Authors: Reinder Wolthuis, senior consultant/project manager cybersecurity at TNO. Can be reached at reinder.wolthuis@tno.nl. Frank Fransen, senior scientist cybersecurity at TNO. Can be reached at frank.fransen@tno.nl.

SOCCRATES - Vision & Roadmap for SOC & CSIRTs

SOCCRATES (SOC & CSIRT Response to Attacks & Threats, based on attack defence graphs Evaluation Systems) is a European innovation project, co-funded by the Horizon2020 program and led by TNO. It brings together some of the best European expertise in the field to develop, implement and evaluate an automated security platform to support SOC analysts. This third article on the project provides a summary of the 'vision, roadmap and guidance for SOC' booklet that was recently published by SOCCRATES.

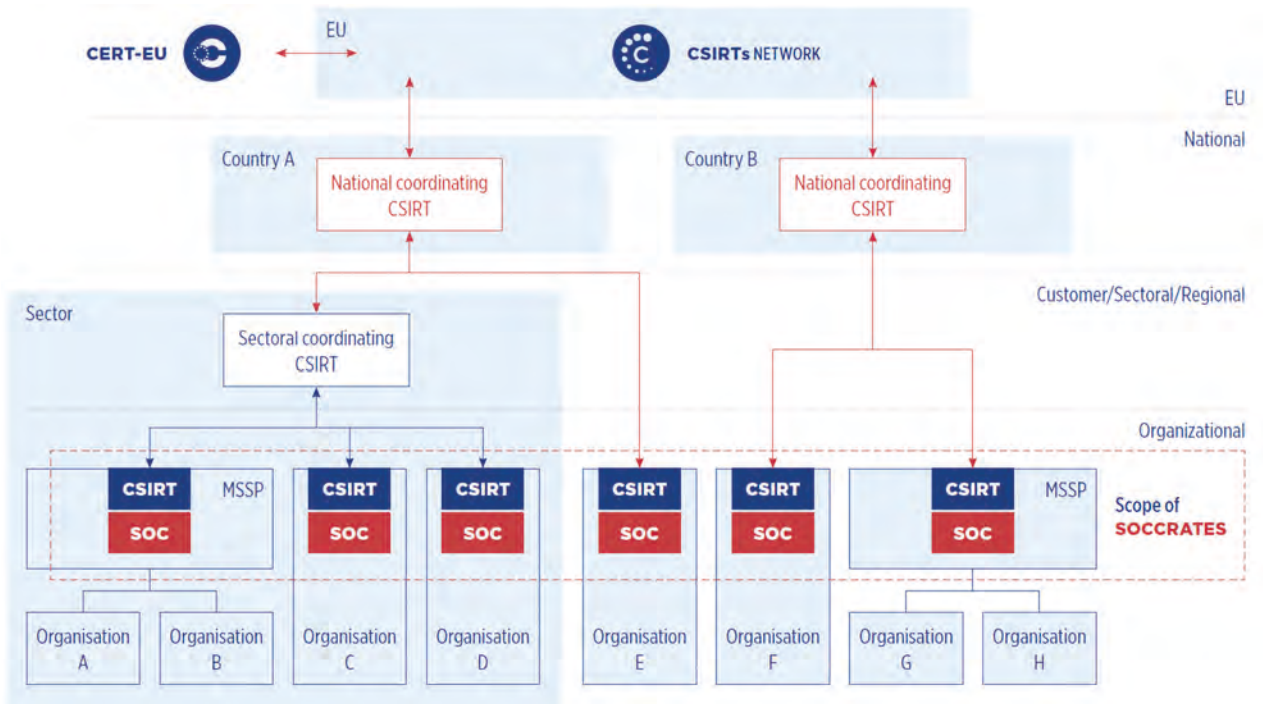

Figure 1 - SOC and CSIRT clustering and layering.

The SOCCRATES project was introduced in two previous articles (iB-Magazine 4 and iB-Magazine 5 2021). The first article gave an overview of the challenges that Security Operations Centers (SOCs) and Computer Security Incident Response Teams (CSIRTs) face, and how the SOCCRATES project addresses these challenges by developing a security automation and decision support platform, 'the SOCCRATES platform'. The second article described in more detail how the SOCCRATES platform is providing security automation for SOC and CSIRT processes. How it provides situational awareness and option awareness to the SOC analyst and enables (semi) automated response execution. This third article elaborates on the SOC and CSIRT capabilities, and the vision of the SOCCRATES project on the future needs for SOCs and CSIRTs.

SOC & CSIRT

The increasing dependency of organisations, and society as a whole, on IT systems and networks as well as the increase of cyber security incidents with major impact, has led to organisations (and governments) increasing their spending on cyber security. Many organisations have established a Security Operations Centre (SOC) and Computer Security Incident Response Team (CSIRT) to protect the organisation against cyber-attacks, or they contracted a Managed Security Service Provider (MSSP) to perform these operational cyber security services for them. Both a SOC and a CSIRT are thus expert teams (often also formally embedded in an organisational unit), that provide operational security services. It is quite common to use alternative names for similar types of such organisational units, such as Cyber Defence Centre (CDC) for SOC, and Computer Emergency Response Team (CERT) for CSIRT. Moreover, the terms SOC and CSIRT are also applied at different levels, as can be seen in figure 1.

In the lowest layer of the picture we see the organisations that have their own SOC and organisations that make use of SOC and/or CSIRT services provided by commercial MSSPs. These actually are the core focus of the SOCCRATES project, i.e. SOCCRATES enhances the SOC/CSIRT capabilities of SOCs and CSIRTs that are run by organisations and MSSPs.

In the higher layers, we mainly see so called coordinating CSIRTs that provide services to a set of organizations known as the constituency (e.g. organisations in a sector, region, country, etc.). These services include typically coordination of security incidents that affect several organisations within the constituency, acting as a single point of contact for the sector/region/country, distribution of cyber threat intelligence and providing security incident analysis and forensic services.

Note that coordinating CSIRTs sometimes operate at a same level

	SOC	CSIRT
Current Capabilities	• Monitoring and detection • Event analysis • Information security incidents analysis (triage) • Vulnerability analysis • Security Awareness Creation	• Information security incidents analysis (incl. event correlation) • Threat analysis (incl. collection, sharing and processing of CTI) • Artefact and forensic evidence analysis • Information security incident coordination • Vulnerability coordination • Awareness Building
	• Event detection (through alerting and/or hunting) / threat hunting • SOC tool life-cycle support (incl. operations and maintenance, tuning tools (e.g. detection sensors), engineering and deployment, and R&D)	
New Capabilities	• Security Orchestration, Automation and Response (SOAR) • Automated security reasoning (real-time threat & impact assessment) • Automated generation and assessment of response actions (both pro-active & reactive) • Automated response execution	

Figure 2 - Current & needed SOC / CSIRT capabilities.

as the Information Sharing and Analysis Centres (ISACs), not shown in the figure, which facilitate gathering and sharing of information on cyber threats while CSIRT activities may go beyond the ISAC activities, e.g. through security incident coordination.
In figure 1 the following coordinating CSIRTs are depicted:

- Sectoral or Regional CSIRTs - Dedicated CSIRTs that collect and analyse threat intel, translate this to the specific context of the sector or region and distribute it within the sector/region organizations. Examples of sectoral CSIRTs in the Dutch context are Z-Cert (healthcare) and IBD (Informatie Beveiligings Dienst, municipalities).
- National CSIRTS – These CSIRTs have the task to enhance a nation's resilience in the digital domain, prevent or limit the failure of the availability or the loss of integrity of information systems of vital operators and central government, and to handle severe computer attacks against critical infrastructure and information within the nation. On Dutch national level we of course have the Dutch National Cyber Security Centre (NCSC).
- CERT-EU - This is a specific CSIRT on EU level and is the Computer Emergency Response Team for the EU Institutions, bodies and agencies

Additionally, a community initiative across coordinating CSIRTs has been started, the 'CSIRT network', which provides a forum where members can cooperate, exchange information and build trust. Although not really common, some of these coordinating CSIRTs maintain monitoring and detection capabilities on a regional or national level.
Of course the different layers are not independent. Especially during serious security incidents, there will be heavy exchange of information between the CSIRTS on the different layers and between CSIRTs on the same layer.

Current SOC / CSIRT capabilities

In ENISA's "How to setup up CSIRT and SOC", from December 2020 (1), a set of services has been identified that are typically provided by a SOC and CSIRT. These typical services are a subset of the CSIRT services framework compiled by the Forum of Incident Response and Security Teams (FIRST) (4). The main difference between the SOC and CSIRT (in practice this separation of duties usually is not quite as strict) is that the SOC provides a real-time monitoring and incident detection service, whereas the CSIRT further analyses an event they receive from the SOC and can coordinate mitigating actions in case the event turns out to be an actual security incident.

New SOC / CSIRT capabilities

Most SOCs and CSIRTs have a good set of capabilities (see also the top of Table 1), but present day SOC and CSIRT capabilities simply do not suffice to deal with the persistence and sophistication of professional threat actors also considering the increasing complexity of ICT infrastructures and shortage of skilled staff.

Therefore we need to increase the speed and effectiveness of detection of and response to ongoing attacks, and the scope, effectiveness and efficiency of proactive analysis of threats to the ICT infrastructure to enhance its cyber resilience.

To achieve this we need to introduce so called Security Orchestration, Automation and Response (SOAR) capabilities. Also we have to introduce automated security reasoning capabilities on the vulnerability, resilience and potential impact of an organisation's ICT infrastructure and automatically generate and assess response actions to ongoing attacks and emerging threats. Furthermore, in order to increase the speed of responding to ongoing attacks and emerging threats, the ICT infrastructure has to be adapted, so it can support automatic instantiating and/or reconfiguring of security controls. Such automated response execution capability will in many environments include a human-in-the-loop, but in modern ICT environments (e.g. programmable infrastructure, cloud-native technology) this may even be fully autonomous response systems. The new SOC / CSIRT capabilities are listed at the bottom of Table 1.

To establish the capabilities that the SOCCRATES project envisions for the future, a variety of technical challenges will need to be overcome:

- Actual machine-readable model of the infrastructure
- Improvement of detection capability and coverage
- Advanced use of Cyber Threat Intelligence
- Real-time Business Impact Assessment
- Recommend Course of Action (CoA) generation
- Automation and orchestration to improve SOC response

In the following sections these challenges are elaborated on by describing the current state and future needs.

Actual machine-readable model of the infrastructure

Although inventory and control of hardware and software assets are essential elements in many cyber security frameworks (e.g. NIST Cyber Security Framework (3), CIS Critical Security Controls for Effective Cyber Defense (2)), many organizations still struggle with keeping their asset inventory up to date. As our IT environments are getting more dynamic, it is becoming an even more challenging task. SOC analysts, however, need to interpret and understand security events in the context of the continuously evolving ICT networks and systems of an organisation. The SOC and CSIRT analysts also need to understand the critical attack surfaces, the attack paths that may lead to a compromise of assets, as well as defence mechanisms present and/or can be enforced to counter an attack. For a human analysts the infrastructure information has to be visualized in a comprehensive manner such that it is easy to understand and security events can be projected on top of the infrastructure to create real cyber situational awareness. Moreover, for automated security reasoning and decision support capabilities, such as automated threat modelling and simulation and real-time business impact assessment, the infrastructure information has to be current, accurate and machine-readable and made available via Open APIs in a standardized format.

There are promising developments (e.g. Software Bill of Materials (SBoM) (5)) and new products entering the market that enable access information on assets (incl. installed software), network topology and vulnerabilities in the infrastructure. But these products do not often provide an API for third party tools to collect a machine-readable standardised model of the infrastructure for (third party) security analysis tools.

The SOCCRATES project foresees the following needs for the near future:

- Improve asset discovery and change detection. There is a need for better asset discovery of a wide range of asset types, both from an internal and external viewpoint, in near real-time, going beyond just IP/port/service detection and into fingerprinting of the make-and-model of all assets. Moreover, uniquely mapping of information about the same asset from different data sources (e.g. Network MAPper (nmap), vulnerability scanner, AD, netflow) to a single object in the data model is challenging.
- Improve access to asset management systems. Access to asset management systems is needed to provide accurate up to date infrastructure information at different levels of detail or granularity (such as the make and model of assets).
- Better visualisation of ICT infrastructures. Visualization is needed with the capability to overlay security status and event information.

- Create the ability to provide historical infrastructure model information. This might be limited to a certain point in time and with gradually decreasing level of detail, but will help understanding historical log events during threat hunting.
- More standardisation of data models. Standardisation of the data models describing the ICT infrastructure in a machine readable manner.
- Improve automatic discovery of security functions in a machine readable manner. This should include information on scope (i.e. what security functions do they provide for which assets?), whether these functions are configurable, via what API, etc.

Improvement of detection capability and coverage

A major activity of a SOC is to respond to the alerts that are generated by detection systems. Approaches to detecting cyber-attacks can be broadly placed into two categories: those that use signatures that describe adversarial behaviour, versus those that aim to detect anomalies that manifest in collected data and could indicate a cyber-attack. For the latter approach, there is increasing interest in applying machine learning algorithms to learn a model of normal behaviour and use this as a basis for detection. The advantage of anomaly-based detection approaches is that novel – previously unseen – attacks can be detected, if the manifestation of their behaviour deviates from a learned norm.

The SOCCRATES project foresees the following needs for the near future:

- Improve detection capability across IT and OT systems. Whilst advancements are being made, OT systems have traditionally not been monitored for adversarial behaviour to the same extent as their IT counterparts. With the integration of these systems, increased attention has been paid to this issue. Although detection systems for deep-packet inspection of industrial protocols (e.g. Modbus, DNP3, OPC UA, etc.) exist, endpoint monitoring and detection on OT devices and infrastructure is still relatively immature or absent. The result is that OT visibility is limited.
- Improve detection of prevailing adversary techniques and procedures. A major challenge for a SOC, is to determine whether a deployed detection posture is able to effectively identify techniques and procedures that are of concern. Knowledge-bases, such as the MITRE ATT&CK Framework, provide insights into the data sources that could be used to detect specific techniques but there is a gap between this information and that needed to determine whether specific procedures that an adversary is using can be detected. This problem is exacerbated by adversaries adjusting their procedures to avoid detection.
- Increase effectiveness of detection of security events in large data sets. The amount of data that can be used to detect security events is growing tremendously. One apparent challenge here is to determine which of all this data is worthwhile paying attention and applying resources to in order to gain useful insights. Put simply, where should one start to detect an attack?
- Decrease number of false positives. Large volumes of data also exacerbate a well-understood problem that is associated with anomaly-based detection systems: false positives, i.e. alerts that indicate malicious behaviour when none exists. The job of the cybersecurity data scientist is to improve detection performance, as much as possible, using techniques such as feature engineering or tuning the hyper-parameters of deep learning models. The goal is to reduce the false positive rate so that SOC analysists do not waste time fielding unwarranted alerts.
- Improve response on detected incident. The obvious advantage of anomaly detection techniques is that one does not need to prescribe the adversarial behaviour to be detected – the norm is learned by a machine learning algorithm and if a sample deviates from this norm, an alert is generated. However, there is arguably a (semantic) gap between what an anomaly detection system generates and insights that can lead to steps to mitigate an attack (i.e. the invocation of a playbook that is related to a specific class of attack). For example, it is not immediately apparent whether a detected anomaly relates to a ransomware attack or perhaps data exfiltration – two types of attack that require distinct responses. Automated support for this activity should help to improve the effectiveness of a SOC, as it aims to realize its KPIs.
- Increase resilience against Adversarial Machine Learning. Machine learning (ML) and artificial intelligence (AI) are finding increasing utility in SOC operations. However, also attackers are exploring the benefits of AI and ML. So-called adversarial machine learning can take many forms. An attacker's goal can

The key to detecting adversary behaviour is procedures.

include model theft and poisoning, for example, and subverting a model's output, in order to cause misclassification. Because machine learning is applied to ever-increasing mission critical applications and adversaries explore this new form of attack, it could become a major future challenge.

Advanced use of Cyber Threat Intelligence

Apart from knowing what you are defending, you also need to know the enemies and their capabilities against which you are defending. This is the goal of Cyber Threat Intelligence (CTI). CTI is evidence-based knowledge about threats that provides situational awareness and actionable decision support. CTI can be further divided into subtypes: strategic, operational, tactical and technical (6). The tactical and technical subtypes are the most relevant for SOC and CSIRT needs. Tactical CTI is knowledge about adversary behaviour, and is referred to as the Tactics, Techniques and Procedures (TTPs) of the adversary. Technical CTI is knowledge about specific malware, tools or infrastructure. Examples are file hashes, IP addresses and domain names observed in an incident and shared as Indicator of Compromise (IoC).

Although IoCs can directly be used to detect or hunt for malicious behaviour, the volume of shared IoCs is very large and they changes quickly. More quickly than the associated TTPs. Detecting adversary behaviour based on the TTPs lets defenders therefore stay ahead of the attackers. Another advantage of tactical CTI is that TTPs can be used for adversary emulation, as is done in Threat Intelligence Based Ethical Red-teaming (TIBER) (7).

Nowadays, IoCs are extensively used by SOCs and CSIRTs in an automated fashion. Threat feeds are downloaded and used to compile a signature for attack detection. Additionally, CSIRTs automate IoC sweeps on logs to find historical intrusion activity that was not detected when the activity took place. The application of tactical CTI is, however, largely a manual process. The underlying reason for this is lack of machine readable standards. MITRE ATT&CK is first and foremost a knowledge base of techniques, linked to adversary groups and software. The tactics in ATT&CK are tactical objectives, not actually tactics. But more importantly, the procedures in ATT&CK are human readable examples, not suitable for processing by a computer.

The key to detecting adversary behaviour is procedures. ATT&CK provides no guidance on how to define procedures in a machine readable format, and the same applies to the standards for sharing CTI (e.g. Structured Threat Information Expression (STIX) and Malware Information Sharing Platform (MISP) formats).
The SOCCRATES project foresees the following needs for the near future:

- Improve quality, relevance and timeliness of technical CTI (i.e. IoCs) to reduce false positive alerts and exhausting limited resources of the SOC and CSIRT chasing non-incidents. New methods are needed to contextualise IoCs to help defenders with prioritisation.
- Increase level of automation for collection, sharing and processing of tactical CTI to enable adversary behaviour detection and assessing the infrastructure with adversary emulation. This includes describing adversary behaviour in a machine readable format, and developing methods and tools for automatically process and use this information for detection and attribution.

Real-time Business Impact Assessment

The impact of attacks on an infrastructure is usually analysed from a technical point of view: the logs and the alerts raised by intrusion detection systems allow a SOC analyst to identify the assets targeted by the attacks and, with the help of attack graphs based tools, predict the potential attack path among the other assets of the infrastructure. This approach is essential, as it greatly facilitates the deployment of courses of action that will both mitigate the attack and correct vulnerabilities. However, this technical analysis does not take into account the operational impact, i.e. to which extent the attack will disrupt the organisation of the company departments. Therefore, in addition to understanding the ICT infrastructure, the SOC analyst needs to be able to assess the potential impact on the business of an ongoing attack or emerging threat. To do so, it is necessary to not only develop a model of the business processes, but also be able to process this model and obtain computable metrics.
In the context of SOC/CSIRT environments, impact analysis on business processes is not usually done. Typically SOC and CSIRTs use predefined lists containing the Business Impact Assessment scores per host, in terms of Confidentiality, Integrity and Availability. More specific analysis of business impact is done manually and in collaboration with the business owner of the particular system., which is time consuming and does not allow the courses of action selection to match the business priorities during an ongoing attack on the infrastructure. Moreover, it does not allow for an assessment of the negative consequences to the business by deploying one or more courses of action. In order for such types of business impact assessments to be performed, a model of the business processes and functions is necessary. Business processes need to be mapped on the ICT infrastructure components, and insight in the consequences of a breach of confidentiality, integrity and/or availability of system resources or information assets needs to be (near real-time) available.

The SOCCRATES project foresees the following needs for the near future

- Improve (automatic) identification of business functions and -processes. It would be extremely useful to at least partially automate the identification of the company's business functions & - processes, as well as their dependencies. Including the dependencies with the assets from the infrastructure that directly support business functions. The main challenge to overcome is the lack of automation solutions in the state of the art. Methodologies to elaborate Business Process Model and Notation (BPMN) models are well known, but usually rely on manual work done beforehand, involving discussions and interviews with various services in the company. However, BPMN almost entirely decorrelates the business view from the technical view, which means that the link between the business entities and the assets must also be defined manually, though without any established methodology.
- Computation of relevant metrics to perform the business impact analysis. The challenge is to design a scalable mathematical model that is able to compute various metrics in real time, all while taking into account things such as asset redundancy and interdependencies and the specificities of the attack. To do so, well known graphical models, such as Bayesian networks, can be exploited, but will often require specific adaptations to match real life situations. Moreover, a

realistic model will need frequent data updates to match the dynamic nature of the business impact. Also, business impact is temporal by nature, the impact would typically be different during business hours compared to weekends, or may depend on seasonal aspects (e.g. point of sale system during the weeks before Christmas), or may depend on particular production orders.

Recommend CoA generation

To be able to automatically suggest optimal courses of actions (CoAs) for improving security in ICT infrastructures we can analyse cause and effect of various possible defence actions related to the infrastructure in a model (in popular terms; a digital twin), before getting into action with implementation. In general, the more detailed this analytic model will be, the better the suggested actions can be. And the model quality depends both on how much "raw data" from the ICT infrastructure is available and how well the model language captures the facts about what actions indeed are efficient security improvements, given different states of the infrastructure. With the model, we can examine the preventive measure optimization, in which we have to weigh and aggregate multiple assumptions made in various scenarios. One thing to assess is the expected shortest time it would take for a simulated attacker to traverse the attack graph connecting the starting and target points. And, with added defensive actions and enabled security controls we expect the estimated time to compromise (TTC) of the selected target(s) to increase, which improves security. By enabling or disabling defences time estimates for different attack vectors varies, and the defender can elaborate on good ways to increase the TTC for the attacker. The challenge we face here though is that the potential action space for the defender is very large, even for just a moderately sized ICT infrastructure. The CoA generator is thus tasked with finding highly effective defense action combinations, sparing the defender the work of trial-and-error simulations of testing different defense strategy hypotheses.

The SOCCRATES project foresees the following needs for the near future

- Improve asset management. As already mentioned, one of the biggest challenges for building an Infrastructure model is the challenge of discovering all the components in an ICT infrastructure. Even though we believe that this will remain a challenge for quite some time we can note that this situation is improving significantly with numerous new tools and tool capabilities. Also, we can note that the challenge is significantly smaller for cloud environments where the infrastructure is deployed from code and does not have to be discovered.
- Improve mapping of the detection space and the security analysis space. If we know that some particular asset has been compromised, an attack simulation with some assumed attacker starting point (such as the internet) will give the easiest attack vector to reach the compromised node. Looking for additional traces of breach along this vector is probably a good starting point to learn more about the incident. In principle we would like to be able to generate attack graphs that also include information on which attack steps can be detected, including the quality of detection.
- Improve visualization and contextualization of CoAs. A great support for a SOC analyst would be the capability to visualize and contextualize the CoAs depending on different threat scenarios and use cases.

Automation and orchestration to improve SOC response

Around 2015 technology started to emerge that we now call Security Orchestration, Automation and Response (SOAR) solutions. Initially these solutions were developed out of convergence from three different technologies: a) security incident case management platform with structured incident response workflows or playbooks, b) threat intelligence platforms that integrate automation for CTI processes, and c) tools for integration of different security tools/technologies in a coordinated way (playbooks). The combination of orchestration and automation for security operations refers to the tasks perform by a SOC analysist collecting information from multiple systems to support the decision-making process. The tools that entered the market could perform mundane repetitive tasks and thereby speed up incident investigations.

Also standardisation to support automation and orchestration of security operations has started. In particularly,

- Open Command and Control (OpenC2) (8), specifications to

enable machine-to-machine communications for purposes of CoAs execution
- Collaborative Automated Course of Action Operations (CACAO) (9), specifications for documenting playbooks for cybersecurity operations and sharing these across organisational boundaries.

The current state of security tools at many organisations can best be described as a plethora of disparate products from different vendors or sources. SOAR solutions can help with the integration and aggregation of the information from the diverse multi-vendor security products and tools, but the diversity and lack of standardised data formats is challenging.

Another challenge when deploying a SOAR solution is the fact that these tools require a significant amount of manual tuning and playbook definition. In addition, it remains to be seen how effective current SOAR solutions are with the increasing number of security events and alerts an organization has to cope with. Note that many of the simultaneous triggers may be related to the same security event. Handling of multiple simultaneous triggers and running different playbooks for related security events needs to be studied further.

The following future needs has among others been identified:

- Increase support for deployment of SOAR tools in SOC and CSIRTs, including integration of diverse security products and tools and sharing of playbooks that can easily be tuned and adopted.
- Improve how to deal with number of playbooks triggered and simultaneously handle potentially on related or even the same security incident.
- Automate playbook generation for execution of dynamically generate response actions. This includes translation from abstract response actions into specific reconfiguration commands for one or more security functions.
- Improve the interaction of the human analysist with SOAR, or security automation in general, will be a topic of concern for the coming years. Since there is a shortage of skilled cybersecurity staff there is much focus on training and education of cyber security personnel. But how will the role of the SOC and CSIRT analyst change in the coming years due to the introduction of security automation?

Concluding remarks

It is clear that SOC and CSIRTS need to transform. The SOC/CSIRT capabilities need to be strengthened and expanded, new capabilities are necessary to be able to handle future threats. Building and implementing these capabilities will have impact on all aspects of the SOC/CSIRT operations, including the interaction with the outside world.

Lookout to next articles

In the coming articles (next editions of the PvIB magazine) we will zoom in on the Orchestration and Integration Engine of the SOCCRATES platform and on the pilot evaluation. More info and the vision, roadmap and guidance for SOC booklet are available at www.soccrates.eu. More detailed information regarding this article can be found in the SOCCRATES vision paper:
https://www.soccrates.eu/wp-content/uploads/2022/05/soccrates_vision_paper_downloadable.pdf

SOCCRATES has received funding from the European Union's Horizon 2020 Research and Innovation program under Grant Agreement No. 833481.

References

(1) Edgars Taurins, How to setup up CSIRT and SOC - good practice guide. ENISA. 2020. (https://www.enisa.europa.eu/publications/how-to-set-up-csirt-and-soc)
(2) Klaus-Peter Kossakowski, Computer Security Incident Response Team (CSIRT) Services Framework, version 2.1. November 2019. Forum of Incident Response and Security Teams, Inc. (FIRST.Org). (https://www.first.org/standards/frameworks/csirts/csirt_services_framework_v2.1)
(3) Framework for Improving Critical Infrastructure Cybersecurity. Version 1.1. NIST. 2018 (https://www.nist.gov/cyberframework)
(4) CIS Controls. Version 8. Center for Internet Security, Inc. (CIS). 2021 (https://www.cisecurity.org/controls)
(5) https://www.ntia.gov/sbom
(6) https://nsarchive.gwu.edu/document/17212-united-kingdom-government-threat-intelligence
(7) https://www.dnb.nl/voor-de-sector/betalingsverkeer/tiber-nl
(8) https://www.oasis-open.org/committees/openc2
(9) https://www.oasis-open.org/committees/cacao