigenvectors]_{2*p} X_{p*n} \tag{4.2}$$

Our outcome of the PCA is the projection of our dataset (consisting of numerous multi-dimensional alerts) onto a smaller subspace made of two dimensions that is representative of the original dataset. While these two new dimensions have no real semantics, they are the composition of the most important dimensions in the original dataset. Therefore, in these two dimensions, alerts that are close in the original dimensions are still close in the newly computed ones while alerts that are distant in the original dimensions are distant in the newly computed ones.

4.1.4 Workflow

Figure 4.2 summarizes the workflow of VEGAS. Alerts generated by the IDSes, at the top of the figure, are transmitted through the network to a filter that dispatches them. Initially, the filter only has the *default* rule that sends all the alerts to the Tier 1 analysts for display on the VEGAS interface, This *default* rule is in bluish-gray color at the bottom of the figure. When a Tier 1 analyst identifies a new group of alerts using the interface of VEGAS, he or she performs a quick analysis of it, annotates it, and adds a new filtering rule to the filter to redirect these alerts to a new destination called a bucket to be analyzed by a Tier 2 analyst. From this moment, the group of alerts that has been identified by the Tier 1 analyst disappears from his or her interface and is sent directly to the proper bucket according to the rule matched. Forthcoming alerts matched by the rule will also be sent directly to this bucket and won't be displayed on the Tier 1 analyst interface. All security alerts, belonging to a filtering rule or not, are also sent to a persistent storage. This way they will still be available for forensics.

After the manipulations by Tier 1 analysts and the creation of a new rule, the PCA is recalculated. This new computation takes into account the new alerts which have arrived and the still unfiltered alerts.

4.2 VEGAS interface

We first provide an overview of the interface of VEGAS in Section 4.2.1, and present how Tier 1 analysts interact with this interface. Then we show how alerts are grouped and how to generate new dispatching rules.

4.2.1 Overview of the interface

Figure 4.3 shows the top of the VEGAS interface presented to Tier 1 analysts. This interface is displayed in a web browser and it was designed with a drill down approach in mind, following Shneiderman's mantra [103]: overview first with the scatterplot, zoom and filter, then details-on-demand with the other representations.

At the top, the header provides general information about the current situation: under the "IDS alerts" title, the total number of alerts currently displayed and the number of selected alerts by the analysts are given. Here, the user has selected 644 out of 1 021 alerts. This information is needed to indicate the volume of alerts and to give an idea of the proportion of the selected alerts. The number of new alerts since the last time the analyst has refreshed the interface is also given (currently 0). The "Reset All" link allows to remove all selections and manipulations accomplished by the user. The rules previously created are not impacted by this action. The "Generate rules" link creates filtering rules based on the characteristics of the selected alerts. These two actions are described with more details later.

The time graph, in zone A, then provides the distribution of alerts over time. It first gives information to the Tier 1 analyst about the variations in the volume of alerts that

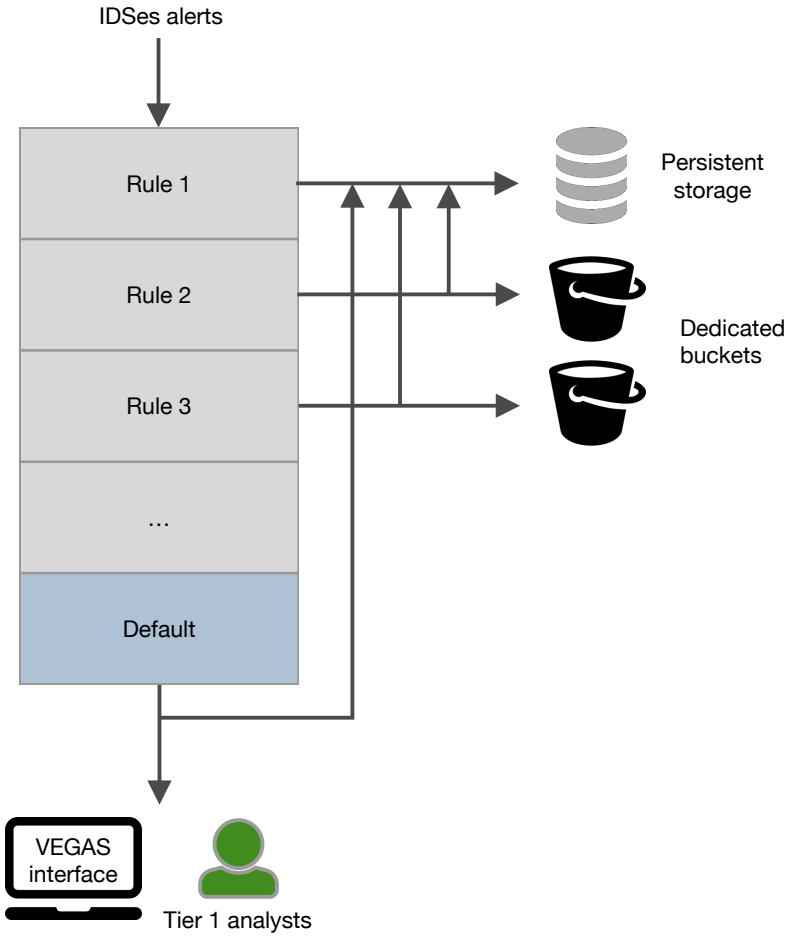

Figure 4.2: VEGAS workflow.

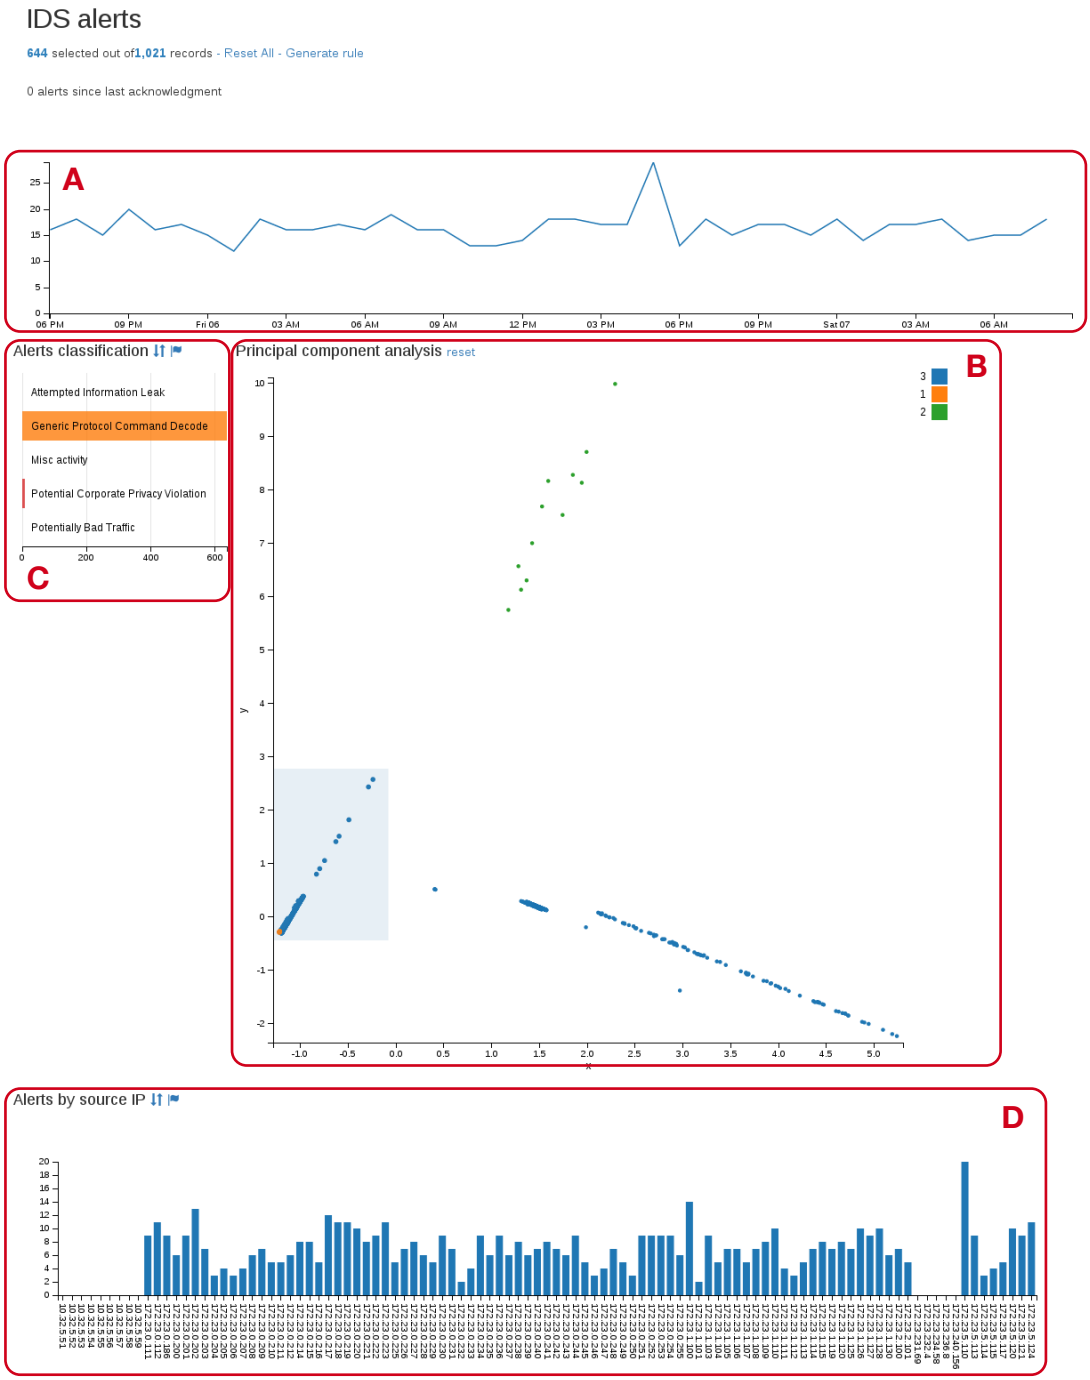

Figure 4.3: VEGAS interface for Tier 1 analysts (beginning).

could indicate changes in the type of attacks (like the start of a DDoS). It also provides information to the analyst about the first and last alerts that have still not been dealt with. On the figure, alerts are distributed between 6 pm on Thursday and 6 am on Saturday. The user can also zoom in to select a specific period.

Below in zone B, the scatterplot displays the result of the PCA that was applied to the alerts. The *Priority* field of the alert is shown on the scatterplot using colors to transmit this information to the user, 1 being the most severe. On the figure, we can clearly see groups of alerts that form several lines. The data used here is from the 2012 VAST Challenge [5], and more explanations about the patterns are given in the evaluation section of VEGAS (Section 4.3.2).

Finally, all the other features of the alerts are shown. *Alerts Classification* in zone C, *Alerts by source IP* in zone D, *Alerts by destination IP*, *Alerts by source port* and *Alerts by destination port* bar charts present the distribution for these features for the alerts that have been selected by the Tier 1 analyst in the scatterplot. Due to space constraint, Figure 4.3 does not display all the bar charts.

At the bottom (still not shown on Figure 4.3), selected alerts are listed in their raw form. This way, Tier 1 analysts can directly view the subset of initially matched alerts.

4.2.2 Analyzing alerts

When Tier 1 analysts launch VEGAS for the first time, the interface displays all the alerts that match the default dispatch rule, i.e., alerts that do not belong to groups that have already been identified and sent to Tier 2 analysts.

First, Tier 1 analysts benefit from an overview of the distribution of alerts over time. Below this time graph, the scatterplot resulting from the PCA computation is displayed. As can be seen in Figure 4.3, patterns can be detected on this representation: alerts that are close in all the original dimensions are close in the two-dimensions space resulting from the PCA computation.

The objective for Tier 1 analysts is to understand the patterns and create meaningful groups. To do so, Tier 1 analysts can select a group of alerts on the scatterplot. All other representations (*Alert classification*, *Alerts by source IP*, *Alerts by destination IP*, *Alerts by source port*, and *Alerts by destination port*) are automatically updated to display only the values exhibited by the selected alerts. After the selection of a group, the goal is to understand the features shared by the grouped alerts. VEGAS offers bar charts representing the values exhibited by the alerts for each field and specific interaction adapted for quick manipulations. Bar charts have been proven to be very efficient to represent categorical fields when they can take numerous different values [83]. Tier 1 analysts can select specific values in each of the fields to look at the distribution of the various other fields of the set of selected alerts. Each time a value or a set of values is selected, the scatterplot and the various bar charts are updated to reflect this selection. This mechanism allows the analyst to better understand the selection and to perform more subtle selections than the ones that would be available in the scatterplot.

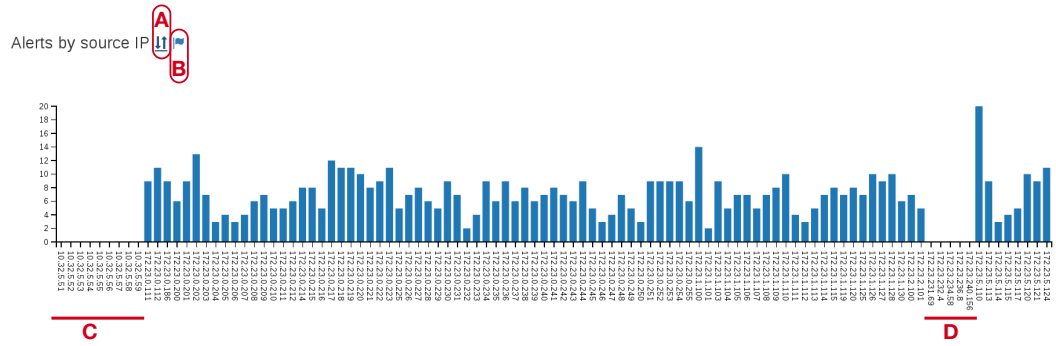

Figure 4.4: Before filtering on the scale.

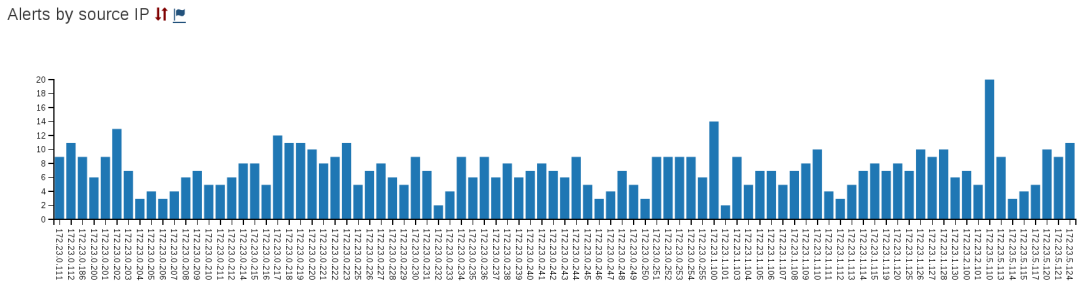

Figure 4.5: After filtering on the scale.

Visualizing the repartition of values is a way to exhibit interesting behavior. Tier 1 analysts can click on the double arrows close to the field name to *filter* the scale on the values present for this field, meaning "the whole set of values that are present in the selected alerts" and not in all alerts. This icon is shown in zone A in Figure 4.4. If Tier 1 analysts want to see the whole scale again, they just have to click another time. When clicked, the icon changes its color in order to recall analysts that they are currently filtering on the scale. This interaction helps analysts in viewing the values for the field in the selected alerts while keeping the possibility to perceive the relative distribution of these values compared to the full set of alerts.

As illustrated in Figure 4.4 and in Figure 4.5 patterns can be detected thanks to this representation. The first observation is that the alerts are coming from internal computers. The IP addresses underlined in zone C are all external to the networks, as shown with their IP addresses (`10.x.x.x`). One can see that five consecutive IP addresses underlined in zone D (`172.23.231.69`, `172.23.232.4`, `172.23.234.58`, `172.23.236.8`, `172.23.240.156`) are not targeted, whereas all other internal addresses are. This observation may be a hint that these addresses have a specific status and the analysts need to investigate it, or that these computers have not been yet infected and may need a specific protection.

4.2.3 Generating relevant filtering rules

Using the scatterplot resulting from the PCA, Tier 1 analysts detect groups of alerts, and perform a quick analysis to understand the main characteristics of this set using the interaction and filtering possibilities described in the previous section. We should emphasize that the interactions offered by VEGAS allow this analysis to be quite freely performed. Therefore, Tier 1 analysts can use all of their skills and knowledge of the context to perform their task.

Tier 1 analysts have to generate the rule to be inserted in the dispatch filter. To do so, they need to select the relevant fields to be included in the rule. While the set of alerts that leads to a rule has a set of values for each of the fields, only some of these fields are relevant. For instance, when analyzing a brute force attack against an `ssh` server, the destination port is undoubtedly of interest, the source and destination IP might be, but the source port is probably of little relevance. This fact is reflected in the distribution of the various fields. Regarding the alert classification, all the alerts should belong to the same classification. All the alerts share a common destination port and if some specific servers are targeted, there is only a few different destination IP addresses. If there are only a few attackers, there are only a few source IP addresses. Therefore, Tier 1 analysts need to select the relevant fields to be included in the rule to be generated by clicking on the flag near the name of the field. The flag is shown in Figure 4.4 in zone B. This flag then changes color and it is possible to revert this selection by clicking on the flag once again. Here, we should emphasize that the selection of the relevant field depends on the Tier 1 analyst *a priori* knowledge and/or way of understanding attacks. We advocate that this added expert knowledge is, in fact, an added value of our proposal.

Once a group of similar alerts has been identified and the relevant fields selected, Tier 1 analysts click on the "Generate rule" link to generate a new rule to be inserted in the dispatch filter. The analysts can then give a title to the new rule and add comments to help other analysts better understand the identified group of alerts.

The filtering rule generation is a simple automated process. All the selected values of the selected fields for the selected alerts are put in a dictionary. For instance, the filter that gets alerts for two servers receiving orders from a botnet using the IRC protocol is shown in Figure 4.6. The Tier 1 analyst has identified a group of alerts coming from the same source port 6667. These alerts are qualified as *Misc activity* by Snort and are targeted two different IP addresses which are 10.32.5.54 and 10.32.5.56. Port 6667 is usually associated with the IRC[4], a protocol to communicate in the form of text. IRC is often used by attackers to transmit commands to infected computers. So the Tier 1 analyst has selected this fields and created a rule. All the alerts matched by the filter are immediately reclassified according to this new filtering rule, including the alerts that were displayed as well as the forthcoming alerts. After the creation of a rule, the PCA is calculated once again.

[4]Internet Relay Chat

```
"name": "command-and-control (C&C) server",
"comment": "IRC traffic for communication",
"filter": {
          "sourcePort": [6667],
          "classification": ["Misc activity"],
          "sourceIP": ["10.32.5.54","10.32.5.56"]
}
```

Figure 4.6: A rule generated to filter alerts for two servers giving or receiving orders to or from a botnet using the IRC protocol.

4.2.4 Viewing filtering rules

Analysts need to know the current rules in the system, and the evolution of alerts matched by them over time. The objective is to see the trends of attacks happening. In order to fit this need, VEGAS proposes a representation of the created filtering rules. Figure 4.7 shows the distribution of all alerts over time, the classified ones and the non-classified ones. It is possible to view the number of alerts filtered by a specific rule by clicking on it. For instance, Figure 4.8 illustrates the evolution of the number of alerts filtered by the rules called *rule2*. This shows that this group of alerts starts to arrive at 7 pm. This is a way to understand the sequence of events occurring on the network. It also helps in understanding when an attack started and if it is still going on. Finally, the list of rules is displayed, as seen in Figure 4.9.

4.3 Implementation and evaluation

4.3.1 Implementation

VEGAS was implemented as two distinct parts, a server and a client. On the server side, the IDSes alerts and rules are stored into Elasticsearch[5], a highly scalable open source search engine with a REST API. Logstash[6] is used to parse the alerts. In a first version, PCA was computed using the programming language R[7] and the package FactoMineR[8]. However, due to performance limitations regarding input/output communications with Elasticsearch, the PCA is now computed server-side with the Python language [9] and the package scikit-learn[10]. Results are directly stored in the Elasticsearch server. We developed a specific plugin in Python to filter incoming alerts and tag them with the

[5] https://www.elastic.co/products/elasticsearch
[6] https://www.elastic.co/products/logstash
[7] https://www.r-project.org/
[8] http://factominer.free.fr/
[9] https://www.python.org/
[10] http://scikit-learn.org/stable/index.html

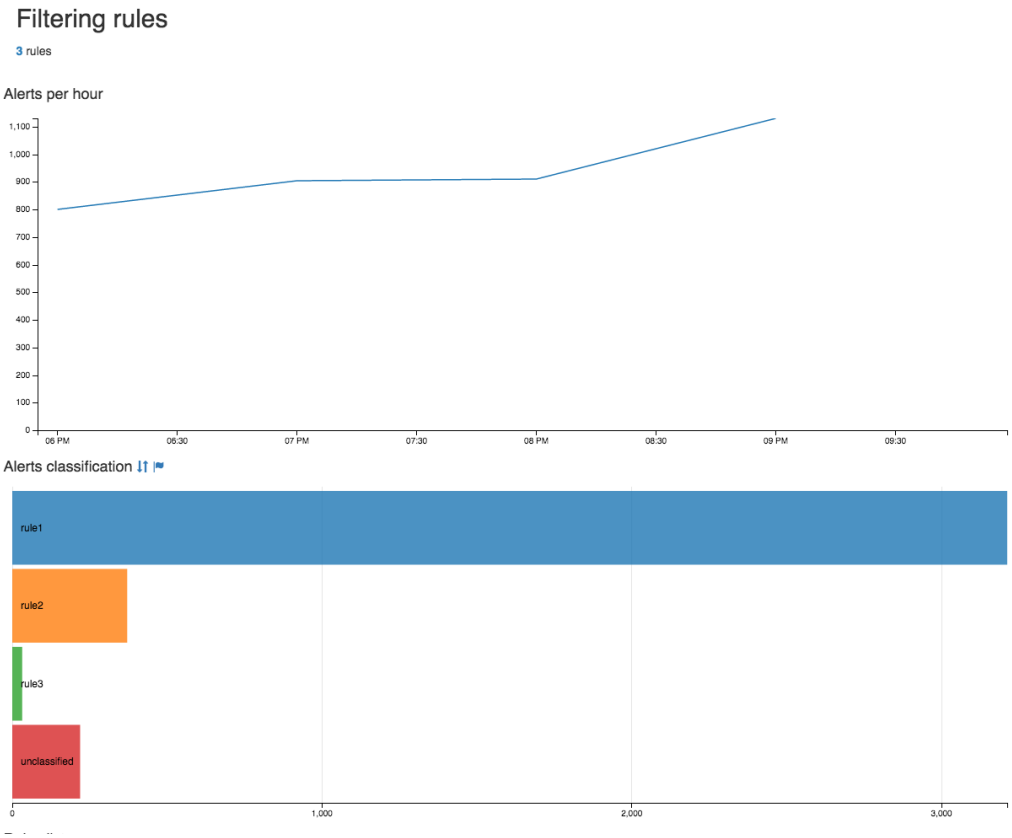

Figure 4.7: Representation of filtering rules over time.

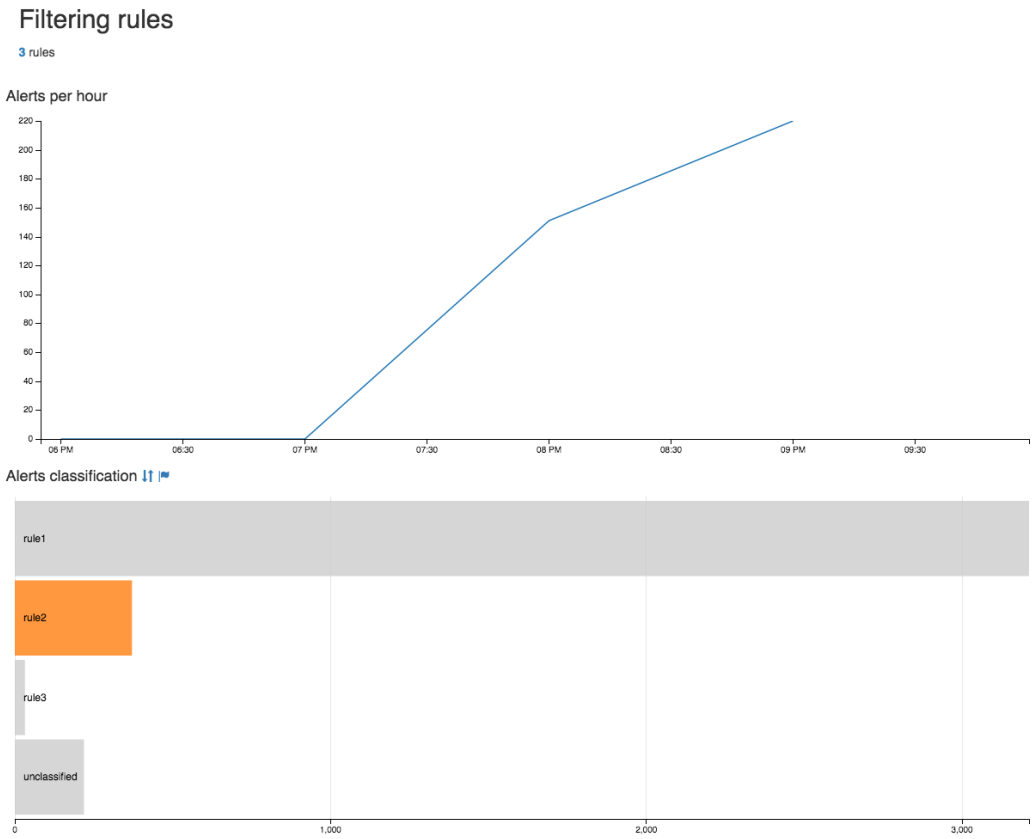

Figure 4.8: Evolution of alerts filtered by rule2.

Rules list

Delete selected rulesMerge selected rules

name	timestamp	content	action
Rules			
rule1	2016-09-19T09:50:03.155Z	{"destPort":["445"],"destIP":["172.23.0.10"]}	Select
rule2	2016-09-19T09:55:21.017Z	{"sourcePort":["6667"],"sourceIP":["10.32.5.56","10.32.5.54","10.32.5.57","10.32.5.51","10.32.5.52","10.32.5.58","10.32.5.59"]}	Select
rule3	2016-09-19T09:58:21.266Z	{"sourceIP":["172.23.240.156"],"sourcePort":["39945","39944"],"destIP":["172.23.0.1"]}	Select

Figure 4.9: Rules list.

matching rules to dispatch alerts correctly. This choice of technology is based on a mature stack which is highly scalable.

The client side of VEGAS is implemented using web technologies: HTML5, Javascript, CSS, and SVG. The representations and charts are built using the D3.js[11] library. Filtering functions are implemented using the Crossfilter[12] library that is very efficient for fast interaction with large datasets.

4.3.2 Evaluation

We used two strategies to evaluate VEGAS. VEGAS was first evaluated using use cases. Then, we performed a field experiment with the experts using VEGAS in close-to-reality conditions.

4.3.3 Use case

First, we used the Snort logs of the 2012 VAST Challenge [5] to perform experiments with VEGAS. The 2012 VAST Challenge focuses on visual analytic applications for both large-scale situation analysis and cybersecurity. It contains a challenge in where unusual events are occurring in one of the Bank of Money's regional offices. The background story of this challenge is the following: some staff members report unwanted messages appearing on their monitors, declare that their systems seem to be running more slowly than usual and that the hard disks seem to always be running. The description of the network is given in Appendix A.

In the challenge, IDSes alerts and firewall logs are provided to find the origins of these events. We only used the IDSes alerts to perform our investigations. During the three days of capture, more than 50 000 Snort alerts have been generated.

The computers we used had the following configurations:

- The server was a quad-core Intel Core i5 at 2.67Ghz with 4 GB of memory.
- The client was a quad-core Intel Core i7-4600U at 2.10GHz with 8 GB of memory running Chromium version 43.

The first 4 000 Snort alerts were used to initiate PCA computation. This number of alerts corresponds to the first four hours of the challenge, from 6 pm to 10 pm. Figure 4.10 shows the representation generated by VEGAS. Three groups of alerts clearly appeared on the scatterplot: group A, group B, and group C. Due to the priority of alerts, given by the color on the graph and knowing that a priority of `1` (high) is the most severe and `3` (low) is the least severe, the group A on the bottom left was the first one to be inspected.

By analyzing it using the interactions and functions offered by VEGAS, we were able to learn that this group is composed of 3 570 alerts (i.e., the vast majority of the 4 000 alerts) and that it could actually be split into two distinct categories:

[11] `http://d3js.org/`
[12] `https://square.github.io/crossfilter/`

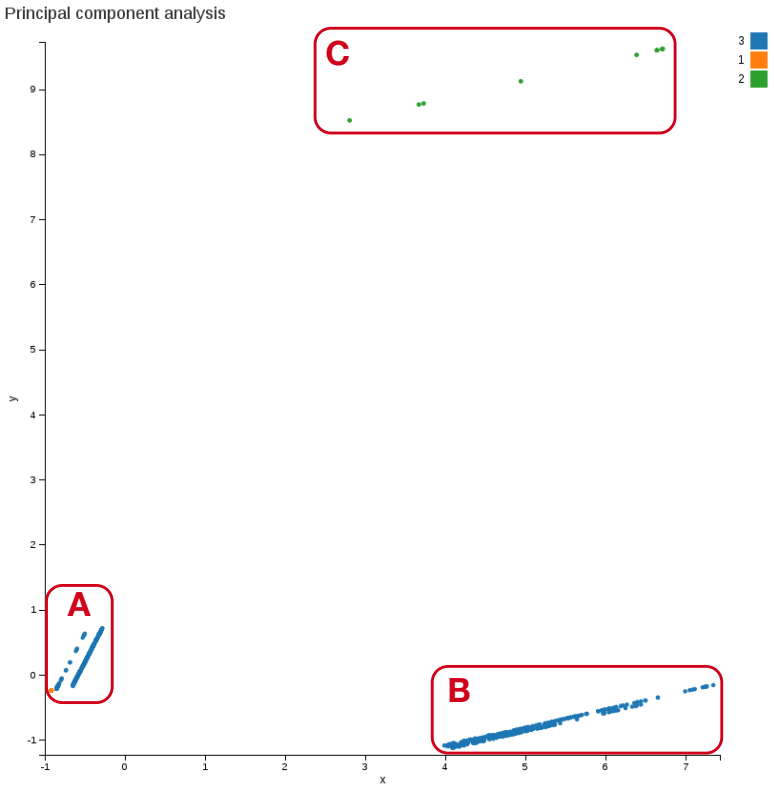

Figure 4.10: Representation after the 4 000 first alerts.

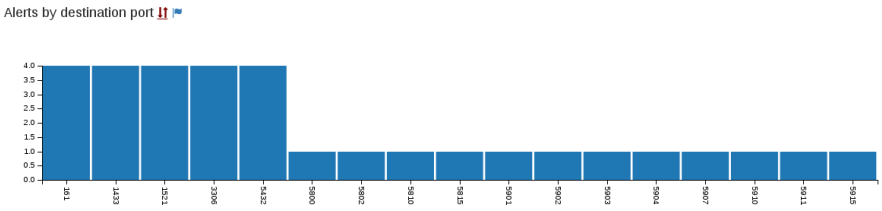

Figure 4.11: Alerts by destination port. Given the distribution, it may be a scan.

- The first one contains only 60 alerts, that were raised during the first hour and was defined by suspicious traffic toward the DNS server on the port `53`. These alerts are shown in orange. Regarding the severe priority, these alerts should be analyzed more in depth by a security analyst to see if this threat was as severe as it seems.
- The second category, with 3 510 alerts, was differentiated from the first one by the destination port: `139` and `445`. They appear in blue. These two ports are used for Microsoft file sharing technologies and are often targeted by attackers [104]. The time graph (not shown) indicated that contrary to the traffic toward port `53`, alerts of this type were still arriving at the rate of 900 alerts per hour.

These two different sub-groups were visually close on the scatterplot because all the source IP addresses are internal IP, meaning that workstations had very probably been compromised. This is consistent with the observations of the staff.

We then studied the two other groups of alerts:

- Group B is a flow of alerts beginning at 7 pm. Seven external IP had been communicating with many internal IP using the port `6667`. These was probably C&C[13] connections through IRC giving orders to the compromised workstations.
- Group C at the top was composed of 32 alerts, which arrived around 9 pm. These alerts were characterized by a single source IP address, `172.23.240.156`, and a single destination IP address, `172.23.0.1`. Using contextual information provided for the challenge (see Appendix A), we could infer that a workstation inside the bank network was targeting the firewall. Using the zoom interaction, we obtained Figure 4.11. Given the number of different destination ports, the repartition of the alerts with a few of them per port, and the classification given by Snort, we believed that a scan of the services on the firewall was currently happening.

At this point, we created four rules with the interface of VEGAS, according to the four different groups we had discovered and the meaningful fields used to describe them. Thanks to the proposed interaction, this only took a few seconds and allowed to quickly dispatch alerts to the security analysts.

During the first iteration of the analysis, other alerts arrived waiting to be displayed on VEGAS, so we refreshed the interface to display them. We performed a few iterations,

[13] Command-and-Control

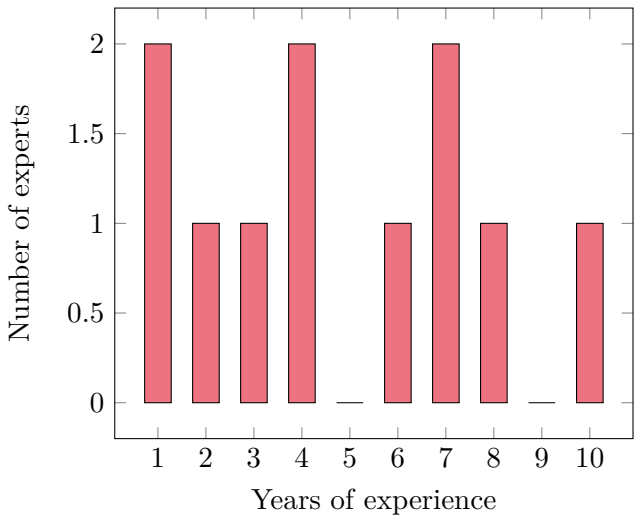

Figure 4.12: Years of experience of our panel for each participant.

each of them using the alerts that had not otherwise been dispatched by the already created filters. For instance, thanks to the scatterplot and the new groups which appeared, and the interaction proposed by VEGAS, we discovered that two other external IP were creating suspicious IRC traffic.

4.3.4 Evaluation by experts

Our test of VEGAS with the case study was a first evaluation and showed that this tool was promising. We took the evaluation one step further by meeting eleven experts, working or having worked in a French defense agency, in SOCs, or in the French army. They were all male with one to ten years of experience in the field as shown in Figure 4.12. While being a relatively small panel, although it is consistent with much of the information visualization literature [91]. Moreover, age experience and background of the experts' panel were diverse.

Methodology

As in our test, we used the second challenge from VAST 2012. First, the context of the mini challenge was explained to the participants. Then we presented VEGAS to the users and its possible interaction.

For the evaluation, the user had access to VEGAS with the first two hours of the challenge. When the majority of alerts had been filtered, we inject new data, as it happens in a real situation.

At the end of the session (approximately half an hour), five questions were asked:

- Q1: Is the problem of alerts triaging relevant?

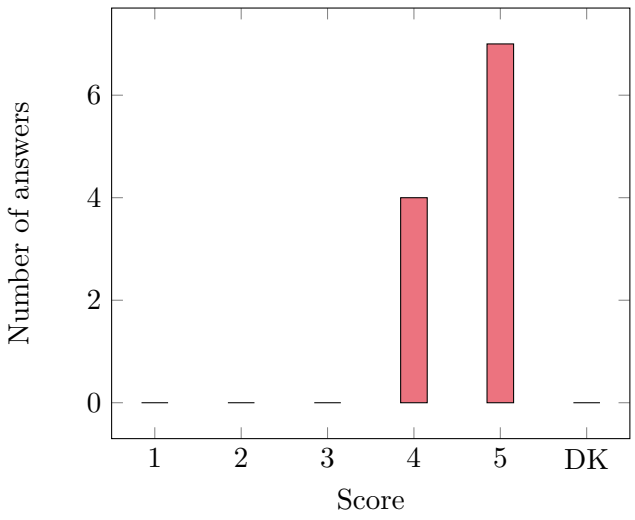

Figure 4.13: Q1: Is the problem relevant?

- Q2: Are the proposed visualization relevant to answer this problem?
- Q3: Is the proposed interaction relevant?
- Q4: Is VEGAS usable?
- Q5: To what extent VEGAS improve the productivity of Tier 1 analysts?

The questions had an interval scale ranging from 1 to 5, with 1 being not at all and 5 very good. DK, for do not know, was also a possibility. These questions are based on the work of Staheli et al. [105] and Angelini et al.[106]. The first three questions evaluate the relevance of the addressed problem (Q1), VEGAS adequateness to this problem (Q2) and the relevance of interactiveness (Q3). The last two questions are aimed at assessing VEGAS usability (Q4) and effectiveness (Q5). Participants of this evaluation can also report comments and suggestions.

The first question was the need to develop a tool with the objective to help Tier 1 analysts perform alerts triaging. We can see the result in Figure 4.13 that all users agree to say that the problem we try to solve is relevant. A user told us that the job of the Tier 1 analyst was a boring one and changes need to be done in order to improve the situation.

The majority of the users found that the visualizations are relevant for our problem. The main positive point is the feedback about the scatterplot based on PCA. After a few minutes, experts understood what the patterns they saw were about, and why some alerts were grouped and aligned. Because the PCA is regularly calculated, the two principal dimensions can change. This change of coordinate system has provoked a surprise for some users; however, they understood it after several iterations with VEGAS. Some remarks have been made regarding the scalability of the bar charts. When two

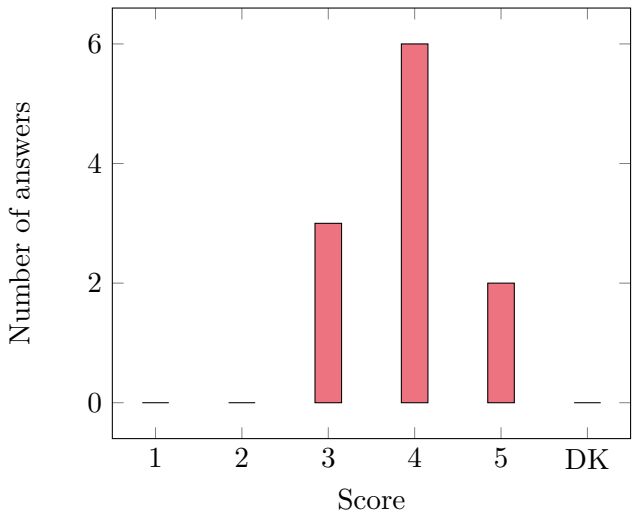

Figure 4.14: Q2: Are the proposed visualizations relevant for our problem?

many values, like IP addresses, are shown, the legend can be unreadable. We argue that if filters are created correctly, the number of incoming alerts will be low and so this situation should not appear. Another type of visualization could also be an answer to this problem, with a representation of the subnetworks for instance.

The results for the question related to the relevance of the proposed interaction are given in Figure 4.15. They are similar to those about visualization, so the experts positively evaluate the simplicity and the basic interaction in VEGAS. There are still some improvements to be made. The user spends a good amount of time scrolling through the page to look or flag specific fields (one said 'I have to scroll all the time'). This is linked to the number of variables to show and the fact that there is no remainder of the current fields flagged so the user has to check the color of the flags for the fields. Adding the name of the current fields flagged on the top of the interface may be a solution to this.

The results for the question related to the usability of VEGAS are shown in Figure 4.16. Once again, the answers are positive. Users added that a brief period of adaptation was needed to understand the interface and the interaction, but after that it was good. The fact that VEGAS can be used easily proves that VEGAS is a major improvement for Tier 1 analysts.

Finally, the last question was about the productivity improvement which may be brought by VEGAS. As shown in Figure 4.17, experts acknowledge our approach for the triaging of IDSes alerts and think that VEGAS can enhance the productivity of Tier 1 analysts. One expert added that Tier 2 analysts could use this tool for a more in-depth analysis of security data. Thanks to the visualization and filtering capabilities, analysts may add other data sources than IDSes alerts and split them into meaningful groups. This expert cited the example of quickly sorting applications logs.

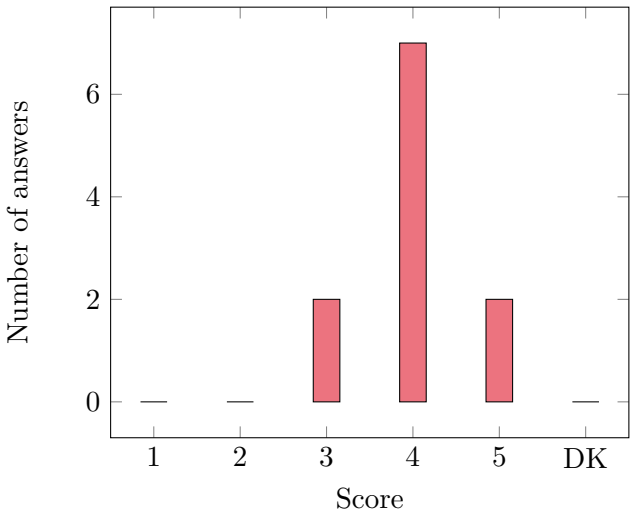

Figure 4.15: Q3: Are the proposed interaction relevant for our problem?

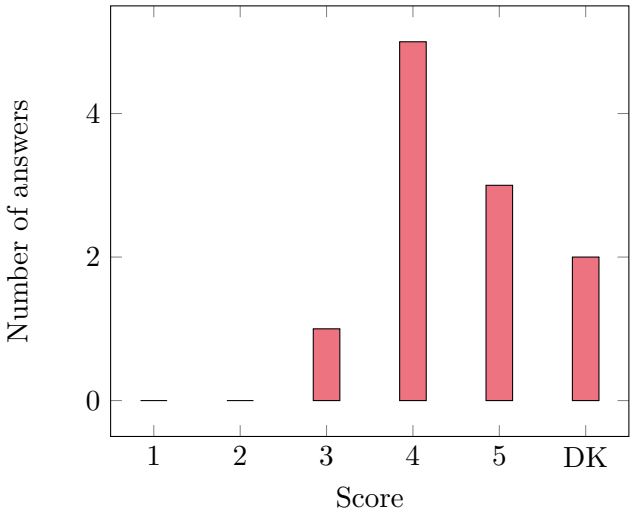

Figure 4.16: Q4: Is VEGAS usable?

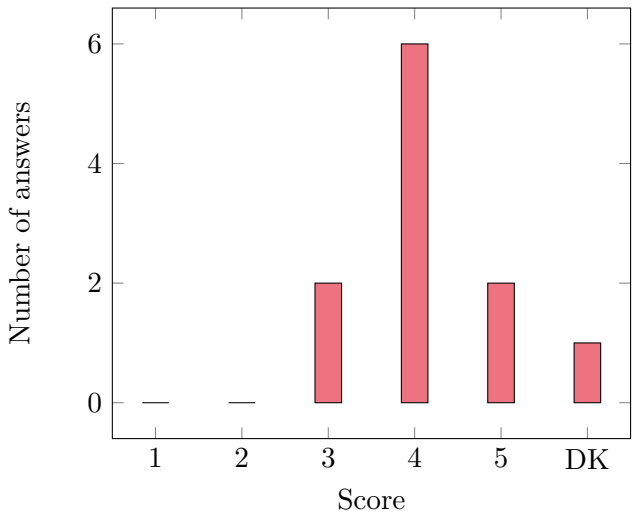

Figure 4.17: Q5: Will VEGAS improve the productivity of Tier 1 analyst?

4.4 Conclusion

In this chapter, we presented VEGAS, a visualization and classification tool that allows Tier 1 analysts to manage important flows of IDSes alerts. VEGAS uses principal component analysis to produce a two-dimensional scatterplot representation that can be used to visually correlate and group IDS alerts. VEGAS allows Tier 1 analysts to explore the various fields of similar alerts to analyze them quickly and to generate meaningful dispatching rules that cause similar alerts to be appropriately forwarded to Tier 2 analysts. This combination of data summarization and visualization is the central point of our proposition. Evaluations with a case study and field testing have demonstrated that VEGAS is useful to quickly detect similar IDSes alerts and group them efficiently and therefore reduces the load of Tier 1 analysts.

With VEGAS we developed an answer to the technical limitations related to the high numbers of security events. Moreover, VEGAS avoids the repetition of the same task for Tier 1 analysts and the creation and quick analysis of groups of alerts mitigates the problem of the lack of creativity.

5 Visualization for collaboration between security analysts

Contents

In Chapter 2, we exhibited the limitations currently impacting SOCs. With VEGAS (detailed in Chapter 4) we have developed an answer for Tier 1 analysts to the technical limitations related to the high number of security events, alongside the repetition of the same tasks and the lack of creativity. This chapter presents TheStrip [4], a new process and visualization tool to enhance collaboration inside SOCs. We detail our proposition for the remaining process limitation we have identified which is the lack of feedback between Tier 1 analysts and Tier 2 analysts. Our proposal also addresses the technology limitations of the rhythm of networks and the threat escalation.

The process we propose is established by constructing rules to define security meta-events and creating a specific feedback loop between Tier 1 analysts and Tier 2 analysts. This workflow makes the work of Tier 1 analysts easier while keeping them under the supervision of Tier 2 analysts. To support this new process and answer the remaining technology limitations, we then present a visualization tool. TheStrip, as a combination of this process and tool, enables collaboration around security incidents and security events inside SOCs, and offers the following features:

- *Quick perception of the context.* Thanks to the division in security meta-events and the view as a timeline, the current situation of the information system can be quickly seen by security analysts. Tier 1 and Tier 2 analysts are able to understand what were the latest changes.

- *Visual reconstruction of attack scenarios.* Tier 2 analysts can easily link security meta-events together to create attacks scenarios and show the evolution of threat.

- *Visual correlation of incidents and security events.* Current trends for incidents and the rules are shown, and the timeline helps security analysts visually find relations between security events over time.

This chapter begins with the lessons we learned from our interviews with experts. Then our new process is detailed with the definition of rules for security meta-events and the creation of a specific feedback loop between the two groups of security analysts. The visualization developed for the process is then described.

5.1 Current process in security operations centers

During the interviews we performed with twelve security analysts, we asked them about their experience to define how SOCs in practice are different from the literature. This section describes the current process happening in SOCs with a focus on collaboration and their limitations.

Analysts we interviewed established that the fundamental task inside a SOC is the real-time analysis of data feeds. When SOCs grow in experience and skills, other capabilities can be covered without compromising SOCs primary capabilities. This is consistent with the division between primary and secondary capabilities seen in Section 2.2. Security analysts added that it is essential to improve the basics of a SOC before meeting all capabilities and adding more missions to achieve. Having advanced SIEM software products with trendy functions was also declared less important than benefiting from a proper implementation of the primary tasks, which, for Tier 1 analysts, evolve around triaging security events.

We used the knowledge gained from the interviews and the information provided by the experts to identify the workflow that is currently in place in actual SOC organizations. This workflow is described in Figure 5.1. The CLUSIF organizational model (detailed in Chapter 2) is used for comparison because it is representative of the literature we reviewed. The main difference between this model for the SOC incident management zone and the experience from the experts is in the relation between Tier 1 analysts and Tier 2 analysts. When Tier 1 analysts receive security events, they do a quick triage according to a knowledge base. If the event is documented, Tier 1 analysts follow the procedure which leads to a qualified incident or a false positive according to the result of the procedure. If the event is not documented, Tier 1 analysts report a suspicious event to Tier 2 analysts.

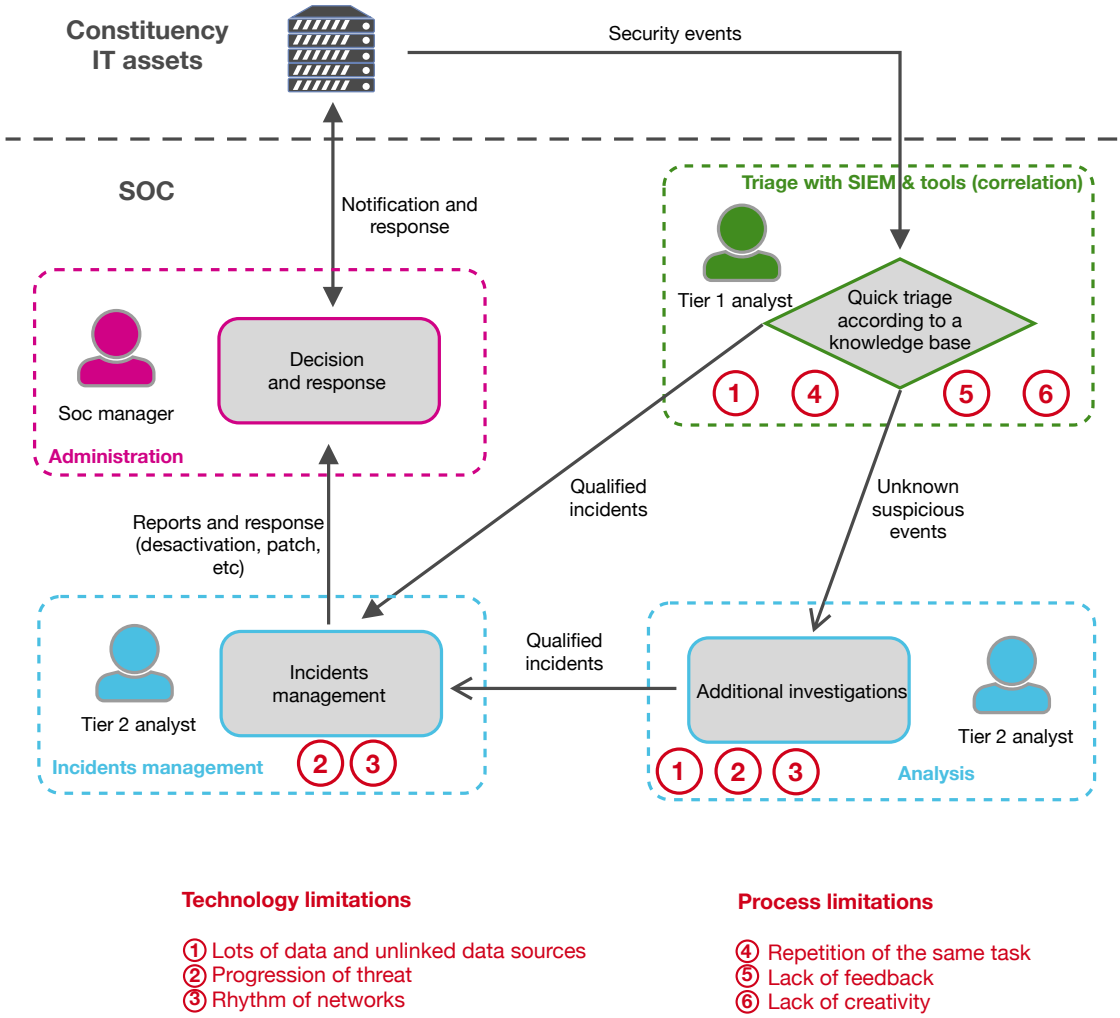

Figure 5.1: Current organization of SOCs with their limitations.

Tier 2 analysts perform two main tasks, depicted by the two blue rectangles on Figure 5.1. On one hand, they analyze unknown events that are suspicious and, following the result of their investigation, create a new qualified incident if needed. On the other hand, they manage the incidents and the creation of an appropriate response to them. A given Tier 2 analyst can be responsible for both tasks or be dedicated to one. Our representation of the organization highlights this separation of tasks for Tier 2 analysts. Also, Tier 2 analysts have to report periodically to the SOC manager, and the response to a threat is taken according to him or her.

During the interviews we conducted, we asked experts about the collaboration happening in SOCs. One noteworthy element is that sometimes Tier 1 and Tier 2 analysts are entirely separated and work in different locations. As face-to-face collaboration between them is not possible, this task is sometimes done with basic and not well-adapted tools. One of the experts talked about spreadsheets to exchange information between analysts, and recognized that it was not the right tool for this task. We point out that more complex tools exist to exchange information between analysts. However, they do not seem to be widespread among the experts. Only one expert spoke about TheHive[1], which is a tool dedicated to incident response.

This representation of SOCs allows to dispatch the limitations we identified and to clearly see for each analyst and for each task what are the related limitations. We see in Figure 5.1 that process limitations (number 4, 5, and 6) are impacting the work of Tier 1 analysts. Technology limitations (number 1, 2, and 3) are affecting Tier 2 analysts, and Tier 1 analysts are impacted by the quantity of data (number 1). The separation of tasks of Tier 2 analysts shows us that the technology limitations affecting them are different. The limitations related to the numerous data sources and the fact that they are not linked only concerns the analysis task. These limitations dictate our proposition of a new collaboration process between security analysts inside SOCs.

5.2 The process

In this section, we propose a new collaboration process which interacts adequately with the rules done by Tier 1 analysts with VEGAS. This process includes a feedback loop between Tier 1 analysts and Tier 2 analysts regarding the triage of security alerts.

5.2.1 Security meta-events and their rules

In VEGAS, rules created by Tier 1 analysts are used to group security events. It prevents repetitions and thus annoying work for Tier 1 analysts, these analysts now creating rules. We propose to enable cooperation with Tier 2 analysts and modify the rules structures so that they are include time characteristics. These elements make us introduce the concept of security meta-events.

[1] `https://thehive-project.org/`

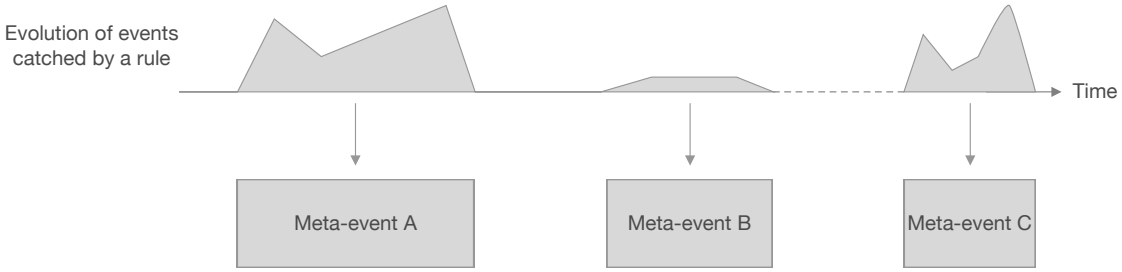

Figure 5.2: Division of security events filtered by a rule in security meta-events.

Security meta-events

Security events detected by the same rule in VEGAS could be linked to different attack scenarios over time and be unrelated. For instance, several attackers could use the same attacks or the same mechanisms to target the monitored information system. The rules being active at all time, all the matched events will be grouped even if they do not share a common underlying cause. The rule is still true; however, the matched events should be split into several groups to reflect their differences and their belongings to distinct scenarios.

This is the reason why we propose our concept of security meta-events to improve the granularity between rules quickly created by Tier 1 analysts and the events filtered by them. Rules are no more simply catching events but are also grouping events in meta-events. The repartition between meta-events is done according to the time of their arrival. If a given amount of time happens between two events caught by the rule, a new meta-event is created. A security meta-event can include one or more events and is marked in a time frame according to the first event and last event it includes. This division in meta-events allows Tier 2 analysts to inspect already triaged security events in meaningful groups, while Tier 1 analysts continue to create rules as before.

Meta-events are illustrated in Figure 5.2. On top of this figure, we can see the volume of events filtered by a rule over time. The first two peaks of events are grouped in meta-events A and B. Sometime later, new events captured by this rule arrive and result in the creation of the meta-event C. Tier 2 analysts need to analyze the three distinct meta-events instead of the whole set of events.

Thanks to this concept, an attack scenario can now be described as a single or a series of security meta-events. In contrast with meta-events, an attack scenario can contain meta-events coming from different data sources.

Rules for security meta-events

Security meta-events are defined by the same rules that are created with VEGAS. The rules are enriched with some fields. As a reminder, the fields initiated by Tier 1 analysts are:

- a *name* field, describing the rule;

- a *comment* field, used to explain more precisely the rule;
- a *filter* field, stating which security events should match.

When manipulating rules, Tier 2 analysts should be able to indicate the people in charge of this rule so that dedicated analysts can directly be notified. In order to have a feedback for their work, the Tier 1 analyst responsible for the creation of the rule should be informed in case of modification. This leads to the creation of the following fields:

- an *in charge* field, specifying the Tier 2 analyst or the team in charge of the analysis and remediation linked to this rule;
- a *creator* field, indicating the Tier 1 analyst who created the rule with VEGAS.

Rules in VEGAS are defined by Tier 1 analysts without any consideration for time. During their inspection Tier 2 analysts can add the period of time needed to split events in different meta-events and the lifespan of a rule. We believe that these features should solely be under the responsibility of Tier 2 analysts. The fields linked to these features are:

- a *start date* field, which is the date of the first captured event by default;
- an *end date* field if needed, since it is possible to disable the rule at a given date if it is known to only be useful for a specific period;
- an *lastUpdate* field to indicate the time of the last modification of the rule;
- an *interval* field that is the minimum time needed between two matched events to create a new security meta-event.

Tier 2 analysts have two distinct tasks, according to the nature of the information given by Tier 1 analysts. To take this into account we add to the rule:

- a *label* field to indicate the current state of meta-events linked to the rule.

This value of the *label* field enables Tier 2 analysts to indicate the state of the rule for the meta-events. The different values are *suspicious meta-event*, *qualified incident*, and *noise*. In our system, we differentiate between suspicious meta-events and qualified incidents, because there is no need for Tier 2 analysts to analyze again security events already inspected in the past. Suspicious meta-events are composed of security events currently happening in the system, e.g., a DDoS attack which is still occurring. A Tier 2 analyst has not looked at this meta-event, and a response has not yet been found. By contrast, after an examination by a Tier 2 analyst, the rule describing this type of meta-events can become a qualified incident or a false positive, depending on the result of the examination. The rule can be used to create future qualified incidents if the security analyst estimates that it is essential to know when new events matching this rule are detected. The rule can define noise if the events should not appear to security analysts in order to reduce the flow of irrelevant security events.

```
'name': 'command and control server',
'comment': 'IRC traffic to communicate',
'filter': {
  'sourcePort': [6667],
  'classification': ['Misc activity'],
  'sourceIP': ['10.32.5.54' ,'10.32.5.56']},
'label': 'qualified incident',
'in charge': 'John Doe',
'startDate': '2017-06-12T12:04:10.345Z',
'endDate': '2018-07-12T12:05:12.345Z',
'lastUpdate': '2018-05-01T17:45:22.541Z',
'interval': '1h'
```

Listing 5.1: An example rule for a security meta-event.

If the creation of the rule is done by Tier 1 analysts, all the fields we described can be created or changed by Tier 2 analysts if needed. This ensures that Tier 1 analysts are responsible for a better triage of alerts while being kept under supervision by Tier 2 analysts.

We give an example of a rule in listing 5.1. A Tier 1 analyst has created a rule that matches security events with a source port of `6667`, a classification of `Misc activity`, and `10.32.5.54` or `10.32.5.56` as source IP. Every security event with these characteristics will be grouped in security meta-events according to that rule. This is the rule presented in Chapter 4. We can see that a Tier 2 analyst has inspected the related events, found that they were symptomatic of a real threat, and therefore has assigned the value `qualified incident` for the label. The Tier 2 analyst has selected himself as the person in charge and he will be notified when new events matching this rule arrive. The value for the interval has been set as one hour.

5.2.2 Proposed workflow

We now present the workflow we designed that uses the concept of security meta-events and implements a feedback loop to empower Tier 1 analysts. The new workflow is illustrated in Figure 5.3. The differences with the current workflow used in SOCs (cf. Figure 5.1) are shown in bold and brown.

Tier 1 analysts are now sending suspicious meta-events instead of single events when faced with unknown suspicious security events by creating rules with VEGAS. By using this mechanism, significant time can be saved. After analyzing the meta-event, Tier 2 analysts can modify the rule if they estimate that it can be improved. Whatever the result, Tier 1 analysts keep an access to the rules. As we will see later, this will have a positive impact on their work and their feeling of empowerment.

In the beginning, there is no rule inside the system. With the regular creation and modification of rules, Tier 1 analysts see a reduction of the rate of irrelevant events so

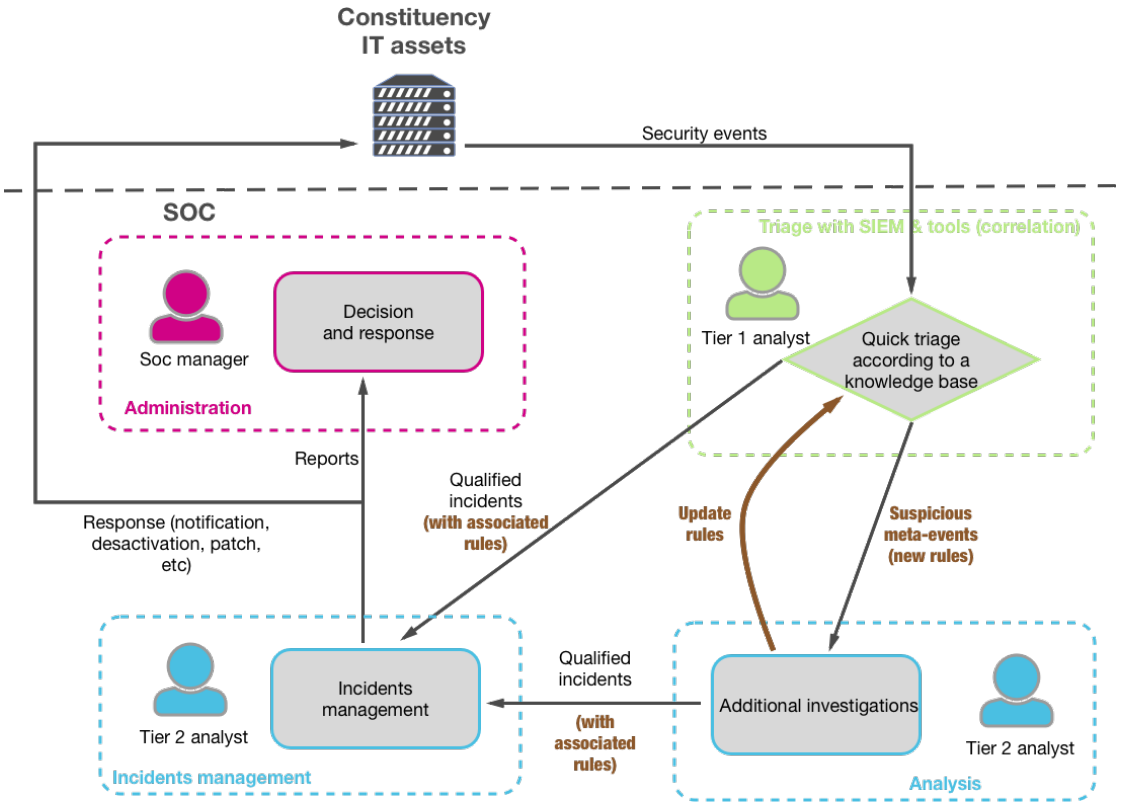

Figure 5.3: Proposed workflow for a SOC.

they can be more efficient in accomplishing their task. The addition of meta-events can also be used for case escalation and the incidents that have already been qualified.

The progression of security events coming from the IT assets in this workflow is shown in Figure 5.4. First, security events are matched against rules for meta-events defined for known noise. If matched, security events are automatically redirected to persistent storage and will not be displayed to security operators (neither Tier 1 nor Tier 2 analysts). We underline the fact that they will still be available for further analysis or forensics if needed.

Remaining security events are matched against the rules defining qualified security incidents. If a security event is part of a qualified incident, it is directly sent to the dedicated Tier 2 analysts for this incident, using the *in charge* field of the rule. An event matching this type of rule means that a previous similar incident took place and that an answer has already been found, so a dedicated Tier 2 analyst can directly be notified.

Any remaining security events are subsequently tried against by the remaining rules that describe suspicious meta-events. These events are supposedly part of an attack that is currently occurring and that Tier 2 security analysts are currently investigating.

Finally, Tier 1 analysts only receive security events that are new, i.e., that do not belong to any existing meta-event. This avoids the hassle of dealing with repetitive events so they can spend more time creating new rules corresponding to suspicious security meta-events.

This workflow facilitates the work of Tier 1 analysts while still keeping them under supervision by Tier 2 analysts thanks to the fact that Tier 2 analysts can modify and enhance the rules created by Tier 1 analysts.

5.3 Visualization for collaboration

We designed a visual tool to support this process and to exchange rules and security meta-events between security analysts. This visual tool puts into action the feedback loop between Tier 1 and Tier 2 analysts. Apart from supporting our process, this visual tool also addresses the two technical limitations described in Section 2 regarding the progression of threat escalation and the rhythm of networks. These limitations are impacting the work of Tier 2 analysts. In order to answer these limitations, we propose through TheStrip a quick perception of the context, a visual correlation of qualified incidents and security meta-events, and a visual reconstruction of attack scenarios.

5.3.1 Interface components

TheStrip primary objective being the improvement of collaboration in SOCs, the first step of a user is to identify himself or herself. The login page of TheStrip is shown in Appendix B. Three distinct types of profile are available in TheStrip according to those found in SOCs. The user can log in as the SOC manager, a Tier 1 analyst, or a Tier 2 analyst. The interaction and manipulations offered by the tool depend on the role selected by the user.

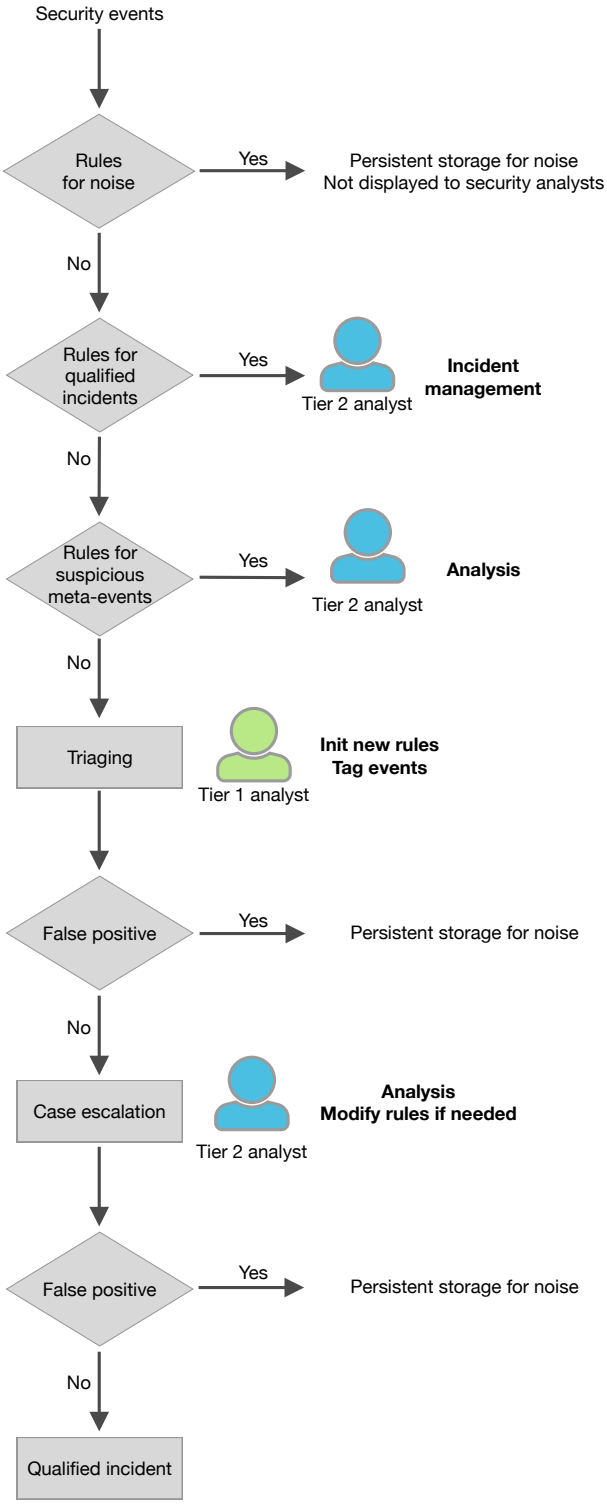

Figure 5.4: Progression of security events.

TheStrip as a visual tool is designed around three distinct views: a timeline view, a rules view, and a scenarios view. Each view is dedicated to a particular type of data with specific objectives in mind. The timeline view is the main view of the tool: its objective is to provide an immediate perception of the context. The rules view is dedicated to the manipulations and correlation of rules. The scenarios view is used to visualize security meta-events which have been linked together by Tier 2 analysts. All these views are presented in details in the remainder of this section.

When the user is logged in, he or she has access to the different views. The different views are accessible from the navigation menu on top of the application. This is shown in the zone A in Figure 5.5. The missions of Tier 1 analysts being the triage of security events, the button shown in the zone B gives a direct access to VEGAS. This way, Tier 1 analysts can do short back and forth between VEGAS and TheStrip in order to accomplish their missions of triaging, and stay updated regarding the context and the feedback.

In the top right corner of each view, in the zone C, the user can control the screen updates of the tool. Based on the experts' feedback, we identified that deterministic screen updates independent from underlying data streams were a necessity. Security events and incidents are coming at a high rate, and the analysts should not be overwhelmed with constant screen updates. So the security analysts can manually refresh the view if needed or can use an auto-refresh function with a timer. One expert said that by using this functionality of auto-refresh on the timeline view, which propose a quick representation of the context, TheStrip can be projected as a background in the SOC room to have a permanent display of the context for all analysts. The timer can be stopped or resumed if the user wants more time to focus on a specific point.

Finally, the name of the current user is shown in the zone D (here `Mary Watson`). By clicking on it, the user can close his or her session of TheStrip.

5.3.2 The timeline view

The objective for a new user after the login is to know what is the current context in regards to the security of the information system. The timeline shown in Figure 5.5 is the central view of the application and fulfill this objective. This view provides high-level awareness of what is happening on the network and enables visual correlation of incidents and security meta-events.

The timeline is divided into three different sections, according to a gradient of gravity. All events, meta-events, and qualified incidents are displayed according to the time of their arrival. The timeline view shows the period started from the oldest unclassified security events or suspicious meta-event to the current time. Axes indicating the time are used between the three sections.

On the bottom zone E, the unclassified alerts are represented on a time chart, giving an idea of the volume of events arriving from the information system. The variations in the volume of alerts could indicate changes in the type of attacks (like the start of

Figure 5.5: Timeline view.

a DDoS). These events will be grouped by Tier 1 analysts into meta-events using the VEGAS interface.

Once Tier 1 analysts have grouped events in meta-events, the meta-events that are now qualified as *suspicious meta-events* are shown in the zone F. These are the meta-events created by a specific rule which need to be inspected by Tier 2 analysts. Tier 2 analysts will now look at the meta-events, and will qualify them as *noise* or *qualified incident* according to our process. In Figure 5.5, we can see that there are currently 5 suspicious meta-events to be analyzed. Two of them have been created by the same rule according to the name (`Web server sending IRC packets`).

The qualified incidents are then shown in the zone G. Here we can see that there are two qualified incidents captured by the same rule regarding a scan of ports, and another incident regarding a tentative of brute force on the SSH server of the workstation 1.

In order to know the relation between the different meta-events and the incidents, and to understand the current context, the color of the meta-events and the incidents on the timeline are indicators of the related attack scenarios. This mechanism is primarily intended to link together qualified incidents to a specific scenario; however, suspicious meta-events can also be attached to a given scenario if needed. Grey is dedicated to the meta-events and incidents not yet linked to a specific scenario. In Figure 5.5, this is the case for the suspicious meta-events called `Data extraction` between 1 pm and 2 pm. On the contrary, the three suspicious meta-events related to a SQL injection and a web server sending IRC packets have been linked to the same scenarios by Tier 2 analysts. The display in a timeline form associated with the creation of attack scenarios enables the analyst to understand the time relation between the security events and redraw the story behind them.

The interaction offered on the timeline are thought for quick manipulations and understanding. By hovering the mouse on a meta-event or a qualified incident, information related to it (number of events, the related rule, and so forth) is displayed. This is shown in Figure 5.5 in the zone H. We can see that 3 events have been captured by this meta-event, linked to a remote file inclusion. There is a Tier 2 analyst currently in charge of the remediation of this meta-events (`John Watson`).

After the analysis of this meta-event, the Tier 2 analyst in charge found that this meta-event is linked to the two qualified incidents displayed in orange. The attackers seem to have scanned the firewall and then try to use a vulnerability to gain access to the information system. Tier 2 analysts can simply update the state of the meta-event by dragging it and dropping it in the desired sections. This interaction makes the reconstruction of attack scenario easy. The new state after the drag-and-drop of the meta-event linked to a remote file inclusion is shown in Figure 5.6.

Useful information is not always captured in security events. For instance, changes of configuration can have important consequences regarding the security of the information system and are not necessarily generating a security event transmitted to the SOC. This is also the case with physical events that are not captured by the information system but are a part of the context. To mitigate this challenge, Tier 2 analysts can also manually add an incident or a security meta-event to the timeline if needed. This is done with

Figure 5.6: Timeline view after some manipulations.

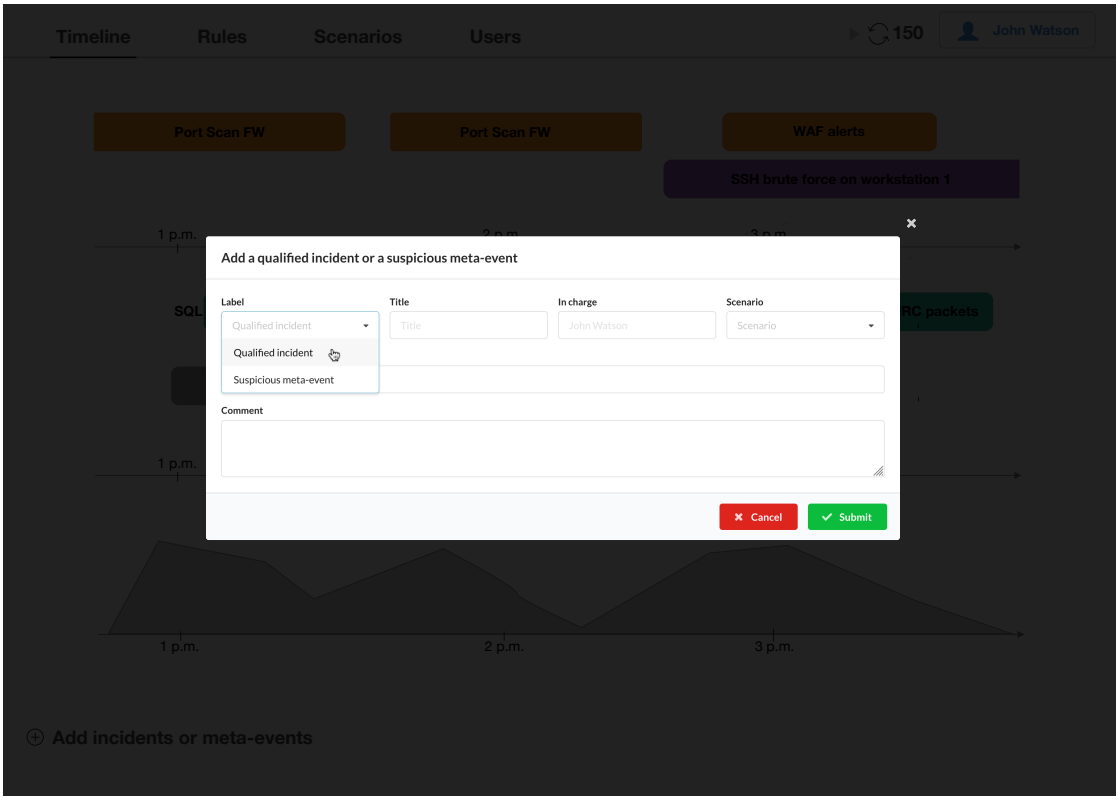

Figure 5.7: Manual addition of a qualified incidents or a suspicious meta-events.

the button `Add incidents or meta-events` at the bottom of the timeline view. A modal window [2] appears as shown in Figure 5.7. The user has to fill the different fields corresponding to the new meta-event or incident he or she wants to add.

5.3.3 The rules view

After the analysis of suspicious meta-events and the quick understanding of the context thanks to the timeline view, Tier 2 analysts may need to modify the rules defining those meta-events. Tier 2 analysts may also want to compare rules between them to find potential relations. This is the objective of the rules view, shown in Figure 5.8.

Rules defining the meta-events are displayed as lists. Tier 2 analysts have access to their main characteristics. By clicking on a rule, a modal window is displayed where the Tier 2 analyst can view all the characteristics of the rules and modify them if needed. When a rule, like the one called `Databases authentication attempt`, has an empty field regarding the person in charge, it means that no Tier 2 analyst has started to look at this rule.

[2] A modal displays content that temporarily blocks interactions with the main view.

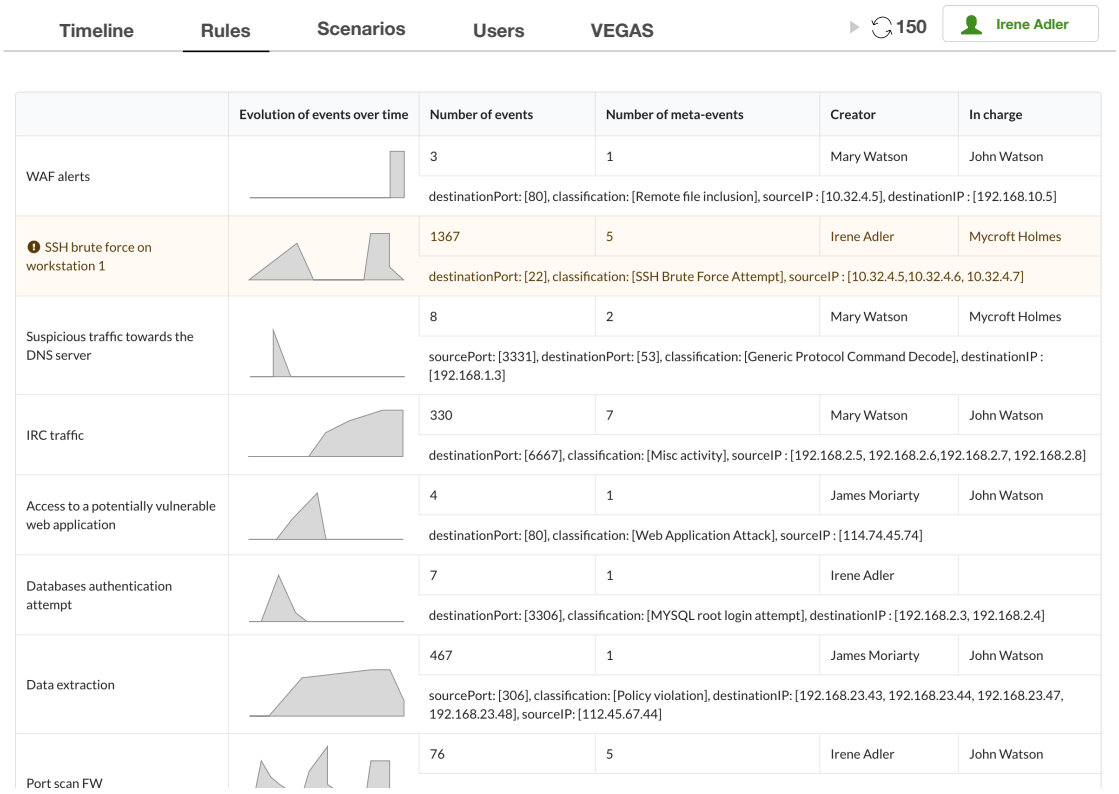

Figure 5.8: Rules view (beginning).

When looking at the rules, Tier 2 analysts need to understand the current trends. The evolution of the number of security events captured of the rules is visually displayed with sparklines used in small multiples. Sparklines are a simple and efficient way to visualize trends [56] and their superposition as a series of similar graphs with the same scale and axes enables the analysts to compare them easily.

According to the distinct tasks of Tier 1 and Tier 2 analysts, Tier 1 analysts have only a read access to the rules view. Modifications and additions of rules are not possible for them. For Tier 1 analysts, this view is where feedback is given to them as described in our process. This information is given with a notification and a different background color when a rule they are responsible for has been modified. In Figure 5.8, the rule `SSH brute force on workstation 1` has been modified. The Tier 1 analyst can look at the rule and view what the comment said so he or she can learn and improve. By viewing the modified rule, the notification disappears. With this mechanism, Tier 1 analysts have to understand the changes made by Tier 2 analysts to their rules. We advocate that this improvement help Tier 1 analyst stay motivated, accomplish their task more easily and results in improvements in the efficiency of the SOC.

5.3.4 The scenarios view

After the quick understanding of the current situations regarding the information system, the objective of Tier 2 analysts is to analyze the suspicious meta-events and reconstruct potential attacks scenarios. On the timeline, only a short period is displayed (from the oldest unclassified security events or suspicious meta-events to the current time), so scenarios which have not been recently active are not displayed. As a consequence, the scenarios view is used to show all scenarios, without any consideration for their time of apparition. The scenarios view is shown in Figure 5.9.

Scenarios are displayed as a list, with the same layout as the rules view. For each scenario, analysts have access to their characteristics: the name of the scenario, the number of rules and meta-events composing the scenario, the number of security events and the current trend regarding the events in this scenario. Tier 2 analysts can click on a scenario to modify it or delete it if needed. New scenarios can also be added using the button at the bottom of the page.

Based on the same principles as the rules view and according to the distinct tasks of Tier 1 and Tier 2 analysts, Tier 1 analysts have only a read access to the scenarios view. Modifications and additions of scenarios are not possible for them.

5.4 Implementation

TheStrip is divided as two distinct parts, a server and a client. The server side is shared with VEGAS, rules and events are directly stored in the Elasticsearch server. Elasticsearch is a highly scalable open source search engine with a REST API which has been proved efficient during the tests of VEGAS. The plugin in Python used to filter incoming alerts has been adapted to the concept of meta-events.

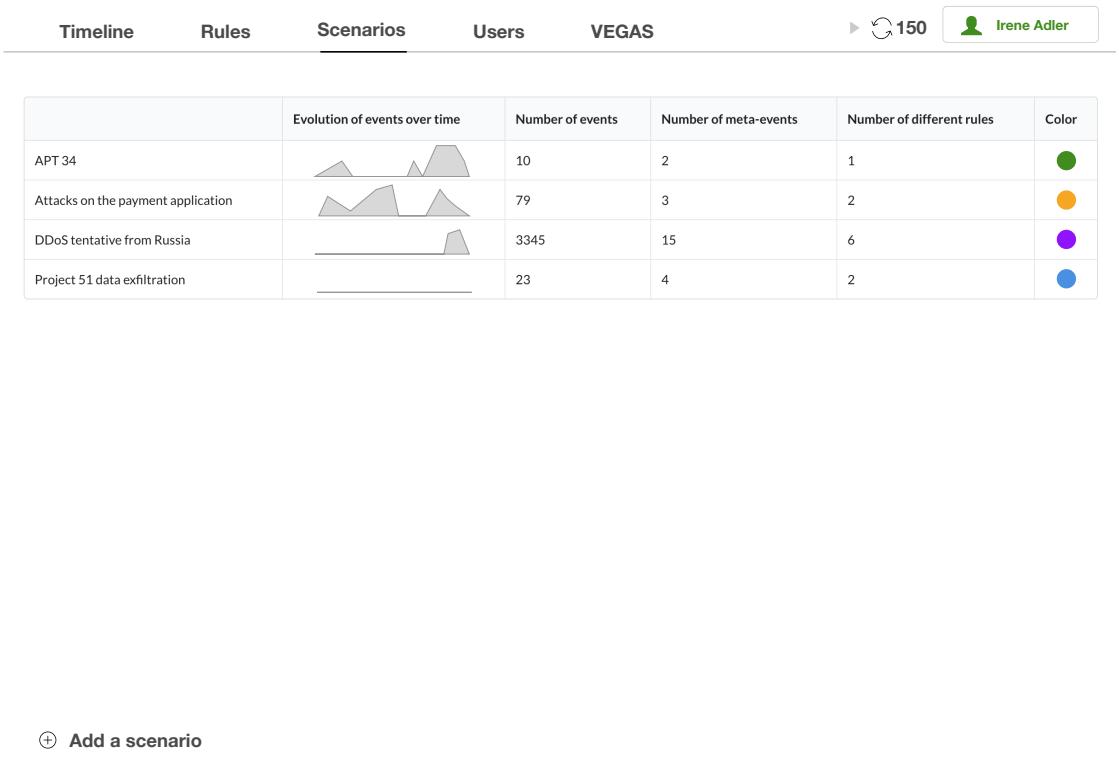

Figure 5.9: Scenarios view.

The client side of TheStrip has been implemented using web technologies, and more specifically React[3] and the D3.js library. React is used to generate the backbone of the application, while D3.js is used for the visualization of the timeline and the sparklines.

5.5 Discussion

We have identified several situations which can happen with the use of TheStrip. With the evolution of threats and attacks, the number of rules defining security meta-events in our system will increase. Tier 1 analysts are constantly adding rules to deal with the new types of attacks they are seeing and this may lead to a situation with a high number of rules to manage for Tier 2 analysts. Correlations and relations between them will be harder to find, and the reconstruction of scenarios may also be impacted. We argue that thanks to TheStrip, Tier 2 analysts can modify the rules to make them more general, for instance by merging two rules detecting the same threat. The end date is also a way to help security analysts cope with the number of alerts. Once a threat has been answered, with a patch for instance, the rule symptomatic of this attack can be removed. Nonetheless, the management of rules with TheStrip is a critical aspect.

The management of the team or the person in charge of a rule is also a critical point. The `in charge` field of a rule is used to indicate the Tier 2 analysts responsible for it. Tier 2 analysts may work in shift, and given the time of the day he or she may be not be present at the moment when the event happened. This can be a problem if the event should be quickly be analyzed to mitigate a theat. The same issue can appear if a Tier 2 analysts leave his or her job, so the events linked to this rule are without any analysts to look at it. A proper management of people has to be set up.

Another point of discussion is related to the period displayed on the timeline. Currently, the timeline shows the time after the last unclassified events or suspicious meta-events. This can lead to an imperfect representation if the Tier 1 analysts and Tier 2 are really efficient and there is nothing to been displayed, so the representation of the context is lost. On the contrary, if the Tier 1 or Tier 2 analysts feel behind, the timeline will show a long period which may become too complex for a quick understanding of the situation. A mechanism to modify the time limits of the timeline could be a response to this situation.

Finally, TheStrip has been developed to answer the limitation related to the rhythm of networks. This limitation has been partially solved with the timeline showing the last elements of context and sparklines on the other views to quickly identify the different trends. With this work, we focus on the last period of time, and the understanding of the current rhythm of network. The creation of attack scenarios leads the security analysts to better apprehend the threats happening and their dynamics. However, periodic activities over a long time period cannot be perceived. For instance, the pattern created by a malware which is active every Friday will not be detected with our work.

[3] `https://reactjs.org/`

5.6 Conclusion

In this chapter, we started by presenting the current workflow and the lack of collaboration happening in SOCs thanks to our interviews with experts. The limitations we exhibited are impacting the security analysts at different levels according to their tasks. To answer the limitations related to the lack of feedback, the rhythm of networks and the threat escalation, we have presented TheStrip, a new process and a visual tool to enhance collaboration inside SOCs.

The process is based on the concept of meta-events and rules to define them. The rules include the work done with VEGAS, and this process enables a feedback loop between Tier 1 and Tier 2 analysts.

We have designed and developed a tool to support the process and to be used by the different analysts working in SOCs. The main view proposes a timeline of the current context and enables a quick knowledge of what is happening in the information system regarding its security. Quick interaction enables threat escalations, and specific views are dedicated to the manipulation of rules and scenarios.

6 Conclusion and perspectives

Contents

6.1 Summary

This thesis makes contributions in the field of information systems security monitoring by addressing several problems met by SOCs. The initial motivation for these contributions stems from the current situation encountered by security operations centers which are the control towers for the security of information systems. Security analysts working in SOCs are exposed to a high number of irrelevant alerts, resulting in failure to react in time to real attacks. Triaging such numbers of irrelevant alerts is a challenging task in the field of security monitoring.

As discussed in Chapter 2, this situation in SOCs results from multiple limitations. We have classified these limitations into two categories. The first category is the technology, SOCs have to deal with limitations regarding lots of unlinked data and data sources, the progression of threat, and the rhythm of networks. The second category is linked to the process happening in SOCs, with the repetition of the same task, the lack of feedback, and the lack of creativity. Tier 1 analysts are particularly affected by these limitations, resulting in a high burnout rate and strenuous working conditions.

Visualization is a relevant solution in cybersecurity which brings back human in the loop. In Chapter 3, we have focused on the existing visualization techniques related to the tasks accomplished in SOCs. We have identified that visualization tools are often not in adequation with the needs of security analysts working in a SOC, showing a lack of comprehension of their work and constraints. Moreover, we have proposed the addition of collaboration between security analysts as a use for visualization in SOCs, a key aspect which has not been taken into account yet.

With VEGAS (Visualizing, Exploring and Grouping AlertS) [2, 3], we have addressed the problem regarding the number of data, the lack of creativity, and the repetition of the same task targeting Tier 1 analysts. VEGAS is a visualization and classification tool that allows Tier 1 analysts to manage the important flow of IDS alerts. We have compared different dimension reduction techniques and VEGAS uses principal component analysis to produce a two-dimensional scatterplot that can be used to visually correlate the alerts. Our tool allows Tier 1 analysts to explore the various fields of similar alerts to analyze them quickly and to generate meaningful rules. These rules will cause similar alerts to be appropriately forwarded to security analysts. This combination of data summarization and visualization is the core of our contribution. Evaluations with a case study and experts have demonstrated that VEGAS is useful to quickly detect similar IDS alerts and group them efficiently.

After addressing the problem of IDSes alerts triaging, we have tackled the remaining process limitation of the lack of feedback alongside the challenges of threat escalation and rhythm of networks. We have proposed TheStrip [4], a new process with a visual tool, to enhance the collaboration between security analysts. The process is established with the improvement of rules to define security meta-events and the creation of a specific feedback loop between Tier 1 analysts and Tier 2 analysts. Security meta-events enable this addition of time and collaborative features to the rules capturing the IDSes alerts.

To support this new process, we have designed and developed a visualization tool. This tool is organized around a timeline view, providing a quick perception of the context and easy reconstruction of attack scenarios. Several others views are available in TheStrip to help with the manipulations of meta-events and attack scenarios.

6.2 Future research directions

Visualization for cybersecurity brings highly interactive tools to identify and analyze suspicious events. Research in this field is primarily targeting monitoring, exploration, inspection and more recently forecasting. The contributions for the security analysts often propose new techniques thanks to the improvement of computing power and a better comprehension of their tasks. However, they are often lacking a way to easily communicate and report the findings to managers, shareholders, or the public. Communication and reporting could also benefit from new visualization techniques. The analysts should not only be guided to identify interesting findings, but the system should also give them the power to efficiently explain their reasoning process and the result to other people.

Thanks to techniques used for data visualization in VEGAS, Tier 1 analysts are now able to create rules. A further step is to help Tier 1 analysts build them from data by suggesting the rules and exploring them with data mining techniques. Indeed, the goal of data mining [107] is to automatically look for regularities and patterns within data. A well-investigated method since the seminal work of R. Agrawal, H. Mannila and their colleagues [108] is the discovery of association rules. Association rule learning is a rule-based machine learning method for discovering interesting relations between attributes from data. This method could extract rules already known or easily created

by Tier 1 analysts (for instance a scan incrementally checking ports) but also new and unexpected rules that can be precious to handle new attacks. However, data in SOCs mix numeric and qualitative values and arrive in streams. Even if there are methods [109] and even software [1] to discover information from streams, mining association rules in this context remains a challenging task [110].

Finally, we would like to conclude with the implications for privacy of security monitoring and more generally cybersecurity. Security analysts need novel techniques to monitor and protect the information systems and their users from malicious actors. Like most technologies, visual tools and monitoring solutions can be used with good intentions or bad motivations. Illegal mass surveillance of people is an example of the abuses which are now possible due to novels technologies. We believe that more research is needed regarding how to improve cybersecurity without compromising privacy, how protection of computers networks can be balanced with the right of each individual.

[1] `https://moa.cms.waikato.ac.nz/`

Appendices

A VAST 2012 network

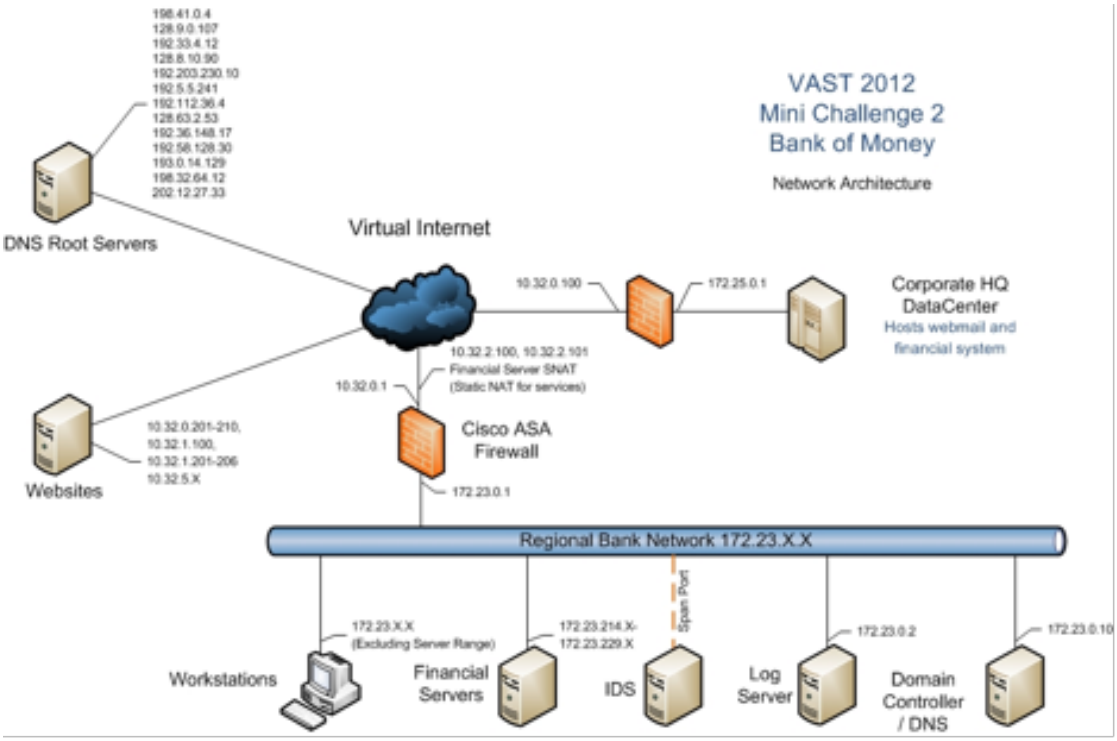

Figure A.1: Bank of Money Regional Headquarters Network.

B Login page of TheStrip

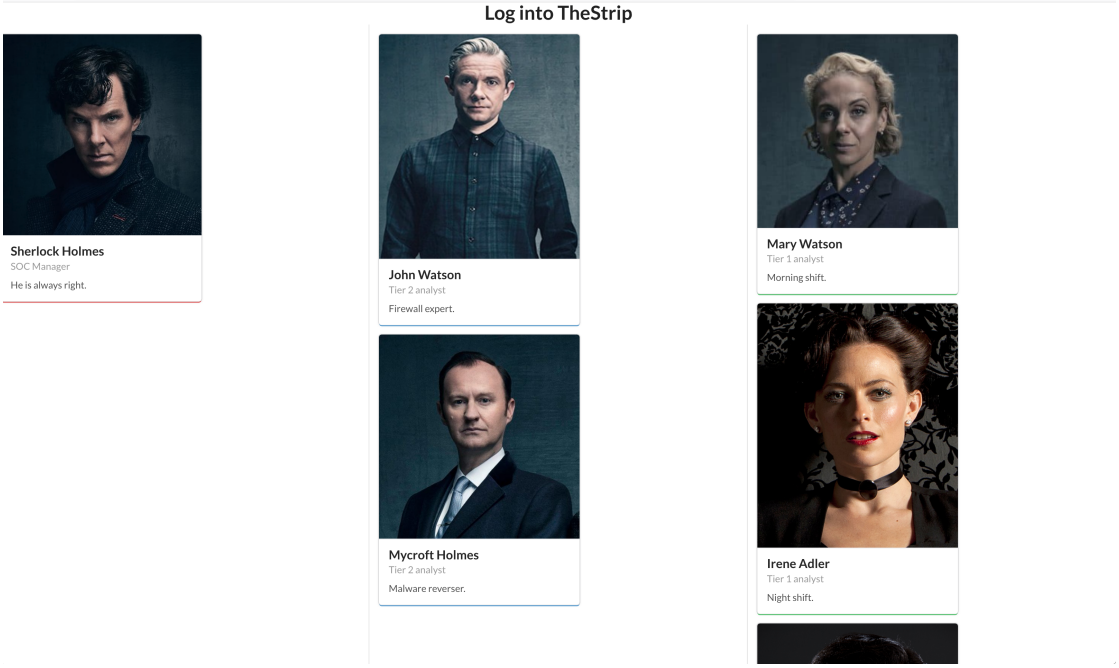

Figure B.1: Login page of TheStrip.

Glossary

APT Advanced Persistent Threat. An advanced persistent threat is a broad term used to describe a set of continuous computer hacking processes, often orchestrated by a intruder or team of intruders targeting a specific entity. 17

AS Autonomous System. An autonomous system is a network or a collection of network controlled and supervised by a single entity. 36

C&C Command and Control. 58

DDoS Distributed Denial-of-Service. 50, 77

DNS Domain Name System. 45, 58

IDS Intrusion Detection System. iii, v, vi, 2, 3, 6, 9, 16, 17, 21, 23, 27, 28, 30, 37, 42–44, 47, 53, 56, 61, 63, 86

information system an organized set of resources (hardware, software, personnel, data and procedures) for processing and communicating information. iii, 2, 3, 8, 9, 19, 22, 27, 30, 37, 38, 44, 66, 69, 75, 77, 81, 85

IOC Indicator of compromise. 7

IoT Internet of Things. 1

IRC Internet Relay Chat. xvi, 52, 53, 58, 59

MDS Multidimensional Scaling. 45

PCA Principal Component Analysis. 32, 42, 45–47, 50, 52, 53, 56, 60

security event An identified and observable occurrence of a system, service, process or network state that may be security relevant. 2, 3, 65, 66, 68–71

security incident A single or a series of unwanted or unexpected information events that have a significant probability of threatening information security. iii, 1, 2, 5, 9, 23, 65

SIEM Security Information and Event Management. 6, 8, 66

situational awareness "the perception of the elements in the environment within a volume of time and space, the comprehension of their meaning, and the projection of their status in the near future" [58]. vi, 3, 4, 22–24, 26, 38

SLA Service-Level Agreement. A SLA is a commitment between two or more parties, usually a service provider and a client. Particular aspects of the service (quality, availability, responsibilities) are agreed between the parties. 13

SOC Security Operations Center. iii, v–ix, xv, xvi, 2–19, 21–24, 37, 39, 41, 45, 59, 65–68, 73, 75, 77, 84, 85, 87

Bibliography

[1] Mathieu Vidard. *La cybersécurité*. Oct. 8, 2018 (cit. on p. v).

[2] Damien Crémilleux et al. "VEGAS: Visualizing, Exploring and Grouping Alerts". In: NOMS 2016 - 2016 IEEE/IFIP Network Operations and Management Symposium. Apr. 2016 (cit. on pp. v, viii, 3, 4, 41, 86).

[3] D. Crémilleux et al. "VEGAS: Visualizing, Exploring and Grouping Alerts (Poster)". In: 2016 IEEE Symposium on Visualization for Cyber Security (VizSec). Chicago, IL, USA, Oct. 2016 (cit. on pp. v, viii, 3, 4, 41, 86).

[4] Damien Crémilleux et al. "Enhancing Collaboration between Security Analysts in Security Operations Centers". In: Conference on Risks and Security of Internet and Systems. Lecture Notes in Computer Science. Arcachon, France: Springer, Oct. 17, 2018 (cit. on pp. v, ix, 3, 4, 65, 86).

[5] *VAST Challenge 2012*. 2012 (cit. on pp. viii, 50, 56).

[6] Robert S. Mueller. "Combating Threats in the Cyber World: Outsmarting Terrorists, Hackers, and Spies". Mar. 1, 2012 (cit. on p. 1).

[7] Karthik Selvaraj et al. *WannaCrypt Ransomware Worm Targets Out-of-Date Systems*. May 12, 2017 (cit. on p. 1).

[8] Brian Fung. "Equifax's Massive 2017 Data Breach Keeps Getting Worse". In: *Washington Post* (Mar. 1, 2018) (cit. on p. 1).

[9] ANSSI. *Prestataires de détection des incidents de sécurité*. Référentiel d'exigences. ANSSI, Dec. 21, 2017 (cit. on pp. 1, 7, 9, 11, 17).

[10] Carson Zimmerman. *Ten Strategies of a World-Class Cybersecurity Operations Center*. The MITRE Corporation, Oct. 2014 (cit. on pp. 2, 6–8, 13).

[11] Risto Vaarandi. "Real-Time Classification of IDS Alerts with Data Mining Techniques". In: IEEE Military Communications Conference, 2009. MILCOM 2009. Boston, MA, USA: IEEE, Oct. 2009 (cit. on p. 2).

[12] Klaus Julisch and Marc Dacier. "Mining Intrusion Detection Alarms for Actionable Knowledge". In: KDD '02. New York, NY, USA: ACM, 2002 (cit. on p. 2).

[13] Hugo Gascon, Agustin Orfila, and Jorge Blasco. "Analysis of Update Delays in Signature-Based Network Intrusion Detection Systems". In: *Computers & Security* 30.8 (Nov. 1, 2011) (cit. on p. 2).

[14] *2017 Cost of Data Breach Study*. IBM Security, June 2017 (cit. on p. 2).

[15] Matt Burgess. "Massive Yahoo Database Reportedly Sold for £240,000 on the Dark Web". In: *Wired UK* (Dec. 15, 2016) (cit. on p. 2).

[16] Sathya Chandran Sundaramurthy et al. "A Human Capital Model for Mitigating Security Analyst Burnout". In: Symposium on Usable Privacy and Security. SOUPS '15. Ottawa, Ontario, Canada: USENIX Association, July 2015 (cit. on pp. 3, 5, 18).

[17] A. D'Amico and K. Whitley. "The Real Work of Computer Network Defense Analysts". In: ed. by John R. Goodall, Gregory Conti, and Kwan-Liu Ma. Mathematics and Visualization. Springer Berlin Heidelberg, 2008 (cit. on p. 3).

[18] Narges Mahyar and Melanie Tory. "Supporting Communication and Coordination in Collaborative Sensemaking". In: *IEEE Transactions on Visualization and Computer Graphics* (2013) (cit. on p. 3).

[19] Eric Cole. *SOC Automation-Deliverance or Disaster*. SANS Institute, Nov. 2017 (cit. on p. 5).

[20] Hewlett-Packard. *Security Operations Building a Successful SOC*. Business white paper. Hewlett-Packard, Nov. 2015 (cit. on pp. 5, 7).

[21] Hewlett-Packard. *5G/SOC: SOC Generations*. White Paper. Hewlett-Packard, May 2013 (cit. on pp. 6, 7).

[22] Kevin D. Mitnick and William L. Simon. *The Art of Deception: Controlling the Human Element of Security*. Indianapolis, Ind: Wiley, 2002. 352 pp. (cit. on p. 6).

[23] CERT Division. *1999 CERT Incident Notes*. Carnegie Mellon University, 1999 (cit. on p. 6).

[24] *US-CERT | United States Computer Emergency Readiness Team*. 2018 (cit. on p. 6).

[25] *SB-1386 Personal Information: Privacy*. 2012 (cit. on p. 6).

[26] Michael Connell and Sarah Vogler. *Russia's Approach to Cyber Warfare*. Center for Naval Analyses Arlington United States, Mar. 2017 (cit. on p. 6).

[27] Ralph Langner. *To Kill a Centrifuge*. The Langner Group, Nov. 2013 (cit. on p. 6).

[28] McAfee. *Creating and Maintaining a SOC*. White Paper. McAfee, 2013 (cit. on p. 7).

[29] Fortinet. *Rock the SOC 101*. White Paper. Fortinet, Dec. 23, 2016 (cit. on p. 7).

[30] Marcel Hoffmann. "How to Build a Successful SOC". Sept. 2014 (cit. on pp. 7, 9).

[31] Splunk. *Building a SOC with Splunk*. Tech brief. Splunk, 2017 (cit. on p. 7).

[32] IBM. *Strategy Considerations for Building a Security Operations Center*. IBM, 2013 (cit. on p. 7).

[33] Joseph Muniz, Gary McIntyre, and Nadhem Al-Fardan. *Security Operations Center: Building, Operating, and Maintaining Your SOC*. Indianapolis, Indiana: Cisco Press, Oct. 29, 2015. 424 pp. (cit. on pp. 7, 9, 10).

[34] EY. *Security Operations Centers against Cybercrime*. EY, Oct. 2013 (cit. on p. 7).

[35] Deloitte. *A New Approach to Cyber Security*. Deloitte, June 2017 (cit. on p. 7).

[36] Tirath Singh. *Building Your Security Operations Center and Taking It to the Next Level*. White Paper. Tata Consultancy Services, 2016 (cit. on p. 7).

[37] David Nathans. *Designing and Building a Security Operations Center*. Waltham, Massachusetts: Syngress, 2015 (cit. on p. 7).

[38] Diana Kelley and Ron Moritz. "Best Practices for Building a Security Operations Center". In: *Information Systems Security* 14.6 (Jan. 1, 2006) (cit. on p. 7).

[39] Saâd Kadhi. "CERT, CSIRT et SOC en pratique : comment s'organiser et quels outils mettre en place". In: *MISC* 94 (Nov. 2017) (cit. on p. 7).

[40] Alissa Torres. *Building a World-Class Security Operations Center: A Roadmap*. Whitepaper. SANS Institute, 2015 (cit. on p. 7).

[41] CLUSIF. *Comment réussir le déploiement d'un SOC*. Mar. 2017 (cit. on pp. 7, 13).

[42] Paul Cichonski et al. *Computer Security Incident Handling Guide : Recommendations of the National Institute of Standards and Technology*. NIST SP 800-61r2. National Institute of Standards and Technology, Aug. 2012 (cit. on p. 7).

[43] Pierre Jacobs, Alapan Arnab, and Barry Irwin. "Classification of Security Operation Centers". In: 2013 Information Security for South Africa. Aug. 2013 (cit. on p. 7).

[44] Moira J. West-Brow et al. *Handbook for Computer Security Incident Response Teams (CSIRTs)*. Pittsburgh, Pennsylvania, USA: Carnegie Mellon University, Apr. 2003 (cit. on p. 8).

[45] MWR InfoSecurity. *Why Are SOCs Failing?* White Paper. 2016 (cit. on p. 16).

[46] McAfee. *McAfee Labs Threats Report*. McAfee, Dec. 2016 (cit. on p. 16).

[47] Hervé Debar, Marc Dacier, and Andreas Wespi. "A Revised Taxonomy for Intrusion-Detection Systems". In: *Annales Des Télécommunications* 55.7-8 (July 1, 2000) (cit. on p. 16).

[48] Daniel M. Best, Alex Endert, and Daniel Kidwell. "7 Key Challenges for Visualization in Cyber Network Defense". In: VizSec '14. New York, NY, USA: ACM, 2014 (cit. on pp. 17, 19).

[49] Sathya Chandran Sundaramurthy et al. "A Tale of Three Security Operation Centers". In: SIW '14. New York, NY, USA: ACM, 2014 (cit. on p. 17).

[50] Prashanth Rajivan and Nancy Cooke. "Impact of Team Collaboration on Cybersecurity Situational Awareness". In: Lecture Notes in Computer Science. Springer, Cham, 2017 (cit. on pp. 18, 19).

[51] Kelly Jackson Higgins. *Death of the Tier 1 SOC Analyst*. Nov. 16, 2017 (cit. on p. 19).

[52] Sathya Chandran Sundaramurthy. "An Anthropological Study of Security Operations Centers to Improve Operational Efficiency". PhD thesis. Florida, USA: University of South Florida, June 7, 2017 (cit. on p. 19).

[53] Colin Ware. *Information Visualization: Perception for Design*. Third edition. Interactive technologies. Waltham, MA: Morgan Kaufmann, 2013. 512 pp. (cit. on p. 21).

[54] J.R. Goodall. "Visualization Is Better! A Comparative Evaluation". In: 6th International Workshop on Visualization for Cyber Security, 2009. VizSec 2009. Oct. 2009 (cit. on p. 21).

[55] Markus Wagner et al. "A Survey of Visualization Systems for Malware Analysis". In: 2015 (cit. on p. 22).

[56] Edward R. Tufte. *The Visual Display of Quantitative Information*. 2nd ed. Cheshire, Conn: Graphics Press, 2001. 197 pp. (cit. on pp. 22, 81).

[57] L. Harrison et al. "Foreword of the 2015 IEEE Symposium on Visualization for Cyber Security". In: 2015 IEEE Symposium on Visualization for Cyber Security (VizSec). Oct. 2015 (cit. on p. 22).

[58] Mica R. Endsley and Debra G. Jones. *Designing for Situation Awareness: An Approach to User-Centered Design, Second Edition*. CRC Press, Nov. 16, 2011 (cit. on pp. 23, 96).

[59] Ulrik Franke and Joel Brynielsson. "Cyber Situational Awareness – A Systematic Review of the Literature". In: *Computers & Security* 46 (Oct. 2014) (cit. on pp. 23, 27).

[60] Anita D'Amico and Michael Kocka. "Information Assurance Visualizations for Specific Stages of Situational Awareness and Intended Uses: Lessons Learned". In: vol. 0. Los Alamitos, CA, USA: IEEE Computer Society, 2005 (cit. on pp. 24, 25).

[61] Raffael Marty. *Applied Security Visualization*. Upper Saddle River, NJ: Addison-Wesley, 2009. 523 pp. (cit. on p. 24).

[62] Christopher Humphries. "User-Centred Security Event Visualisation". PhD thesis. Université Rennes 1, Dec. 8, 2015 (cit. on p. 24).

[63] Hadi Shiravi, Ali Shiravi, and Ali A. Ghorbani. "A Survey of Visualization Systems for Network Security". In: *IEEE Transactions on Visualization and Computer Graphics* 18.8 (Aug. 2012) (cit. on p. 26).

[64] Fabian Fischer. "Visual Analytics for Situational Awareness in Cyber Security". PhD thesis. Konstanz: Universität Konstanz, Apr. 21, 2016 (cit. on p. 26).

[65] D. A. Keim. "Information Visualization and Visual Data Mining". In: *IEEE Transactions on Visualization and Computer Graphics* 8.1 (Jan. 2002) (cit. on p. 26).

[66] Stephen Few. *Information Dashboard Design: The Effective Visual Communication of Data*. 1st ed. Beijing ; Cambride [MA]: O'Reilly, 2006. 211 pp. (cit. on p. 26).

[67] Hideki Koike and Kazuhiro Ohno. "SnortView: Visualization System of Snort Logs". In: VizSEC/DMSEC '04. New York, NY, USA: ACM, 2004 (cit. on pp. 27, 28).

[68] K. Abdullah et al. "IDS rainStorm: Visualizing IDS Alarms". In: IEEE Workshop on Visualization for Computer Security, 2005. (VizSEC 05). Oct. 2005 (cit. on p. 27).

[69] Kiran Lakkaraju, William Yurcik, and Adam J. Lee. "NVisionIP: Netflow Visualizations of System State for Security Situational Awareness". In: VizSEC/DMSEC '04. New York, NY, USA: ACM, 2004 (cit. on p. 27).

[70] H. Koike, K. Ohno, and K. Koizumi. "Visualizing Cyber Attacks Using IP Matrix". In: IEEE Workshop on Visualization for Computer Security, 2005. (VizSEC 05). Oct. 2005 (cit. on p. 27).

[71] Jamie Rasmussen et al. "Nimble Cybersecurity Incident Management Through Visualization and Defensible Recommendations". In: VizSec '10. New York, NY, USA: ACM, 2010 (cit. on pp. 27, 29).

[72] Hadi Shiravi, Ali Shiravi, and Ali A. Ghorbani. "IDS Alert Visualization and Monitoring through Heuristic Host Selection". In: ed. by Miguel Soriano, Sihan Qing, and Javier López. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2010 (cit. on p. 28).

[73] Ying Zhao et al. "IDSRadar: A Real-Time Visualization Framework for IDS Alerts". In: *Science China Information Sciences* 56.8 (Aug. 1, 2013) (cit. on p. 28).

[74] S. Foresti and J. Agutter. "VisAlert: From Idea to Product". In: ed. by John R. Goodall, Gregory Conti, and Kwan-Liu Ma. Mathematics and Visualization. Springer Berlin Heidelberg, 2008 (cit. on pp. 28, 29).

[75] Ben Shneiderman. "Tree Visualization with Tree-Maps: 2-d Space-Filling Approach". In: *ACM Trans. Graph.* 11.1 (Jan. 1992) (cit. on p. 30).

[76] Fabian Fischer, Johannes Fuchs, and Florian Mansmann. "ClockMap: Enhancing Circular Treemaps with Temporal Glyphs for Time-Series Data". In: Eurographics Conference on Visualization (EuroVis 2012). May 7, 2012 (cit. on p. 30).

[77] Fabian Fischer et al. "BANKSAFE: Visual Analytics for Big Data in Large-Scale Computer Networks". In: *Information Visualization* 14.1 (Jan. 1, 2015) (cit. on p. 30).

[78] J. Ortiz-Ubarri et al. "Toa: A Web Based Network Flow Data Monitoring System at Scale". In: 2015 IEEE International Congress on Big Data. June 2015 (cit. on p. 31).

[79] Daisuke Inoue et al. "DAEDALUS-VIZ: Novel Real-Time 3D Visualization for Darknet Monitoring-Based Alert System". In: VizSec '12. New York, NY, USA: ACM, 2012 (cit. on p. 31).

[80] Peter Curtis et al. “A Tool for Rapid Visual Interrogation & Triage of Alerts”. In: 2014 (cit. on p. 32).

[81] Alberto Boschetti et al. “TVi: A Visual Querying System for Network Monitoring and Anomaly Detection”. In: VizSec ’11. New York, NY, USA: ACM, 2011 (cit. on p. 32).

[82] N. Anh Huynh et al. “Uncovering Periodic Network Signals of Cyber Attacks”. In: 2016 IEEE Symposium on Visualization for Cyber Security (VizSec). Oct. 2016 (cit. on pp. 32, 33).

[83] Christopher Humphries et al. “CORGI: Combination, Organization and Reconstruction Through Graphical Interactions”. In: VizSec ’14. New York, NY, USA: ACM, 2014 (cit. on pp. 33, 34, 50).

[84] Simon Walton, Eamonn Maguire, and Min Chen. “Multiple Queries with Conditional Attributes (QCATs) for Anomaly Detection and Visualization”. In: VizSec ’14. New York, NY, USA: ACM, 2014 (cit. on pp. 33, 34).

[85] Orestis Tsigkas, Olivier Thonnard, and Dimitrios Tzovaras. “Visual Spam Campaigns Analysis Using Abstract Graphs Representation”. In: VizSec ’12. New York, NY, USA: ACM, 2012 (cit. on p. 35).

[86] Timothy R. Leschke and Alan T. Sherman. “Change-Link: A Digital Forensic Tool for Visualizing Changes to Directory Trees”. In: VizSec ’12. New York, NY, USA: ACM, 2012 (cit. on p. 35).

[87] H. Kim et al. “Firewall Ruleset Visualization Analysis Tool Based on Segmentation”. In: 2017 IEEE Symposium on Visualization for Cyber Security (VizSec). Oct. 2017 (cit. on p. 35).

[88] Tamara Yu et al. “EMBER: A Global Perspective on Extreme Malicious Behavior”. In: VizSec ’10. New York, NY, USA: ACM, 2010 (cit. on p. 36).

[89] Francesco Roveta et al. “BURN: Baring Unknown Rogue Networks”. In: VizSec ’11. New York, NY, USA: ACM, 2011 (cit. on p. 36).

[90] Florian Mansmann, Timo Göbel, and William Cheswick. “Visual Analysis of Complex Firewall Configurations”. In: VizSec ’12. New York, NY, USA: ACM, 2012 (cit. on p. 37).

[91] J. R. Goodall and M. Sowul. “VIAssist: Visual Analytics for Cyber Defense”. In: 2009 IEEE Conference on Technologies for Homeland Security. May 2009 (cit. on pp. 37, 38, 59).

[92] Christopher P. Lee and John A. Copeland. “Flowtag: A Collaborative Attack-Analysis, Reporting, and Sharing Tool for Security Researchers”. In: VizSEC ’06. New York, NY, USA: ACM, 2006 (cit. on p. 37).

[93] Siming Chen et al. “OCEANS: Online Collaborative Explorative Analysis on Network Security”. In: VizSec ’14. New York, NY, USA: ACM, 2014 (cit. on pp. 37, 39).

[94] Diane Staheli et al. “Collaborative Data Analysis and Discovery for Cyber Security”. In: Symposium on Usable Privacy and Security. Denver, CO: USENIX Association, 2016 (cit. on p. 38).

[95] Hervé Debar and Benjamin S. Feinstein. *The Intrusion Detection Message Exchange Format (IDMEF)*. RFC 4765. IETF, Mar. 2007 (cit. on p. 42).

[96] H. Hotelling. “Analysis of a Complex of Statistical Variables into Principal Components.” In: *Journal of Educational Psychology* 24.6 (1933) (cit. on p. 45).

[97] I. Borg and P. J. F. Groenen. *Modern Multidimensional Scaling: Theory and Applications*. 2nd ed. Springer Series in Statistics. New York: Springer-Verlag, 2005 (cit. on p. 45).

[98] Joshua B. Tenenbaum, Vin de Silva, and John C. Langford. “A Global Geometric Framework for Nonlinear Dimensionality Reduction”. In: *Science* 290.5500 (Dec. 22, 2000). pmid: 11125149 (cit. on p. 45).

[99] Sam T. Roweis and Lawrence K. Saul. “Nonlinear Dimensionality Reduction by Locally Linear Embedding”. In: *Science* 290.5500 (Dec. 22, 2000). pmid: 11125150 (cit. on p. 45).

[100] Mikhail Belkin and Partha Niyogi. “Laplacian Eigenmaps and Spectral Techniques for Embedding and Clustering”. In: NIPS’01. Cambridge, MA, USA: MIT Press, 2001 (cit. on p. 45).

[101] Christopher Humphries et al. “ELVIS: Extensible Log VISualization”. In: VizSec ’13. New York, NY, USA: ACM, 2013 (cit. on p. 46).

[102] Susan Garavaglia and Asha Sharma. “A Smart Guide to Dummy Variables: Four Applications and a Macro”. In: 1998 (cit. on p. 46).

[103] B. Shneiderman. “The Eyes Have It: A Task by Data Type Taxonomy for Information Visualizations”. In: , IEEE Symposium on Visual Languages, 1996. Proceedings. Sept. 1996 (cit. on p. 47).

[104] Heng Du. *Windows TCP 139 and 445 Vulnerability*. Feb. 10, 2010 (cit. on p. 58).

[105] Diane Staheli et al. “Visualization Evaluation for Cyber Security: Trends and Future Directions”. In: VizSec ’14. New York, NY, USA: ACM, 2014 (cit. on p. 60).

[106] Marco Angelini, Nicolas Prigent, and Giuseppe Santucci. “PERCIVAL: Proactive and Reactive Attack and Response Assessment for Cyber Incidents Using Visual Analytics”. In: 2015 IEEE Symposium on Visualization for Cyber Security (VizSec). Oct. 2015 (cit. on p. 60).

[107] Usama M. Fayyad et al., eds. *Advances in Knowledge Discovery and Data Mining*. AAAI/MIT Press, 1996 (cit. on p. 86).

[108] Rakesh Agrawal et al. “Fast Discovery of Association Rules”. In: AAAI/MIT Press, 1996 (cit. on p. 86).

[109] Graham Cormode and Marios Hadjieleftheriou. “Methods for Finding Frequent Items in Data Streams”. In: *VLDB J.* 19.1 (2010) (cit. on p. 87).

[110] Nan Jiang and Le Gruenwald. “Research Issues in Data Stream Association Rule Mining”. In: *SIGMOD Record* 35.1 (2006) (cit. on p. 87).

UNIVERSITE BRETAGNE LOIRE / MATHSTIC

CentraleSupélec

Titre : Visualisation pour la supervision de sécurité des systèmes d'information

Mot clefs : sécurité informatique, système de détection d'intrusion, SOC, collaboration, visualisation

Resumé : Le centre opérationnel de sécurité, SOC, est un élément central pour la sécurité des systèmes d'information. Dans cette thèse, nous nous intéressons à ses limites et proposons un nouveau processus et deux outils visuels pour y répondre. Nos contributions permettent à la fois une meilleure collaboration entre les analystes travaillant au sein des SOCs, ainsi que de faciliter visuellement le triage des événements de sécurité au sein des systèmes d'informations.

Title: Visualization for information system security monitoring

Keywords: cybersecurity, intrusion detection systems, SOC, collaboration, visualization

Abstract: A security operations center, SOC, is a key element for the security of information systems. In this thesis, we exhibited the limitations of SOCs and proposed a process associated with two tools to answer them. Our contributions enable a better collaboration between the security analysts working in SOCs and facilitate security events triage thanks to visualization.