SOCCRATES

Vision, Roadmap & Guidance for SOC

SOC & CSIRT Response to Attacks & Threats based on attack defence graphs Evaluation Systems

TABLE OF CONTENTS

- **FOREWORD** 4
- **INTRODUCTION** 6
  - SOC, CSIRT and MSSP 6
  - Positioning SOC and CSIRT 11
  - SOCCRATES vision 14
- **ACTUAL MACHINE-READABLE MODEL OF THE INFRASTRUCTURE** 20
  - Current state 21
  - Future SOC and CSIRT needs 22
  - Research roadmap 28
- **IMPROVEMENT OF DETECTION CAPABILITY AND COVERAGE** 29
  - Current state 30
  - Future SOC/CSIRT needs 31
  - Research roadmap 35
- **ADVANCED USE OF CYBER THREAT INTELLIGENCE** 38
  - Current state 40
  - Future SOC/CSIRT needs 41
  - Research roadmap 44
- **REAL-TIME BUSINESS IMPACT ASSESSMENT** 45
  - Current state 45
  - Future SOC/CSIRT needs 46
  - Research roadmap 49

RECOMMEND COA GENERATION 51

- Current state 53
- Future SOC/CSIRT needs 54
- Research roadmap 57

AUTOMATION AND ORCHESTRATION OF SECURITY OPERATIONS 59

- Current state 59
- Future SOC/CSIRT needs 62
- Research roadmap 66

OVERALL PERSPECTIVE AND RECOMMENDATIONS 68

- Overall perspective 68
- Recommendations 74

REFERENCES AND ENDNOTES 76

APPENDIX A SOCCRATES PROJECT 79

ABBREVIATIONS 82

CONTRIBUTING EXPERTS 83

FOREWORD

When asking CISOs of large organisations what keeps them awake at night, two different concerns always come to the fore. The first challenge is how to deal with the continuous rise of the number of attacks on their networks. This is consistently competing for the #1 spot with the problem of how to get enough security staff in their SOCs and CSIRTs to ward off these attacks.

The gap between attackers and defenders in the cyber arena is growing at an ever-increasing pace. Our opponents are creative, have ample funds at their disposal, and are not hampered by any legal or ethical obstructions. The time that a CISO mainly had to worry about the odd script kiddy is long gone. On top of that, the adversaries embrace novel techniques just as enthusiastically as any IT aficionado and implement them with professional zeal.

The obvious response to these challenges is straightforward in theory. To make optimal use of the limited number of cybersecurity personnel, we simply have to increase their effectiveness and efficiency. And when it comes to the exponential growth of the number of attacks, that are more often than not boosted by intelligent attack tooling, the logical answer lies in countering automation with automation.

In all honesty, a fair number of point solutions is available to cater to the needs of many organisations. The burgeoning cybersecurity market offers anything you might require, ranging from in-depth security training at any proficiency level for your staff to self-learning AI powered security solutions to fend off attacks. However, using point solutions as a stopgap is not sufficient anymore, as linear improvements in defense will always be outpaced by exponential growth of adversarial capabilities. A more holistic approach is called for.

That is why I am proud of the results of the SOCCRATES project. It is unique in its approach to encompass all elements involved in the detection of and response to ongoing attacks, and the proactive analysis of threats. Leading-edge automation reduces the relative dependence on the human factor, effectively leveraging its effectiveness and efficiency and freeing up scarce resources within the organization.

I highly recommend the output of the project. This vision paper conveniently outlines stepping stones for further development, and I am looking forward to the next steps that the research community will take based on this work. Also, I would like to underscore the opportunities for the private sector embodied in the results. Commercial parties are strongly encouraged to follow up on the concepts proposed in this vision paper, because there are markets to be expanded and fortunes to be made. SOCCRATES provides a range of exciting new research and commercial possibilities. Let us all step up to the plate to jointly create a safe and secure digital future.

Martin E. Pekárek
Research team lead
Dutch National Cyber Security Centre (NCSC-NL)
Member of the SOCCRATES Advisory Board

INTRODUCTION

Over the past years, the cyber threat landscape has greatly evolved. Advanced cyber-attacks are now conducted by professional threat actors that have substantial resources and (technical) capabilities. Such attacks are often targeted in nature and may involve a great degree of automation, persistence and (technical) sophistication. Meanwhile, the dependency in all sectors on Information and Communication Technology (ICT) and thus the potential impact of any cyber-attack is ever increasing. A complicating factor in defending against such cyber-attacks is the continuous evolvement of ICT infrastructures with diverse and emerging technologies (e.g. Cloud technology, Internet of Things (IoT)). In summary, defending against cyber-attacks has become a challenging task for most organisations.

To deal with these challenges, many organisations have increased their efforts in security monitoring and incident response. Some (typically large and mature) organisations established internal (self-reliant) Security Operations Centres and Computer Security Incident Response Teams. Other (typically smaller or less mature) organisations outsourced these tasks to a Managed Security Service Provider (MSSP). In recent years, these security operations activities have been extended with structural Cyber Threat Intelligence (CTI) activities, to anticipate upcoming threats and take appropriate precautions.

SOC, CSIRT and MSSP

Both a Security Operations Centre (SOC) and a Computer Security Incident Response Team (CSIRT)* are expert teams that provide operational security services to an organisation. In ENISA's "*How to setup up CSIRT and SOC*", from December 2020 (Edgars Taurins, 2020), a

* CSIRTs are known under several different names. Another common name is Computer Emergency Response Team (CERT). The scope and services provided may differ, as described on page 11.

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) (FIRST, 2019). 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 is a team that is assigned to handle cyber security incidents. In other words, incident monitoring & detection is done by the SOC and incident response is coordinated by the CSIRT. 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. The IETF RFC2350 (N. Brownlee, 1998) is an international standard that expresses the general Internet community's expectations of these CSIRT teams. This standard describes how and for what purpose other CERT/CSIRT organisations can approach the CSIRT in case of incidents.

The particular set of services provided by a SOC or CSIRT may differ among organisations. It is also quite common to use alternative names for similar types of organisational units, such as Cyber Defence Centre for SOC, which may also offer a super or subset of the services. It is good to have insight in the services provided by the SOC or CSIRT but to further develop the services offered, we need to establish which capabilities a SOC or CSIRT should have. In (Zimmerman, 2014) the author established a comprehensive list of SOC capabilities. An example of such a capability that is of particular interest for the SOCCRATES project, is the SOC tool life-cycle support. This includes operations and maintenance of all SOC systems, sensor tuning, engineering and deployment of (new) SOC systems, and even research and development of own SOC systems and tools.

In SOCCRATES, we prefer to use capabilities instead of services, since we develop technologies to enable or enhance SOC and CSIRT capabilities. These SOC and CSIRT

capabilities can be used to provide services internally or externally (e.g. by MSSPs) to organizations. In Table 1 the SOCCRATES view on typical SOC and CSIRT capabilities is shown. Please note that the table depicts the capabilities that can be found in *current* SOCs and CSIRTs.

As described further on in this document, the innovative automation solutions that SOCCRATES is developing will come with additional capabilities that SOCs and CSIRTs will need to develop.

Table 1 **SOCCRATES view on current SOC and CSIRT capabilities**

	SOC	CSIRT
Capabilities identified by ENISA	• Monitoring and detection • Event analysis • Information security incidents analysis (triage) • Vulnerability analysis • Security Awareness Creation	• Information security incidents analysis (incl. event correlation) • Threat analysis, including collection, sharing and processing of CTI and threat hunting • Artefact and forensic evidence analysis • Information security incident coordination • Vulnerability remediation coordination • Awareness Building
Additional capabilities	• 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)	

Organisations that decide to set up a SOC or CSIRT are confronted with many questions. One of the main questions is whether to run their own in-house SOC and CSIRT, or to outsource (parts of) the operational tasks to a commercial Managed Security Service Provider (MSSP). Some organisations choose an in-house SOC and CSIRT due to the sensitive nature of their data and incidents, and due to the strategical importance of managing the cyber threats for their business. Other organisations may choose to outsource (part of) their SOC and CSIRT tasks, e.g. because they may not have the expertise, nor the resources to build and maintain their own in-house SOC and CSIRT capability. Another reason for outsourcing could be that they require 24/7 security monitoring and response operations and are not capable to set this up themselves. In many cases SOCs are operating on 24/7 basis, but CSIRT teams usually are not 24/7 and generally work business hours while being on-call available during non-business hours.

In recent years there has been an increase of cloud-based security products being offered as Security as a Service (SECaaS). The security products and tools used by SOCs are more and more being provided as SECaaS offerings. Many organisations are, however, reluctant to deploy parts of the SOC technical infrastructure within a third-party cloud environment, as the data that a SOC collects and processes is security and privacy sensitive. SOCCRATES distinguishes four deployment options for SOC technical infrastructure, see Figure 1.

The SOCCRATES main pilot was organized on two pilot sites (and an external CTI provider). On both pilots sites the pilot was conducted on-premise, in one case in an MSSP environment (Norway) and in the other in an end user SOC (Poland). Thus the SOCCRATES pilot covers the deployment scenarios depicted in the first row of the figure.

Figure 1 **Four deployment options for SOC technical infrastructure**

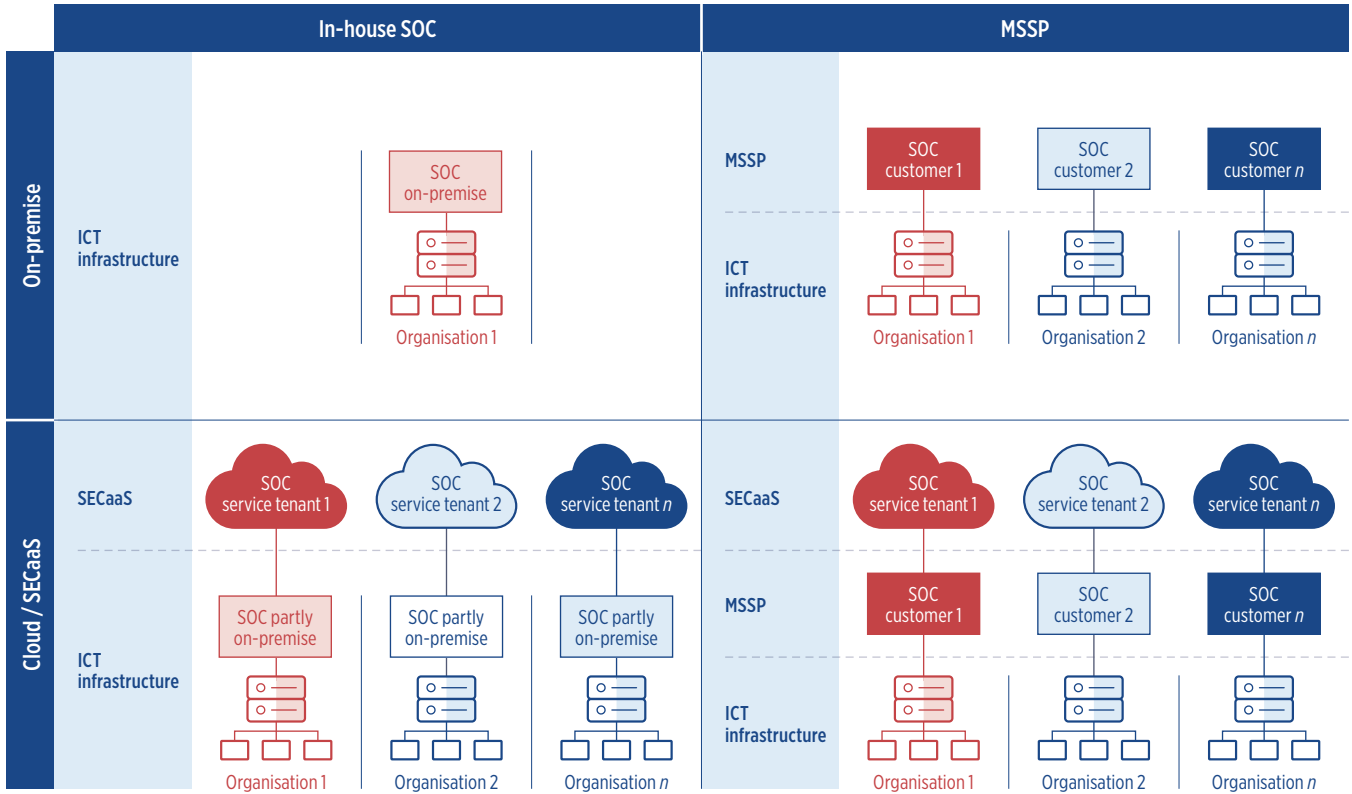

Positioning SOC and CSIRT

There are many developments on (regulation of) CSIRTs and SOCs in Europe. The terms SOC and CSIRT are applied at different levels, as can be seen in Figure 2.

In the lowest layer of the figure we see the SOC in different set-ups:

- Organizations that maintain their own (in-house) SOC and CSIRT
- Organizations that make use of SOC and/or CSIRT services provided by commercial MSSPs.

This layer actually is 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, best practices, and training material, and
- Providing security incident analysis and forensic services.

Note that coordinating CSIRTs sometimes operate at a same level as the Information Sharing and Analysis Centres (ISACs). ISACs are non-profit member driven organizations that facilitate gathering and sharing of information on cyber threats (relevant for the members). ISACs have been created around particular private sectors (e.g. finance, energy, health). They may facilitate two-way sharing of information between the private and the public sector, and serve as a focal point on the national level to gather information about cyber incidents and analyse it (ENISA, 2018). In that sense, ISACs play an important role within EU Cybersecurity Strategy.

In Figure 2 the following coordinating CSIRTs are depicted:

- Sectoral or Regional CSIRTs - Some sectors or regions have 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. Sectoral CSIRT activities may go beyond the ISAC activities, e.g. through security incident coordination at the sectoral level. Examples where such sector CSIRTs have been set up are healthcare, government and finance.
- 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. Examples of such national CSIRTs are the Dutch National Cyber Security Centre (NCSC) or the Norwegian NorCERT, hosted by the Norwegian 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. It is composed of IT security experts from the main EU Institutions. CERT-EU helps to keep the ICT infrastructures of all EU institutions, bodies and agencies secure and cooperates with other CERTs/CSIRTs in the Member States and with specialised IT security companies in order to respond to information security incidents and cyber threats.

Additionally, a community initiative across coordinating CSIRTs has been started, the 'CSIRT network'. This is a network composed of EU Member States' appointed national CSIRTs and CERT-EU ("CSIRTs Network members"). The European Commission participates in the network as an observer. ENISA is tasked to actively support the CSIRTs cooperation, provide the secretariat and active support for incident coordination upon request. The CSIRTs Network provides a forum where members can cooperate, exchange information and build trust. Members will be able to improve the handling of cross-border incidents and even discuss how to respond

Figure 2 **SOC and CSIRT clustering and layering**

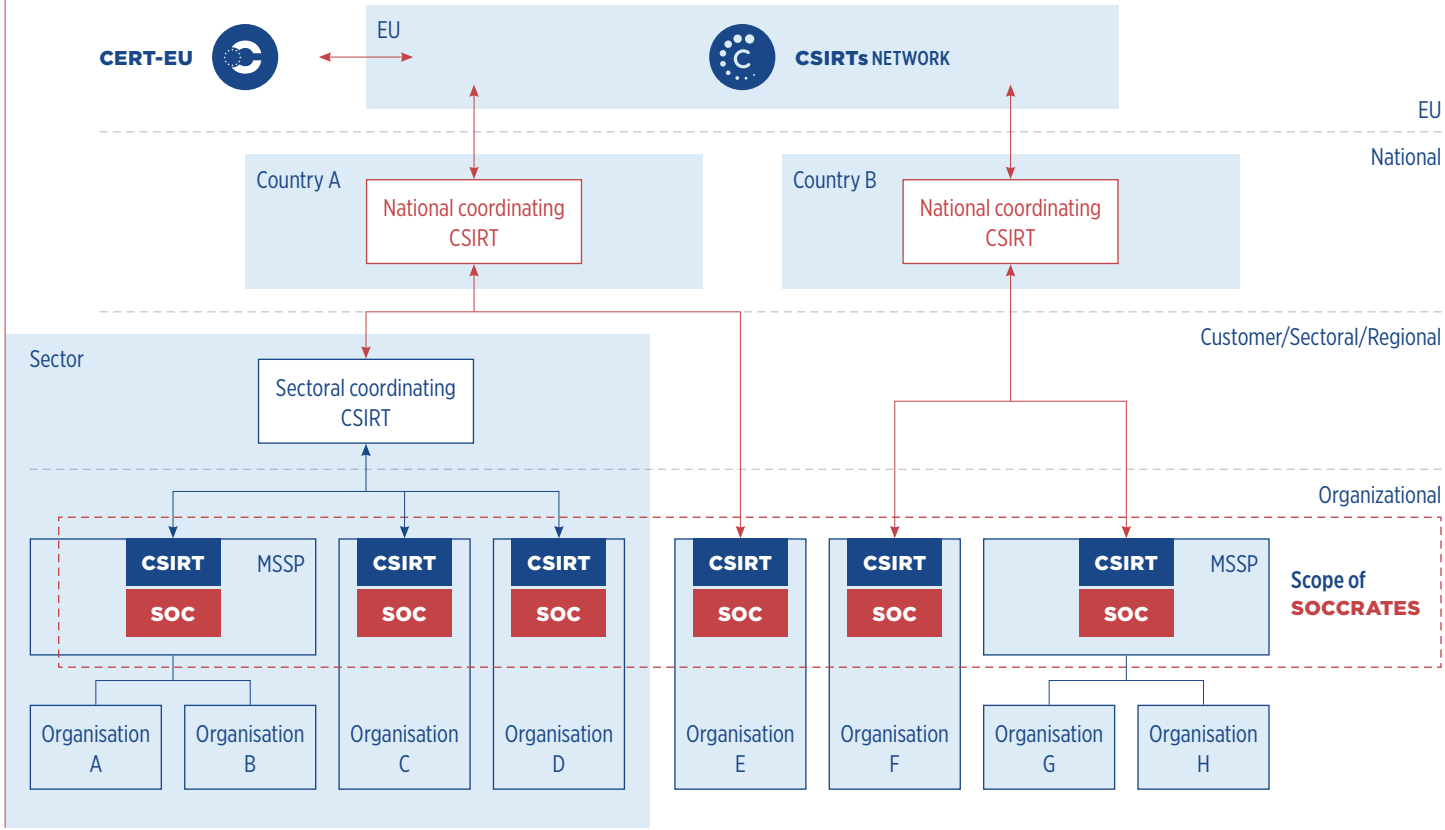

in a coordinated manner to specific incidents. The NIS Directive in Article 12 establishes the CSIRTs Network "to contribute to developing confidence and trust between the Member States and to promote swift and effective operational cooperation".

Note that some of the higher level CSIRT initiatives also maintain monitoring and detection capabilities on a regional or national level. This is, however, not really common considering the sensitive nature of monitoring data and specific (sensitive) knowledge required of monitored infrastructure.

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 CSIRT on the same layer. This not directly in the scope of SOCCRATES, but we have taken this into account in the design of the SOCCRATES platform and interfaces.

SOCCRATES vision

The increasing dependency of organisations and society as a whole on IT systems and networks and the increase of cyber security incidents with major impact, has led to organisations (and governments) increasing their spending on cyber security. As already indicated in the previous sections, many organisations have established a SOC and CSIRT to protect the organisation against cyber-attacks or contracted a MSSP to perform these operational cyber security services for them. However, despite the increased spending and new investments in cyber defences, most organizations are unable to keep pace with the ongoing evolution of threats and attack methods. Present day SOC and CSIRT capabilities simply do not suffice to deal with the persistence and sophistication of professional threat actors. As it stands now, the gap between defenders and attackers will only increase further in the coming years. The four following challenges contribute greatly to this trend.

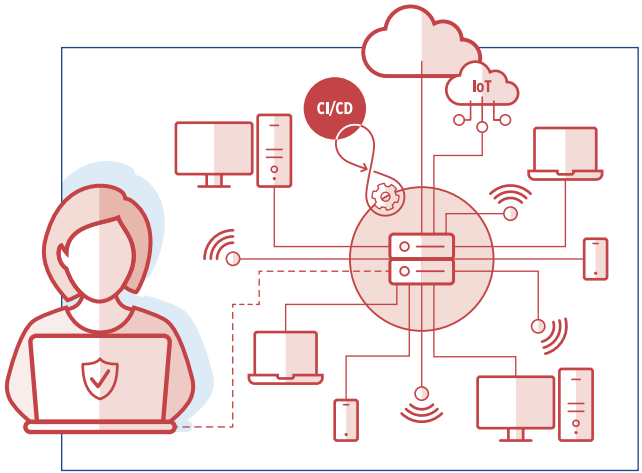

ICT is increasingly complex, changing and interdependent

The ICT infrastructures of organisations are continuously changing at an ever increasing pace. Introduction of new technologies such as IoT, cloud services and CI/CD, demand that security teams adapt and learn about their security implications.

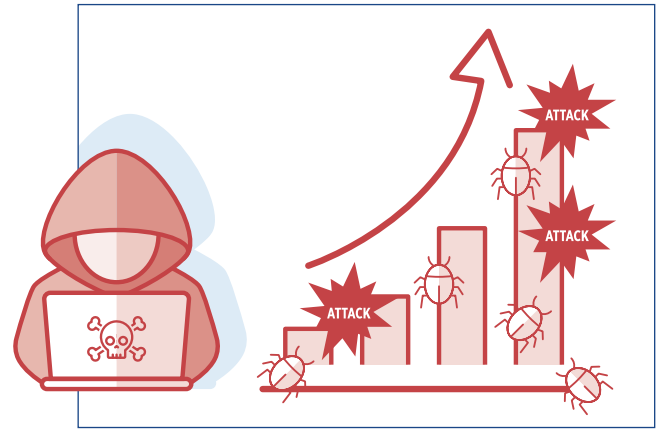

Increasing number and complexity of attacks

The threat landscape is continuously changing and security incidents have been on the rise. Cyber-attacks are becoming more complex, targeted and automated. Security teams have to continuously learn about these new threats and attack techniques and be able to quickly adapt so they can defend against them.

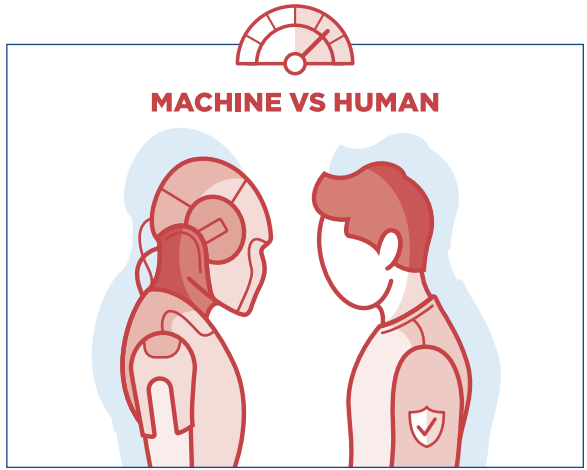

Attacks get more automated, response is human-centred

Threat actors are increasingly using automated attack tools. The response handling in SOCs and by CSIRTs, however, is still very human-centred. This will lead to an uneven battle: machine speed versus human speed. In order to reduce the time to detect, respond and recover, security teams have to increase the level of automation in incident response.

Shortage of skilled cybersecurity staff

The quality and effectiveness of security operations heavily depends on skilled cybersecurity experts, while at the same time there is a clear shortage of skilled cybersecurity workforce for several years. Recruiting, hiring and retaining cybersecurity staff for SOC and CSIRT is challenging.

The ever increasing gap between attackers and defenders can only be stopped through a fundamental *game changer*. The SOCCRATES partners believe that **automation of cyber security** operations can be that game changer. The higher level of automation in SOC/CSIRT operations will 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.

The present market already offers solutions (Security Automation, Orchestration and Response (SOAR) solutions) to automate some of the SOC and CSIRT work, but these mostly focus on playbook driven incident response. The SOCCRATES vision and innovation results go beyond the current SOAR ideas (see also Figure 3). We 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* (Courses of Action, CoA) to ongoing attacks and emerging threats. To further increase the speed of threat & attack response we need an infrastructure that supports automatic instantiating and/ or reconfiguring of security controls. Such *automated response 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.
In summary, we need to develop new capabilities on top of the already existing capabilities (see Table 2).

The SOCCRATES project has identified and worked on the following challenges that will support the implementation of the new capabilities:

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

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

Figure 3 **Proactive and reactive responses**

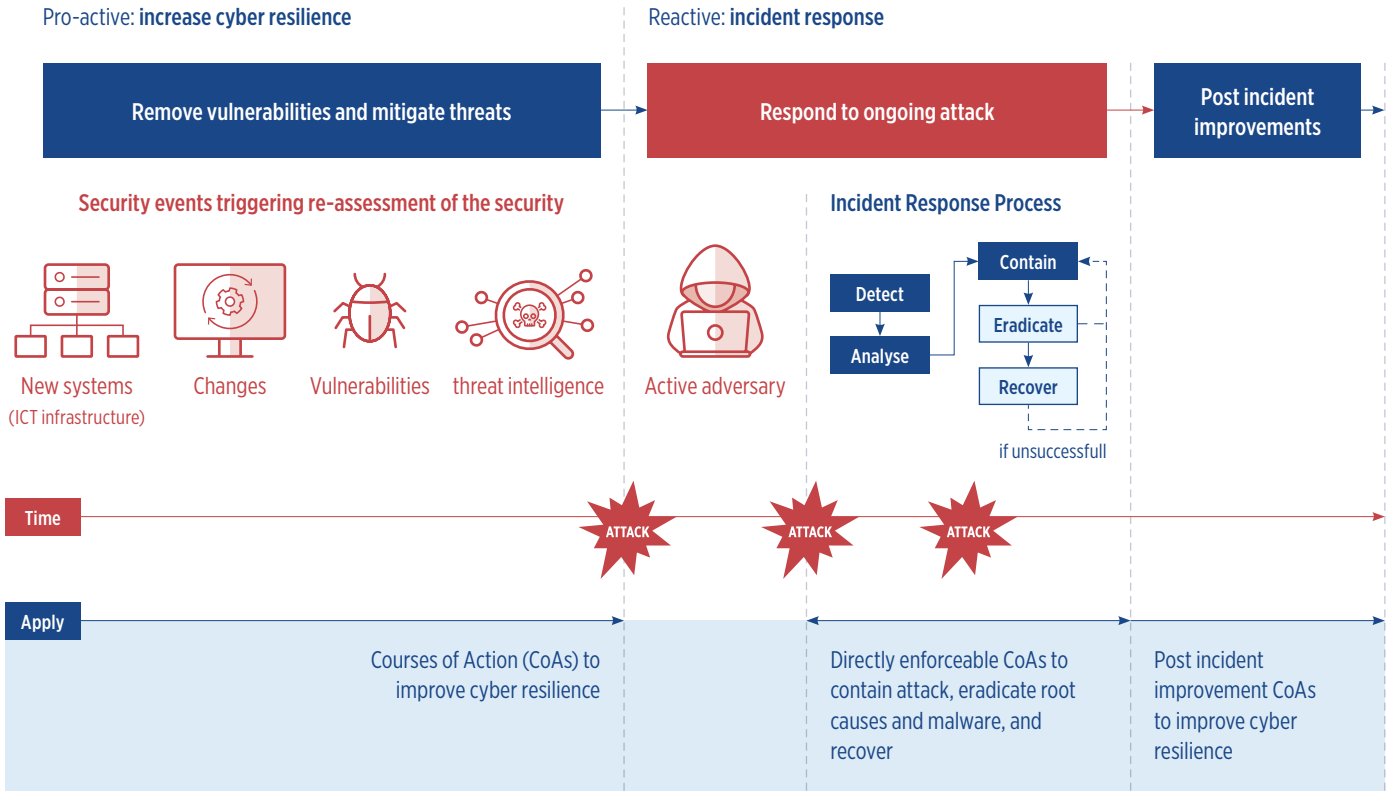

Table 2 **Proactive and reactive responses**

	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, including collection, sharing and processing of CTI and threat hunting • Artefact and forensic evidence analysis • Information security incident coordination • Vulnerability remediation 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 Automation, Orchestration and Response (SOAR) • Automated security reasoning (real-time threat & impact assessment) • Automated generation and assessment of response actions (both pro-active & reactive) • Automated response	

ACTUAL MACHINE-READABLE MODEL OF THE INFRASTRUCTURE

You can't defend what you don't know! This is a well-known statement in IT asset management and a crucial foundation for any security program. Although inventory and control of hardware and software assets are essential elements in many cyber security frameworks (e.g. NIST Cyber Security Framework[1], 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 and Internet of Things (IoT) devices are introduced, 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 critical business assets, as well as know what defence mechanisms are present and/or can be enforced to counter an attack (e.g. slowing down the adversary, containing compromised assets, or evict adversary from the network). 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.

In other words, we need tools and techniques that provide an actual machine-readable model of the infrastructure, that includes information on among others the network topology, the hosts on the network, the soft-

ware or applications running on these hosts, the user accounts and authorisations, known vulnerabilities, and the security controls and their configurations. Note that this should not just be from an internal view point of the IT infrastructure. Due to blurring of network boundaries (working from home, supply chain issues) and increased use of cloud technologies, it is becoming more important to also have an external view of the organisations IT infrastructure. Organisations should ask themselves: *what is the exposure of the organisation's IT assets to the outside world?* And *how do we manage the external attack surface?* These are a critical question that organizations need to be able to answer quickly, to avoid being victims of various high-profile incidents, as has been demonstrated by the HAFNIUM/Exchange[3] and the Apache log4j[4] incidents, not to mention multiple vulnerabilities in VPN products[5] and exposure of remote access services, such as RDP[6].

Current state

There are many (commercial) products and open source tools that can provide information necessary for creating a model of the IT infrastructure. This includes network scanning tools (e.g. nmap), vulnerability scanners, agent based host query tools (e.g. OSquery), asset discovery and management tools. These are, however, typically point solutions. It is becoming more common to support integration with other tools and data sources, but it is usually a cumbersome tasks due to the different data formats, time synchronisation, the large amounts of raw data, etc. Some academic approaches collect and correlate heterogeneous sources in real-time to enable Cyber Situational Awareness include (S. Jajodia, 2011), (S. Noel, 2016), and (Martin Husák, 2021). With cloud infrastructures such as AWS, Azure and tools such as OpenStack and Kubernetes, it is becoming easier to collect data via APIs on

what is deployed in the cloud infrastructure and how the network is configured in a machine-readable manner. There is, however, use of proprietary concepts and data formats and lack of standardisation among the different cloud providers, which will make this task difficult.

A number of external services exist that regularly scan the Internet and can report on external attack surface exposure. For example, SOCCRATES project partner Shadowserver scans the Internet regularly on over 60 different protocols. Any organization can report their IP address space (by CIDR or ASN) or domains to Shadowserver, and automatically gain free access to daily updates of infrastructure exposure and more[7]. Commercial services like Shodan[8] or Censys[9] can offer similar functionality. Others try to assign exposure risk scores and ratings to organizations (for example, Bitsight[10] or SecurityScorecard[11]).

It is clear the market has identified that there is a need for SOCs and CSIRTs to have access to accurate information on assets, network topology and vulnerabilities in the infrastructure they are protecting. New products and services are entering the market focused specifically on providing this information to security teams. Gartner has even identified a new product category on its Hype Cycle for Network Security 2021, called Cyber Asset Attack Surface Management (CAASM) and defined it as "an emerging technology focused on enabling security teams to solve persistent asset visibility and vulnerability challenges. It enables organizations to see all assets (both internal and external) through API integrations with existing tools, query against the consolidated data, identify the scope of vulnerabilities and gaps in security controls, and remediate issues." This is a promising development, but we have to be wary. Many of these products are new and may not provide an API for third party tools to collect a machine-readable standardised model of the infrastructure for (third party) security analysis tools.

Future SOC and CSIRT needs

SOC/CSIRT analysts need to be able to have up-to-date and correct information on the ICT infrastructure. They need to be able to visualize this data and to collect this in a machine-readable manner. Moreover, the ICT infra-

structure information should be provided in a manner that it can be used to generate a model or a digital twin of the ICT infrastructure for automated security analysis.

We foresee the following needs for the near future:

- *Asset discovery and change detection* 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. nmap, vulnerability scanner, AD, netflow) to a single object in the data model is challenging. Unique identification and naming of assets and objects must be addressed in order to build an actual machine-readable model of the infrastructure.
- *Access to an asset management system(s)* that can provide up to date infrastructure information at different levels of detail or granularity (such as the make and model of assets). For some situations it is sufficient to identify host and network topology based on IP addresses and routing information, other times it is necessary to be able to understand what software is installed and running on a host, and what this software actually consist of (i.e. open source and third-party software components used). In the latter case one may want to have access to the Software Bill of Materials (SBoM)[12]. The infrastructure information must be continuously updated to ensure the latest correct information is available.
- *Visualisation of ICT infrastructure* with the capability to overlay security status and event information.
- 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.
- *Standardisation* of the data models describing the ICT infrastructure in a machine readable manner.
- *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.

Infrastructure Modelling Component (IMC)

SOCCRATES developed an Infrastructure Modelling Component (IMC) to produce a *current and accurate model of the ICT infrastructure* and thus enable automated security analysis tools to reason about the security state and supply visualizations that accommodate human comprehension of the IT infrastructure.

As we observed that there is no commonly acknowledged ICT infrastructure modelling standard currently available that fulfils the needs of the SOCCRATES project, the project started to develop an appropriate reference (meta) model based on the SOCCRATES needs. The new reference model adopted relevant parts from the existing initiatives in the cyber-security industry and standards. The IMC consists of adapters to collect infrastructure data from different sources (e.g. vulnerability scanners, Nmap scans, NetFlow, domain controller), graph database, open API, visualization and triggering function. The core of the IMC is based on the ACT platform from mnemonic.

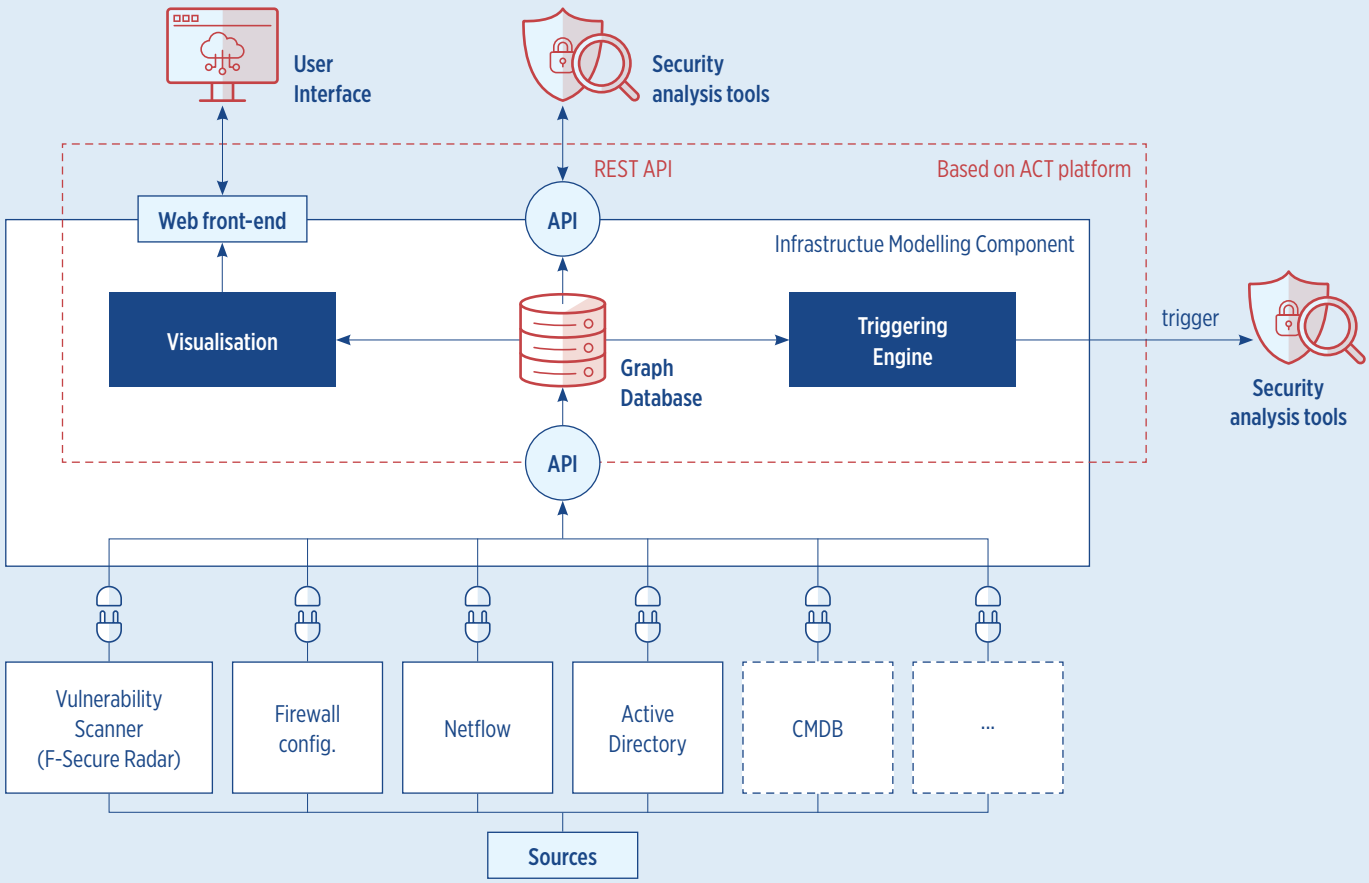
User Interface
Security analysis tools
REST API
Based on ACT platform
Web front-end
API
Infrastructue Modelling Component
Visualisation
Graph Database
Triggering Engine
trigger
Security analysis tools
API
Vulnerability Scanner (F-Secure Radar)
Firewall config.
Netflow
Active Directory
CMDB
...
Sources

Attack Defence Graph (ADG) analysis

SOCCRATES integrated and enhanced the securiCAD product for Attack Defence Graph (ADG) analysis (developed by our partner foreseeti) into the platform for SOC and CSIRT to increase situational awareness by providing insight in the vulnerability of the ICT infrastructure as a whole. The role of ADG in SOCCRATES is to predict how attacks propagate over a given ICT infrastructure. The ADG is based on research from KTH (Pontus Johnson, 2018), (Sotirios Katsikeas, 2020) and can be used to predict, for instance, the potential effect of a new threat, a new vulnerability, or a change in the system configuration. ADG based analysis can also be used to determine if and how the attacker may compromise other systems when an attack is detected. The ADG utilizes the data provided by the IMC and performs an analysis based on

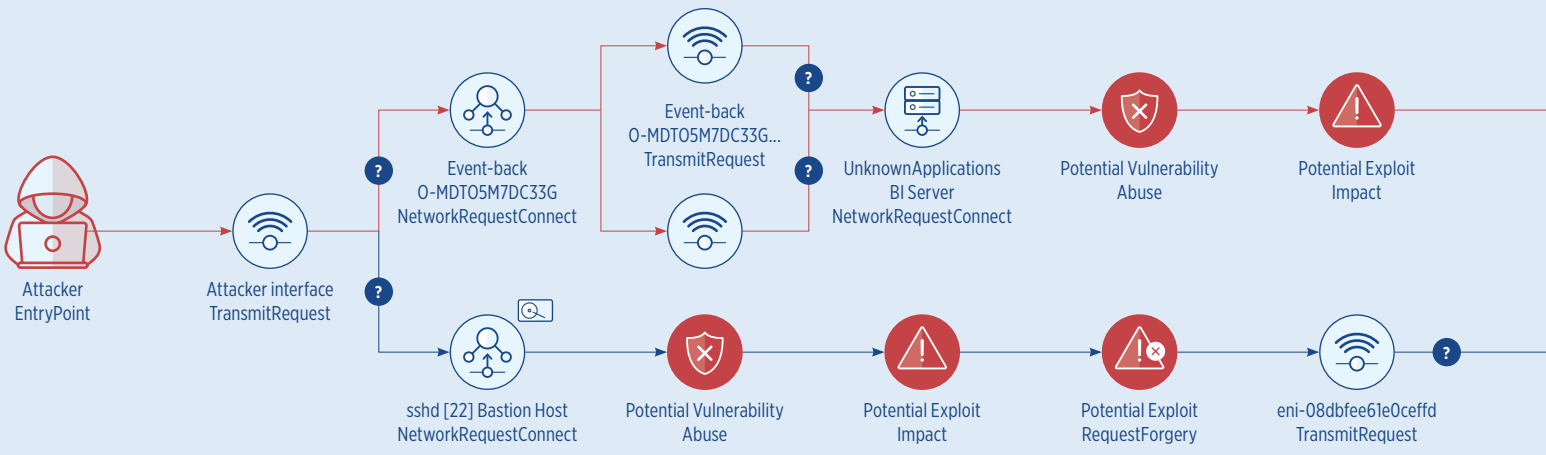

probabilistic attack simulations that are further enriched with cyber threat information from the threat intelligence platform. As a result, the ADG provides quantitative metrics on risk exposure and time to compromise (TTC), the most likely attack paths to high value assets as well as suggested mitigations.

To support the needs of SOCCRATES, securiCAD has been enhanced with the capabilities to handle near real-time modelling and simulation of on-premise and cloud environments based on the data in the IMC and simulations based on cyber threat information to tailor the attacker capabilities to specific threat actors. Furthermore, the underlying attack logic has been extended to support the data types in the IMC as well as mapping it to the MITRE ATT&CK framework.

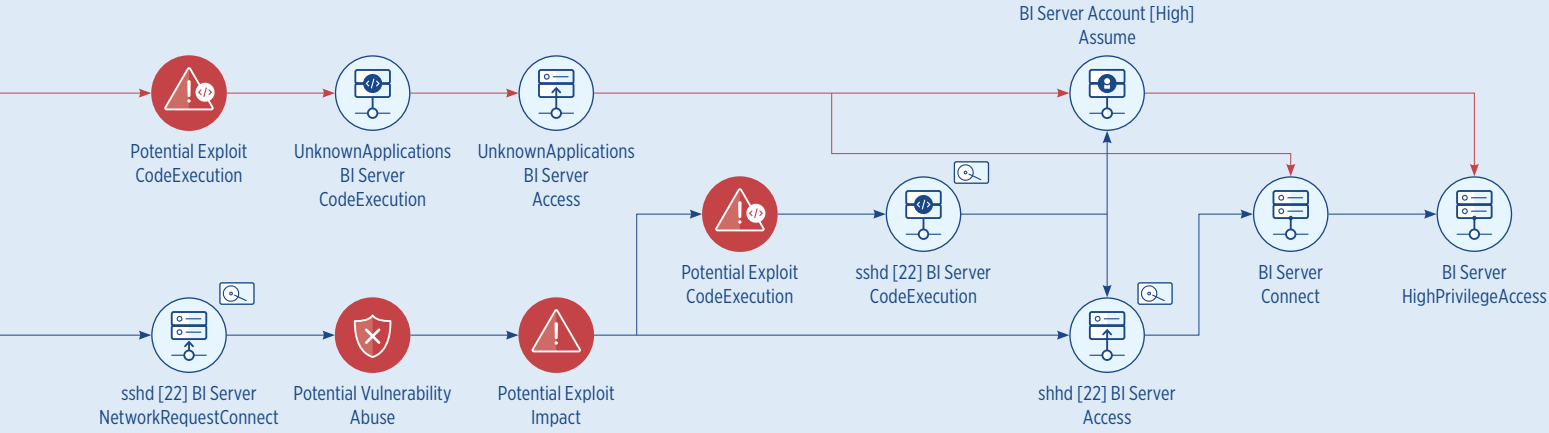

Research roadmap

Based on the identified future SOC/CSIRT needs and the work in the SOCCRATES project, the following research challenges are identified:

Data model for ICT infrastructures: in SOCCRATES we looked into standardised data models for ICT infrastructure modelling. As we could not find a good solution, we developed our own. We, however, identified that it is challenging to uniquely map information about the same asset from different data sources to a single object in the data model. We therefore suggest to start with research on this important topic. As a next step we propose to focus research on extending the data model with among others more details (e.g. integration of SBOM), security functions and domain specific data (e.g. from cloud, Operational Technology (OT) and IoT environments) and assess how to collect the data from data sources. As part of the research on extension of the data model, the need for the addition of the data has to be studied.

Database architecture: In SOCCRATES we use a graph database. Although this is a good basis, we identified some limitations. It is therefore proposed to do research in how to best setup a hybrid database architecture, with both graph and relational databases, for an infrastructure modeling component that is able to support all the requirements. Including ability to provide detailed historical infrastructure model information.

Visualisation: There is further research needed in the visualisation of infrastructure models combined with security status and event information. For instance the ability to plot attack graphs onto the infrastructure model, plot business impact information, plot observed attack steps and compromised host onto the infrastructure model, plot directly enforceable containment of compromised hosts on the infrastructure model or plot adversary capabilities after a successful attack step (e.g color all hosts that have adversary access after a privilege escalation). It is important to take into account the human limitations in processing and interpreting large amounts of visualized data and study, test and validate these visualization in real SOC environments.

IMPROVEMENT OF DETECTION CAPABILITY AND COVERAGE

Organizations rightfully put significant effort into ensuring their systems are secure against cyber-attacks by implementing cybersecurity programmes that ensure a strong defensive posture. However, there are inevitably compromises that need to be made, which result in the potential for adversaries to compromise a target environment. Additionally, adversaries change their tactics, techniques and procedures, possibly rendering security controls ineffective. Consequently, there is a strong need to introduce the capability to detect adversarial behaviour that is targeting an infrastructure. This detection can take many forms and use a variety of data sources, including that found on end-systems (e.g. log data) and from the network. For cyber-physical systems, detection can also be performed using data that is collected from the physical system, such as an electricity network.

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. There are also several applications of machine learning to perform multi-class classification in order to determine the type of adversarial behaviour observed (e.g. a DoS attack versus a port scan).

Current state

A major activity of a SOC is to respond to the alerts that are generated by detection systems. This includes addressing so-called false positives – alerts that incorrectly indicate the presence of an attack when none is present. In general, anomaly detection systems are understood to generate more false positives than those that use signatures for detection. This can represent a significant overhead for a SOC and can unnecessarily consume the time of analysts. Addressing false positives is an activity that is well-suited to being automated.

Building on an alert, a SOC must determine whether it is sufficiently significant, e.g. based on consequence, to be escalated and subject to further investigation. This can involve correlating (or associating) the alert with other information or alerts – this activity is normally supported with a Security Information and Event Management (SIEM) solution (but also many other supporting tools) and is performed by analysts with a greater level of expertise than those tasked with initial triage activities. Increasingly, so-called Security Orchestration, Automation and Response (SOAR) tools are used to support these activities via automation of repeatable playbooks and functionality that can be used to collate information from disparate internal and external sources. The aim of these tools and processes is to move from alerts that are generated by detection systems to actionable insights that can be used to respond effectively to an incident.

A major detection-related issue is concerned with *coverage*, which can relate to two main aspects: i) ensuring that prevailing adversary techniques and procedures can be detected; and ii) with the integration of Information Technology (IT) and Operational Technology (OT), whether the systems in both domains (IT and OT) are being sufficiently monitored. Regarding the former issue, Cyber Threat Intelligence (CTI) (see Chapter “Advanced use of Cyber Threat Intelligence” on page 38) can be used to determine known adversarial techniques and procedures of threat actors. A major challenge for a SOC, however, 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. The second coverage-related issue is concerned with ensuring detection capability across IT and OT systems, which are becoming increasingly integrated. 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. For example, detection systems exist on the market that can perform deep-packet inspection to detect the misuse of industrial protocols (e.g. Modbus, DNP3, OPC UA, etc.). However, arguably, endpoint monitoring and detection on OT devices and infrastructure is still relatively immature or absent. There are several reasons for this, including the capabilities of the devices involved and a general engineering objective to reduce the complexity of these systems, especially for safety-related applications. Moving beyond technical aspects, there are organizational challenges that must be addressed regarding cybersecurity detection, including the establishment and organization of an IT/OT SOC and how it interfaces with the operational parts of an organization.

Future SOC/CSIRT needs

Driven by several trends, there is arguably much more data that can be used to detect cyber-attacks. These trends include:

- An increased acceptance or awareness that monitoring for the purpose of cybersecurity is important, thus resulting in more data being collected;
- The availability of technologies and systems to ingest and store data at higher rates, coupled with ever-increasing long-term storage capacity;
- More devices being connected and instrumented for data collection, e.g. in the context of the Internet of Things and OT environments;
- A perennial increase in adversarial behaviour, which manifests in various forms of system and network data.

These large (or big) data represent a significant opportunity to detect cyber-attacks, if it can be applied usefully. However, there are some challenges that need to be addressed to capitalize on this opportunity. Perhaps the most apparent challenge 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? Systems and applications generate a wealth of logging information that could be used to detect a cyber-attack but it is not immediately apparent which logging information is noteworthy from a detection perspective and how an attack may manifest in the data.

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. Data scientists that apply machine learning to anomaly detection concern themselves with improving detection performance, which is normally presented using metrics such as accuracy or precision. The job of the cybersecurity data scientist is to improve these performance metrics, 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. However, even with excellent performance, based on very low false positive rates, due to the large amounts of data that is being operated on, an analyst can still find themselves overwhelmed with false positives to investigate and subsequently dismiss. This can be time consuming, tying up valuable human resources that could be applied usefully elsewhere.

As mentioned earlier, anomaly detection techniques that are based on machine learning algorithms have received significant research attention and can be found in many commercial products. The SOCCRATES project makes use of several anomaly detection techniques and systems, which aim to identify abnormal behaviour in network and system data and could indicate a cyber-attack. 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.

To address this issue, the events that are generated by anomaly detection systems need to be contextualized, to determine their root cause, e.g. in terms of specific adversarial behaviour that caused the anomaly. Currently, this contextualization is performed by SOC analysists, using expert knowledge and information derived for various tooling (e.g. asset management tools). Automated support for this activity – i.e. to determine the significance of a detected anomaly and determine whether it relates to an incident that is worthy of response – should help to improve the effectiveness of a SOC, as it aims to realize its KPIs.

Machine learning (ML) and artificial intelligence (AI) are finding increasing utility in SOC operations, to detect and classify adversarial behaviour and facilitate decision support systems, for example. However, just as defenders are exploring the benefits of AI and ML, so are the attackers. So-called *adversarial machine learning* can take many forms. An attacker's goal can include model theft and poisoning, for example, and subverting a model's output, in order to cause misclassification. Adversarial machine learning is attracting a lot of attention from the research community, in an effort to understand the nature and significance of the threat. Arguably, this is not a current threat that has significant operational impact; however, as 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.

DGA Detection and Classification

Related to the challenge of performing detection and classification on large datasets and addressing false positives, the SOCCRATES project has explored how to effectively detect and classify Domain Name System (DNS) names that are generated by malware using so-called Domain Generation Algorithms (DGAs). This work has leveraged the large datasets that are available to the project from The Shadowserver Foundation. The solution that has been developed includes a pipeline of detection and classification algorithms that aim to initially detect DGA-generated DNS names and subsequently classify suspected names according to the malware family that generated them. This later stage involves the use of both feature- and non-feature-based machine learning models, which aim to reduce the overall number of incorrect classifications.

The REASENS Framework

The REASENS Framework can be used to determine the root cause of events that are generated by distributed sensors, such as intrusion detection systems. In the SOCCRATES project, REASENS has been applied to inferring prevalent adversarial tactics, based on the alerts that are generated by the various anomaly detection components in the project. The goal in this work has been to take steps towards bridging the semantic gap between alerts that are created by anomaly detection systems and actionable information that can be used to initiate response playbooks. This is primarily achieved by encoding contextual information (e.g. about adversarial behaviour and the target environment) into the Evidential Network formalism that is at the core of REASENS Framework. In short, the system takes as input alerts from anomaly detection systems and yield a set of inferences and their likelihood the adversary tactic that caused them. This information is relayed to a SIEM solution, which can be used by a SOC to initiate specific playbooks that are related to the inferred tactic from REASENS.

Adversarial Machine Learning Study

Appreciating that adversarial machine learning has the potential to cause significant disruption to SOC operations and undermine the utility of intrusion detection systems that are based on machine learning algorithms, the SOCCRATES project has performed a study into this emerging topic. Specifically, the project has explored machine learning-based approaches to generate DGA-like DNS names that are not readily detected by the algorithms proposed in the project. This work has been leveraged in the project to improve the robustness of the models to this form of adversarial machine learning. Moreover, the project explored how so-called adversarial examples can be generated, which undermine the function of deep learning techniques that are used for network intrusion detection. This study indicates that, under certain circumstances, it is possible to generate examples that are incorrectly classified. For the moment, we argue that these findings do not result in any immediate operational risk – further research is needed to explore the practical application of these forms of adversarial behaviour.

Research roadmap

Based on the identified future SOC/CSIRT needs and the innovations that have been created in the SOCCRATES project, the following research milestones can be targeted:

Further refine and harden DGA detection techniques: Currently, there exist two distinct approaches to identifying domain names depending on the input data being used. In feature-based approaches, lexical or lexicographic information (n-grams, vowel ration etc.) is extracted from the domain names to create new data features. On the other hand, in feature-less approaches, a character encoding is applied to the domain name, and the resulting numerical sequence is used as input for a particular algorithm. In the SOCCRATES project, we explored these two directions in parallel, and reached the conclusion that integrating the two methods is the most promising direction to achieving even more efficient classification. On the medium-term, we deem worthwhile to explore multi-channel neural networks for the domain classification task, which can use both the raw domain

name and derived features as input. In this way, we could exploit the full potential of both methods: obtain explainable results due to the feature component, and extract complex patterns with deep learning. At the same time, it is also important to ensure that AI-based DGA detectors remain robust against adversarial examples. To this end, the current methods can be improved through training on enriched datasets, that also include self-made DGAs and extension to GAN (Generative Adversarial Networks) architectures. On the long term, should the appropriate data be collected and made available, DGA detectors could become more customized (trained for a specific language) or threat-centric (specialized in predicting and recognizing a specific malware family). This can better support analysts with investigating specific malware campaigns.

Further bridging the Anomaly Semantic Gap: As discussed earlier, there is arguably a semantic gap between the output from anomaly detection systems and the actionable information that is needed by SOC teams to respond to an incident. In the SOCCRATES project, we have explored the use of causal inference (using the REASENS Framework) to support the closing of this gap. Further research is required on this important topic: at the core of the REASENS Framework is an Evidential Network, which captures expert knowledge about the relationships between alerts and the adversarial techniques they describe; as a medium-term goal, the specification of these models could be improved using data-driven approaches, e.g. to populate the probabilities that certain alerts describe particular activities. Furthermore, based on the insights from REASENS or similar systems, there is the potential to introduce automation that supports the collation of relevant information that can support a SOC team when addressing an incident.

Modelling and hands-on exploration of adversarial ML: To date, there has been some uses of adversarial machine learning to support an attacker's objectives, such as to generate more realistic phishing emails. The research community has explored the potential threat associated with other forms of adversarial machine learning, such as adversarial examples. However, what is required is a well-established and comprehensive threat model that, for example, can be used to identify an adversary's in-

tent, motivation and capabilities in this context. This is necessary for organizations to understand the nature of the risk. Moreover, many of the studies on this topic are rather theoretical in nature, especially those that explore the nature of the threat in the cybersecurity domain (this is in contrast, for example, to studies in the image classification domain, in which this form of research is more advanced). What is needed, as a medium-to-long term milestone is to explore how such attacks could be orchestrated practically using suitable testbeds, for example. In addition, further insights are required on how to address these issues in the supply chain, addressing questions such as what standards and best practices (e.g. related to adversarial training of models) are required from vendors to mitigate these threats.

ADVANCED USE OF CYBER THREAT INTELLIGENCE

The Chapter "Actual machine-readable model of the infrastructure" (page 20) describes an infrastructure model: a model that enables you to know yourself and to know what you are defending. But in order to defend yourself effectively you also need to know your enemies. This is the goal of Cyber Threat Intelligence (CTI). Gartner defines threat intelligence as follows:

> Threat intelligence is evidence-based knowledge, including context, mechanisms, indicators, implications and actionable advice, about an existing or emerging menace or hazard to assets that can be used to inform decisions regarding the subject's response to that menace or hazard.[13]

In short, 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. The tactical and technical[14] subtypes are the most relevant for SOC and CSIRT needs.

- Tactical CTI is knowledge about adversary behaviour: their Tactics, Techniques and Procedures (TTPs) or Modus Operandi (MO). Changing behaviour is expensive and time consuming for an adversary, so tactical CTI has a long shelf life. The de facto industry resource for tactical CTI is MITRE's Adversarial Tactics, Techniques and Common Knowledge (ATT&CK)[15] framework, but specialised threat investigations firms also compile valuable reports about the campaigns and methods of particular actors.
- Technical CTI is knowledge about specific malware, tools or infrastructure. Observables such as file hashes, IP addresses and domain names can be contextualised, e.g. this IP address is assigned to a command and control (C2) server for the Trickbot malware. When contextualised, an observable becomes technical CTI: an Indicator of Compromise (IoC). Changing methods at the level of observables is cheap and fast for an adversary, so IoC-based detection is a never ending race to obtain new indicators and de-

fenders are always one step behind. Despite this limitation, sharing observables remains a valid approach toward detecting threats if consumed and processed in a timely fashion. They are also of immense value when used against historical records, typically to detect and reconstruct past intrusions.

The attractiveness of tactical CTI for attack detection was highlighted by David Bianco in his "Pyramid of Pain" article in 2013[16]. The core point from Bianco is that the ability to detect adversary behaviour lets defenders stay one step ahead of the attackers. Furthermore, this ability enables a defender to obtain new IoCs. When adversary behaviour is detected, the related IoCs can be automatically extracted and shared with other defenders in near real time.

Tactical threat intelligence also has another important use case: adversary emulation. This concept is most commonly used for "red team" activities, essentially the testing of an organisation's cyber defences through controlled intrusion attempts. Since we defend against real adversaries, we try to emulate real adversaries when testing. One example framework for the use of adversary emulation in security testing is TIBER-EU[17]. Adversary emulation can also be used by SOCs and CSIRTs when responding to an intrusion. If you know how an adversary behaves, then you can predict what their next steps might be and select an appropriate course of action (CoA) before the adversary is able to achieve its objectives. Adversary emulation is based on an adversary emulation plan (AEP), and adversary emulation plans are based on MITRE ATT&CK. MITRE Engenuity's Centre for Threat-Informed Defense has published several examples of AEPs[18]. AEPs consist of three main parts: a human readable summary, the operational flow, and a command-by-command detailed plan of how to carry out an attack. The operational flow chains techniques together to show the sequence of events during an intrusion, while the detailed plan shows exactly what the adversary does at each step.

In addition to detection and response, CTI can also be used to predict the behaviour of specific threat actors when they target a network. The more you know about an adversary, the better your ability to predict, prevent, detect, and respond to their actions.

Current state

IoCs are extensively used by SOCs and CSIRTs today, and their collection and processing (e.g. to compile a signature for applicable security monitoring systems) is largely automated. Threat feeds are downloaded, and block lists and alert lists are produced and distributed to security appliances, to enable real-time detection. Additionally, CSIRTs automate IoC sweeps on logs to find historical intrusion activity that was not detected when the activity took place. Nevertheless, there is still room for improvement in the process of timely sharing relevant IoCs.

Conversely, the application of tactical threat intelligence is still largely a manual process. Detection logic for adversary behaviour is created manually, and then mapped to a technique and/or tactic in ATT&CK. If a CSIRT or intel team would want to compile an AEP from such tactical intel, this would presently involve an excessive amount of time, effort and expertise. Tactical CTI is shared in a form that cannot be directly used to detect adversary behaviour: *a human must process the CTI first*.

The underlying reason for this state of affairs is today's technology and standards. 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 same issue is present in the standards for sharing CTI: Structured Threat Information eXpression (STIX)[19] and the Malware Information Sharing Platform (MISP)[20] formats.

Techniques are in most cases too general to be used to detect adversary behaviour. MITRE even states this in the ATT&CK design and philosophy paper:

> The techniques within ATT&CK may have many procedures for how an adversary could implement them — and because adversaries are always changing, it is difficult to know what all those procedures are in advance. That makes discussing coverage of a technique tough, especially when some ways of detecting behavior rely on individual procedures and some may span multiple procedures or even an entire technique.
> [...]
> It is important to always review the threat intelligence on what techniques, sub-techniques, and procedures adversaries have used to understand the details and how variations might affect how you determine coverage. Anyone mapping to ATT&CK should be able to explain the procedures they cover.[21]

The key to detecting adversary behaviour is procedures. In general, a procedure can be detected and mapped to a technique, while a technique cannot be directly detected. 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. In order to facilitate automated detection of adversary behaviour, this should be addressed.

Future SOC/CSIRT needs

SOCs and CSIRTs are flooded with data and information. The amount of new daily data in threat feeds makes it infeasible to manually assess every new entry. Furthermore, the data quality of IoCs is highly variable, and low quality IoCs can cause large amounts of false positive alerts for a SOC, further compounding the problem. New innovations in technical CTI are needed to solve these issues.
We need new methods to contextualise IoCs, in order for defenders to be able to better understand what kind of information they must focus on. Furthermore, we need to ensure that IoCs are relevant to the receiver, and delivered in a timely fashion, to avoid defenders exhausting themselves chasing non-incidents. Improvements that are needed include:

- New methods and tools to improve the quality of automated IOC feeds
- Sharing of “sightings” – observations from other members of the community
- Reporting of “known good” observables, e.g. IP addresses used for scanning by the security community.

Finally, technical CTI is still sometimes shared in ways that hinder automation, such as unstructured human readable text. Initiatives to raise awareness of existing standards and formats for sharing technical CTI could help address this issue.

New innovations in tactical CTI are also needed. Tactical CTI is still largely a manual process. Humans read and process high level reports, and then try to derive the knowledge needed for detection. In order to make tactical CTI more useful to defenders, we need automation. The first step towards automation is to describe adversary behaviour in a machine readable format, at a level of detail that can be used to detect that behaviour. MITRE ATT&CK is too high level for this purpose and does not clarify which sequence of techniques represents the behaviour of an actual real world adversary, so a more detailed, low level description is needed. If we can share machine readable descriptions of adversary behaviour at a sufficient level of detail, then the next innovation would be to develop methods and tools for defenders to automatically process and use this information for detection and attribution.

SOCCRATES INNOVATIONS

DGA Detective

SOCCRATES developed a machine learning classifier that detects domain names generated by a Domain Generation Algorithm (DGA). The classifier enables the detection of several procedures that implement the Dynamic Resolution: Domain Generation Algorithms sub-technique from ATT&CK, and makes it possible to collect and share new IoCs in near real time. The classifier has been adapted to process large sets of domains quickly. By coupling the classifier with Shadowserver's ingests and sandbox executions of hundreds of thousands unique malware samples, potential new DGA's can be shared through Shadowserver's API in a timely manner.

Adversary Emulation Planner

SOCCRATES developed a new method to generate a sequence of attack stages for an adversary group, campaign or incident. The method is based on semantic modelling of technique dependencies and enables automated generation of attack stages from a list of techniques. This method was implemented in an Open Source tool, published on Github. The sequence of attack stages is used as input to the ADG Analyser to enable more realistic simulations of adversary behaviour.

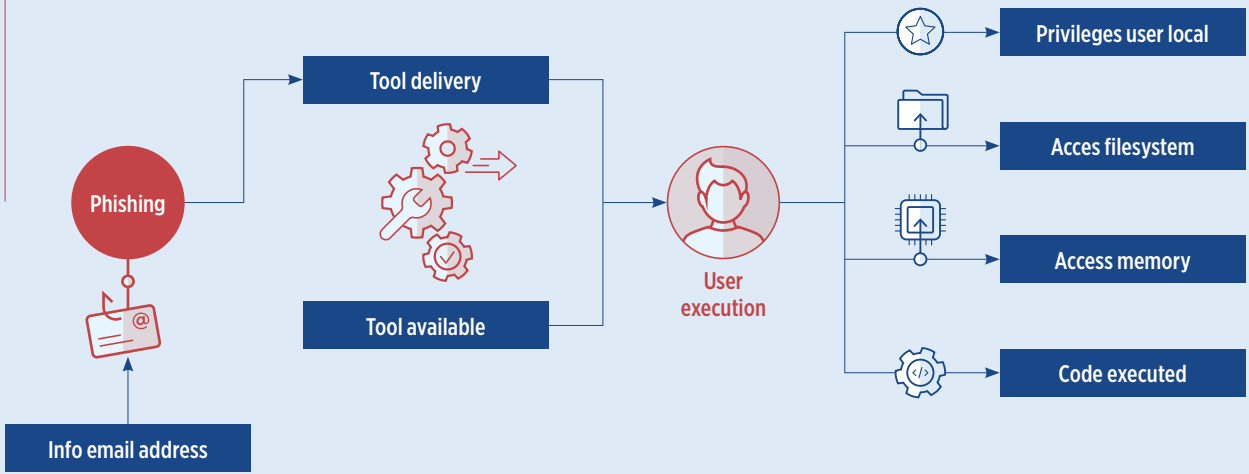

Research roadmap

We propose the following medium and long-term research milestones:

Technical CTI: The SOCCRATES project applied machine learning methods to detect Domain Generation Algorithms, implemented in the DGA Detective tool. Our experience indicates that methods from artificial intelligence and machine learning could have wider use cases and address the needs from the previous section: the application of machine learning classifiers to improve the quality of automated IOC feeds, contextualise IOCs, and discover trends in large data sets. We suggest quality improvements and trend discovery as mid-term research milestones, and IOC contextualisation as a long-term milestone.

Tactical CTI: The AEPs generated by the Adversary Emulation Planner tool developed in SOCCRATES can be used by the ADG Analyser to simulate an adversary, but they do not have a sufficient level of detail to be used for the detection of adversary behaviour. Our view is that the most significant mid-term research challenge in tactical CTI is to develop a format for a machine readable description of adversary behaviour that can be used for detection and attribution. Long-term, such a format can be a starting point to develop tools to automatically transform the description of adversary behaviour to the different domain specific languages (DSLs) that are used for detection.

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.

Assessing the operational impact allows the SOC analyst to get a more global view on the repercussions on the business side, on top of the usual, more detailed view on the impacted assets. By identifying which company activities will cease to work properly, it becomes possible to warn the impacted departments, therefore allowing for smoother and faster reorganization of the business processes until full recovery from the attack damage.

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. This means that modelling the business processes is not enough in itself, and must be followed by some adaptation into a mathematical model adapted for impact computation.

Current state

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. Sometimes this is extended with a flag for privacy (i.e. the host contains personal data), and a flag to indicate whether the host provides a service that is part of critical infrastructure. Some organisations have this information accessible for the SOC and CSIRT via the AD or a CMDB, but sometimes it is an excel sheet. More specific analysis of business impact is done manually and in collaboration with the business owner of the particular system. The latter 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.

Currently, the main approach to obtain a model of the business processes and functions in the company is to rely on Business Process Model Notation (BPMN). A BPMN model makes it possible to draw a precise description of the choreography of business functions within a business process, while identifying dependencies to data objects, such as data stores. However, as of now, there is no reliable automated way to obtain such models, due to the lack of business oriented data sources that can be easily processed. Moreover, BPMN models cannot directly be exploited by a SOC analyst, since aside from data stores, these models do not contain any link or dependency to assets from the ICT infrastructure, which are the direct targets of attacks. Business processes thus 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.

Future SOC/CSIRT needs

The main needs for real-time business impact assessment can be summed up as follows: on one hand, it would be extremely useful to at least partially automate the identification of the company's business functions and business processes, as well as their dependencies.

Including the dependencies with the assets from the infrastructure that directly support business functions. On the other hand, the SOC analyst must be able to compute relevant metrics to perform the business impact analysis.

Regarding the first aspect, the main challenge to overcome is the lack of automation solutions in the state of the art. Methodologies to elaborate 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.

Regarding the second aspect, the challenge is to design a scalable mathematical model able to compute various metrics in real time, all while taking into account both the specificities of the infrastructure, such as asset redundancy and interdependencies, and the specificities of the attack, in terms of confidentiality/integrity/availability impact. 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. Business impact analysis indicators can be obtained by computing metrics such as the impact probability and intensity on business entities, the downtime tolerances, all while considering values such as Recovery Time Objective (RTO) and Recovery Point Objective (RPO). To make matters even more complicated, 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. When we do dynamic risk assessment it is important to include such temporal aspects. Ultimately, the business impact analysis should be fully integrated as part of business continuity and resiliency planning.

Business Impact Analyser

In the context of the SOCCRATES project, a Business Impact Analyser component was implemented, building solely on theoretical research work. This component first builds a graph representation of both the infrastructure assets and the business logic tied together, therefore allowing the SOC operator to evaluate how an attack targeting specific assets can propagate its impact up to the business functions and processes of the company. It computes two values, which can be displayed on a user-friendly graphical interface: the impact probability on each business entity, and the critical time, under which the attack impact on the company remains limited. With these two values, the SOC operator is able to assess the emergency of a situation, and is able to directly contact the company departments associated to highest impact probability values.

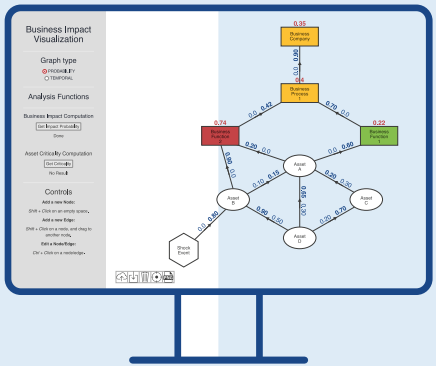

The main challenge of implementing this component in SOC/CSIRT environments is related to the manual work needed beforehand to design a graph of business entities, and evaluating the dependencies and downtime tolerance. However, once this model is designed, the remaining operations of the Business Impact Analyser can be fully automated.

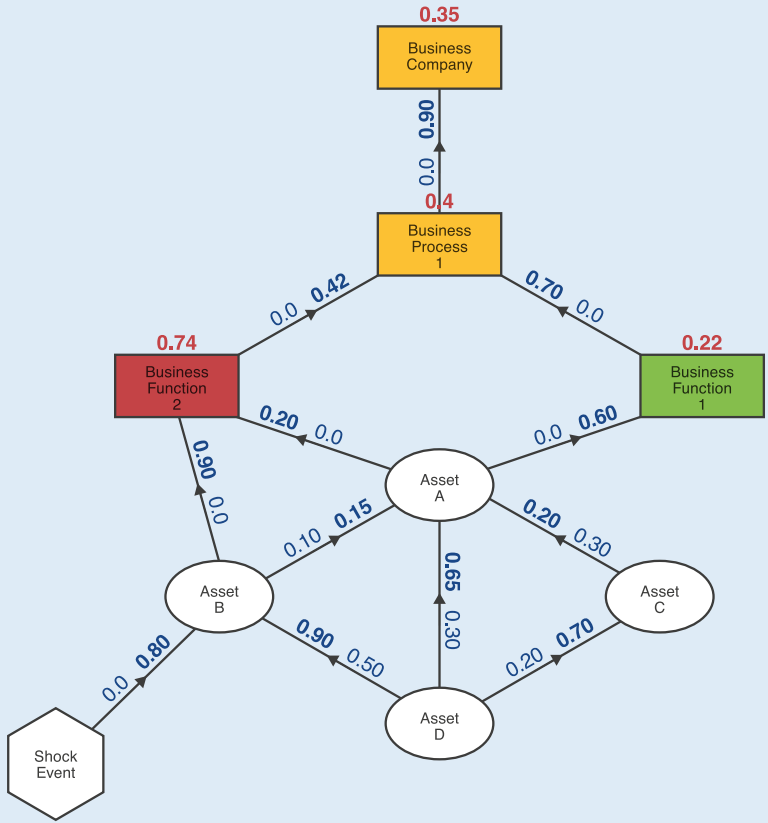

Research roadmap

Real time business impact assessment appears to be a burning research issue, as it has great potential improve the quality of responses to threats and ongoing attacks. Two specific research topics can be highlighted, echoing the main needs expressed in the previous section:

- elaborate theoretical frameworks based on graph theory adapted to business impact analysis
- elaborate a less abstract approach to model business activities, with an explicit link to main infrastructure assets

The first topic can take advantage of existing graphical models such as Bayesian networks, but such models are often subject to theoretical constraints that do not hold true for business impact analysis. Likewise, further research work is necessary on the propagation models of confidentiality/integrity/availability attack impact, since impact on business entities cannot be thought in the same terms as impact on assets (for example, an attack with high confidentiality impact does not induce any ob-

vious business impact, since besides a financial loss, the company can maintain the production without any service disturbance).

Regarding the second topic, the main challenge comes from the fact that a substantial evolution from the well-established BPMN model would be required. The main objective would be to design a model that allows clear identification of the assets involved in business functions, all while introducing new data sources that would make automation possible with respect to model building and dependency quantification.

With advances in both of these topics, it will be possible to achieve the development and deployment of mature and user-friendly business impact analysers in SOC/CSIRT environments with a high level of automation from model building to impact assessment, with much more limited risks of estimation errors.

RECOMMEND COA GENERATION

Within SOCCRATES different types of courses of action have been identified. On the one hand, courses of action to improve the cyber resilience of an infrastructure after changes in the infrastructure, discovery of new vulnerability or new cyber threat intelligence. On the other hand, courses of action to respond to ongoing attacks. The latter are sometimes referred to as Containment CoAs and have to be *directly enforceable* to effectively contain compromised hosts, prevent adversaries from gaining further access and eradicate root cause vulnerabilities that allowed the adversaries to gain initial access. The speed of deployment is less of an issue for the CoAs to improve the cyber resilience. This is also the case for post-incident improvements from the lessons learned from an cyber security incident.

To be able to automatically suggest optimal courses of actions for improving security in ICT infrastructures at large is an advanced, and to some extent utopic, capability. In the SOCCRATES vision this capability is built around the fundamental idea that we 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 depend both on how much "raw data" from the ICT infrastructure is available and how well the model language can distinguish which actions are efficient security improvements, given different states of the infrastructure. In SOCCRATES the raw data is collected in the Infrastructure Model Component (IMC) (see page 24) and the model language is hosted in the Attack Graph Generator (ADG) component. The ADG converts the architectural data from the IMC into possible attack graphs representing the possible attack vectors over the ICT infrastructure. (The graphs could be considered theoretically simulated penetration tests covering a myriad of possible vectors.) The attack graph is the model that is used for doing analysis of security and in the end to generate suggested defence Courses of Actions (CoA).

Once we have a model at hand the next thing we notice is that the security quest is multifaceted and needs to be defined more precisely in different scenarios. What exactly do we mean by security in this context? The best action to take from a defence point of view would likely not be the same if the infrastructure was under confirmed ongoing attack as if we wanted to launch a project to improve the resilience of the infrastructure with respect to some observed relevant cyber threat intelligence. Perhaps we see this most clearly in the ever increasing speed of DevOps cycles. For many companies, in particular those operating in cloud infrastructures, their ICT infrastructures are highly volatile where the software, network connections, and user accounts are changing constantly. To be able to promptly identify configurations that weakens the overall security posture is clearly very important. (And also difficult, as for instance the many stories of accidentally exposed data in cloud containers have witnessed about the last couple of years.)

Starting with examining the preventive measure optimization, this challenge boils down to a matter of risk management related to how to most effectively secure the valuable assets located in the ICT infrastructure. Thus, the same ICT infrastructure can have a very different security posture, depending on what we consider valuable. In addition to understanding the system configuration the system resilience also depend on the threat actors, where they are located and what kind of skills they possess. So how to secure the system depend on the considered threat scenario and can thus not be unambiguously answered with a single answer. Instead we have to weigh and aggregate multiple assumptions made in various scenarios. In terms of the security analysis made in SOCCRATES, the threat scenario is defined by one (or several) starting point(s) and by one (or several) target point(s) and the security of the infrastructure is assessed by 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 according to

our definition, means that security improves. In order to aid the defender, the attack graphs generated from the models also contain defence actions. So, 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 (to avoid the word optimal) defence action combinations, sparing the defender the work of trial-and-error simulations of testing different defence strategy hypotheses.

Now, if we move over to the incident response scenario this turns out to be more complicated. Of course, given an observed breach we might want to prevent further penetration and compromises of valuable assets in the infrastructure. The actions for this follows the above described logic, though the time horizon is likely much shorter and thus the available action space smaller. However, security actions now also include containment, eradication and recovery. These dimensions of security are not appropriate to measure with (and optimize with respect to) TTC. This is not to say that attack graphs are not useful in this context, on the contrary, some ideas about this will be discussed below. Nevertheless, automated generation of these Containment CoAs strongly depend on the capabilities already present in the ICT infrastructure for containment and eradication, as they have to be directly enforceable. Generating Containment CoAs may therefore follow a more direct process of identifying and comparing the options.

Current state

As of today SOCs have a focus on operational and tactic security actions, i.e. acting on ongoing attacks or some imminent threat. Introducing the ADG component into the SOC represents an ambition to make it more strategic in nature. We want the SOC analysts to become more engaged in the question of building a long term resilient infrastructure. To the best of our knowledge support for this type of work is scarce. Even though some SOCs are already moving in this direction, the work is largely

driven and dependent on the expertise and experience of some individuals, tacit in their heads. Also, the challenge of valuing information about ongoing attacks (to avoid alert fatigue) lacks support. Understanding and estimating how severe a compromise is in terms of what this means in subsequent attack vectors is also to a large extent dependent on tacit knowledge. The current state of automatic security mitigation is perhaps most advanced in the area of intrusion detection and prevention systems where actions can be encoded and set to trigger based on some observed behaviour in the product. In general the triggering logic is however quite local in nature. I does not take into account the configuration of a larger ICT infrastructure and is thus to some extent suboptimizing its suggestions. Also the solutions are clearly built for reactive measures rather than preventive ones.

Much of the solutions developed within the SOCCRATES project is advocating and propagating the vision that SOCs need a better understanding of the own ICT infrastructure they are tasked to protect. Reflecting at a distance it is somewhat surprising that the tool support in SOCs are fairly advanced for "outward-looking" activities and fairly rudimentary for "inward-looking" activities. One reason for this is the fact that ICT asset management is still largely an unsolved problem from a practical point of view. Another explanation is of course that outward-looking has served well for reactive security measures but with higher demand on also acting proactively this lack of decision support is a challenge. For instance, as argued above, the SOC needs to identify vulnerabilities and weaknesses as they emerge. But vulnerabilities do oftentimes not make any noise when they appear. This puts high demands on the SOC to have detailed and up to date data of the infrastructure.

Future SOC/CSIRT needs

As already mentioned, one of the biggest challenges in realizing and evolving the SOCCRATES vision in full scale is the challenge of asset management where all the components in an ICT infrastructure needs to be "discovered". It falls back on the old expression "how do you control what you cannot see?". 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.

As the SOCCRATES project has demonstrated it is possible to connect intrusion detection sensors to the attack graphs used as the fundament for preventive security assessment. However, the mapping of the detection space and the security analysis space needs to be addressed in a more holistic way to fully support the SOC and CSIRT. Already within the project we can demonstrate that the attack simulations can also be used "in reverse" to support CSIRT activities. 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 indicators of which attack steps that have detection capabilities connected to them, including their quality (in terms of expected true positive and negative rates) for each individual sensor. We also acknowledge that such "observation layer" to the attack graph does not only include intrusion detection systems, but any type of forensic trace. This integration must both be concretized conceptually, how does for instance multiple inconsistent observations update attack simulations and how can attackers wipe their own tracks, as well as technically, how are forensic capabilities observed and transformed into a model - perhaps by extending the SOCCRATES IMC.

Yet another domain where additional work is needed for realizing the long-term SOCCRATES vision is how to visualize and contextualize the CoAs depending on different threat scenarios and use cases.

CoA Generator

In SOCCRATES we have devised a prototype that automatically suggests defense Courses of Actions (CoA) to the SOC analyst. The CoA is generated from a model of the ICT infrastructure (generated by the "Attack-Defense Graph" generator component in turn built from data from the "Infrastructure Model Component") that the SOC is set to protect, as well as additional assumptions about:

- what are the valuable assets in the ICT infrastructure,
- the attacker location (could be a generic threat origin such as the internet as well as an observed location inside the ICT infrastructure),
- the expected cost of implementing each respective defense available in the defender's portfolio of actions and controls, and
- a defense budget available for the defender (in the scenario).

Response Planner

The Response Planner developed in SOCCRATES performs two tasks: it generates containment CoAs, and it evaluates and ranks CoAs based on either business impact or financial impact, generating CACAO playbooks for the CoAs that will be selected by the SOC operator. The objective of the containment CoAs is to stop the progression of an ongoing attack, and focuses on mitigations such as host isolation and traffic filtering. For the other, more long-term CoAs, the ranking is done based on the Return on Response Investment (RORI) metric, that aims at evaluating if the CoA is financially worth being deployed, and which CoA appears to be optimal according to this criterion.

Research roadmap

In terms of developing the CoA capability per se we see a number of directions for future work (some of them already partly initiated within the project). First and foremost there is an imminent need to expand the CoA generator to also encompass structural changes of models. Technically it is possible to encode all (most) potential defence mechanisms as defence action nodes in attack graphs (according to the standard attack/defence graph formalism). However, practically and more intuitively, we oftentimes would prefer to capture defence actions as model changes. For instance, we want to introduce a new network segment, a new account, remove a privilege and similar which most intuitively corresponds to adding and removing objects in a model (rather than what perhaps can be seen as "tuning a property" when enabling a defence node). Initial work of formalizing structural changes have been done within SOCCRATES by transforming the ADG models into property graphs in graph databases and using graph query languages to specify and identify "security anti-patterns" that can be added to the defender's potential action space. With a good library of anti-patterns and their corresponding mitigation-patterns, the total design space can become quite close to the defender's actual action space. Of course the CoA generator algorithm has to be updated to manage the then dynamically changing attack graphs that the pattern swaps introduce.

In SOCCRATES we have developed a traditional graph-based search algorithm to realize the CoA. Another timely approach to this challenge is to deploy machine learning, and in particular we believe that reinforcement learning is well suited for this task. If we see the attack graphs as the training ground for the defender where we also have the defined set of defensive actions, in terms of both mitigation patterns and atomic defence actions, this can be used for defender training. In order to do this, assets in the model will be associated with rewards and penalties. Rewards are provided to the defender for keeping systems uncompromised, available, and configurations at a low cost, and penalties are provided when the attacker succeeds to reach sensitive targets in the model, i.e. attack steps in the graph. By training defenders through such self-play scenarios, it is our belief that

it is possible to develop good defence strategies that can replace the CoA generator algorithm developed in SOCCRATES. By varying penalties, rewards, and the available action space for the defender, as well as the ICT infrastructure architecture, one can seek to develop both preventive and reactive defence strategies. Some initial experiments with good results of this setup have been conducted by KTH using OpenAI's gym. We also note that a similar setup can be used to train the attacker, which in turn can be used to train the defender. In the meantime attacker behaviour can also be modelled by applying observed attacker behaviour or any synthetically encoded logic with a varying degree of rationality, replacing the fully rational attacker that implements shortest path-based simulations that have been used in SOCCRATES.

Furthermore, developing the aforementioned "observational layer" to the attack graph is not only a matter of development but also partly on the research roadmap. With an observational layer in place it is also possible to deploy the artificially intelligent defender in an actual system environment to fight attackers.

Finally, even though this outlook for future research and development stems from the technology and research frameworks used in SOCCRATES, where system models and attack graphs are central, we lastly acknowledge that the machine learning community in general is trying to remove any intermediary models (other than the neural networks of the AI actors) in much of their applications, so full scale AI-based defence will likely not make use of attack graphs. However as there is so much potential variance in ICT infrastructures and so little available data to learn from, it is our belief that the attack graphs has a strong role to play for the short- and mid-term horizon.

AUTOMATION AND ORCHESTRATION OF SECURITY OPERATIONS

It is without doubt that although the decision and reasoning skills of humans are superior to those of any computer system, the capabilities to perform repetitive tasks in a faster way and with fewer and more controlled errors are greater in machines. This fact leads us to affirm that to improve the response time and reduce the MTTR (Mean Time to Recovery) values in the SOC/CSIRTs, it is necessary to increase the usage and number of processes automated and consequently this implies the need of an enhanced orchestration. This is even more important in a dynamic context where cyber-attacks vectors are constantly changing and where we have to consider not only the execution of simple and predefined if-then scripts, but also automated analysis, threat hunting, generation of Course of Actions adapted to the current attack status, the assets affected, the objective of the adversary, and so on. In any case, we have to be aware of the many challenges that arise for automated response, in particular in SOC/CSIRT environments, and that the human intervention cannot be totally suppressed and needs to be integrated in the incident response workflows.

Current state

Around 2015 technology started to emerge that we now called Security Orchestration, Automation and Response (SOAR) solutions. Initially these solutions were developed out of convergence from three different technologies:

- Security Incident Response Platform: a platform for security incident case management and coordination of incident response. These platforms provided functionality to define incident response workflows, or playbooks.

- Threat Intelligence Platforms: a platform for collection, assessment and sharing of cyber threat intelligence. The platforms started to integrate workflow technology to manage the cyber threat intelligence process.
- Security Orchestration and Automation: orchestration refers to integrating different tools/technologies in a coordinated way, usually through workflows, such that they can function together. Automation is the process of machines executing tasks that were traditionally performed by humans. 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.

The market for SOAR solutions, has matured since. Gartner published a report in 2020 (Claudio Neiva, 2020) comparing different SOAR solutions based on the following features: their interoperability, the security operations covered, technologies included, their capability to create workflows, if they involve the analysts and the type of deployment they use. And comparative lists of commercial SOAR solutions or vendors can easily be found on the Internet[22]. Also some open source SOAR solutions have been developed an published, such as Shuffle[23].

According to Swinlane 2021 SOAR Report (Cybersecurity Insider, 2021) on a survey done to cybersecurity professionals, practically all of these organizations (92%) considered Security Orchestration Automation and Response (SOAR) solutions useful and more than half of them reported benefits using SOAR, identifying as the most valuable SOAR capabilities: automation (71%), orchestration (62%) and response (60%).

Apart from the SOAR products, also standardisation to support automation and orchestration of security operations has started. The Organization for the Advancement of Structured Information Standards (OASIS) has created several technical committees on the subject:

- Collaborative Automated Course of Action Operations (CACAO)[24] is a Technical Committee of OASIS that is working on specification for documenting playbooks for cybersecurity opera-tions and sharing these across organisational boundaries.
- Open Command and Control (OpenC2)[25] is a concise and extensible language to enable ma-chine-to-machine communications for purposes of command and control of cyber defence components, subsystems and/or systems in a manner that is agnostic of the underlying prod-ucts, technologies, transport mechanisms or other aspects of the implementation.

SOCCRATES has adopted both these standardisation activities for the SOCCRATES platform.

The current state of security tools at many organisations can best be described as a plethora of disparate products from different vendors or sources. Typically these tools have their own data representation, interpretation mechanisms and no standardized data exchange capabilities, making it challenging for SOC analysts to create a holistic picture of their organizational security posture (C. Islam, 2019) (Johnson Kinyua, 2020). 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 will remain challenging.

Another challenge that organisations should realise when deploying a SOAR solution is the fact that these tools require a significant amount of manual tuning and playbook definition. These playbooks are security event specific. With the continuous changing threat landscape organisation will need to keep adding and tuning playbooks and automation functions. In addition, the applicability of current generation of SOAR tools for effectively analysing and responding to the so called unknown-unknowns attacks, remains to be seen. Indeed this is not a limitation of the SOAR solution, but mainly of the lack of automated reasoning and courses of action generation, as we have demonstrated in SOCCRATES.

Last but not least, with the increasing number of security events and alerts an organization has to cope with, it remains to be seen how effective current SOAR solutions can handle large amount of simultaneous triggers. Note that many of the simultaneous triggers may be related to the same security event. It would thus be unwise if a SOAR solution would start individual response actions. Handling of multiple simultaneous triggers and running different playbooks for related security events needs to be studied further.

Future SOC/CSIRT needs

To reduce the mean time to response and the overall Mean Time To Recovery (MTTR), SOC and CSIRT will need to use automation and orchestration of security operations. Although the SOAR market seems to be maturing, the research on SOAR solutions has only recently started (C. Islam, 2019) (Johnson Kinyua, 2020). There remains much foundational research work on the topic of security automation and orchestration.

We foresee the following needs for the near future:

- Support to ease the introduction and deployment of SOAR tools in SOC and CSIRTs tool chain. This includes management and support for integration of numerous security products and tools. Sharing of generic playbooks that can easily be tuned and implemented within an organisation's own SOC and ICT infrastructure.
- With the increasing number of security events and alerts, a key goal is reduction of false positive and prioritize the alerts. SOAR tools can play a role in automatically enrich and collect additional information in order to reduce false positives and prioritize alerts. We do have to be careful, however, in managing the number of playbooks we trigger and simultaneously handle. Particularly if the multiple simultaneous triggering events are related to the same security event or incident.
- In SOCCRATES we have tools that dynamically generate response actions, or CoAs. These generated response actions will typically be too abstract and must be translated into specific reconfiguration

commands for one or more security actuators, and structured in a step-by-step machine-readable dynamic security playbook. Automated playbook generation is an approach that uses a library of modular templates that can be automatically combined and configured to create an executable playbook for a specific high-level response action for a specific environment.

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. This is particularly also the case for SOC and CSIRT analysts. As we introduce more automation and orchestration of security operations, we have to be aware that the role of the SOC and CSIRT analyst will change in the coming years. Do we still need to train more people for Tier 1 analysts, as we will automate the triage in the near future? Do we need to train more security automation and orchestration tool engineers, that can efficiently deploy and manage these security automation tools?

Another aspect is the role, or the level of involvement, of the human in the security automation and automated response. When it comes to security automation many organisations do not like to give full control to an automated, or even autonomous, system. Most organisation thus prefer to have a human-in-the-loop and explicitly authorise reconfiguration actions before execution. In the coming years this will, however, change. Automation is much faster than human comprehension. The human should not unnecessarily slow down the automation. In the coming years, however, there will be tasks that are difficult to automate. The human SOC analyst will thus remain necessary and have to collaborate with the automated security systems. In (Yan Shoshitaishvili, 2017) the interplay between human and automation is explored for automated vulnerability discovery. The paper points to the shift from tool-assisted human-centred analysis to human-assisted tool-centred analysis. Such a shift will likely also come to SOC and CSIRT operations.

SOCCRATES Orchestrator and Integration Engine

SOCCRATES has implemented an Orchestrator & Integration Engine (OIE) component based on the open source tools Activiti and Cortex. The OIE is the core of the SOCCRATES Platform. The five use cases defined in the project (see previous innovation point) have been modelled in BPMN and they have been deployed (and the services required implemented) in the SOCCRATES Orchestrator & Integration Engine.

The different use cases are triggered in the SOCCRATES orchestrator through the Activiti REST APIs, invoked by different components depending on the use case (e.g. by the SIEM in the case of a response on a detected ongoing attack, or by the IMC tool if it is detected the deployment of a new system in the infrastructure). The interaction of the SOCCRATES orchestrator with the rest of the SOCCRATES components is performed through the usage of different Cortex analyzers/responders.

The architecture of the SOCCRATES Orchestrator and Integration Engine allows it can be easily extended in the future, mainly in two ways: (i) including new features/ capabilities by the integration of additional workflows or extending the current use cases, (ii) integrating new tools/ components with additional analysers and responders using a common interface.

A web interface has been implemented for the SOCCRATES platform to show information about the status of the different workflows executed in the Orchestrator & Integration Engine and integrate the user interfaces of the different tools in a single portal. Additionally, it also allows the user to interact with the orchestrator for decision making offering him an overview of the outcome from the integrated tools as result of their analysis.

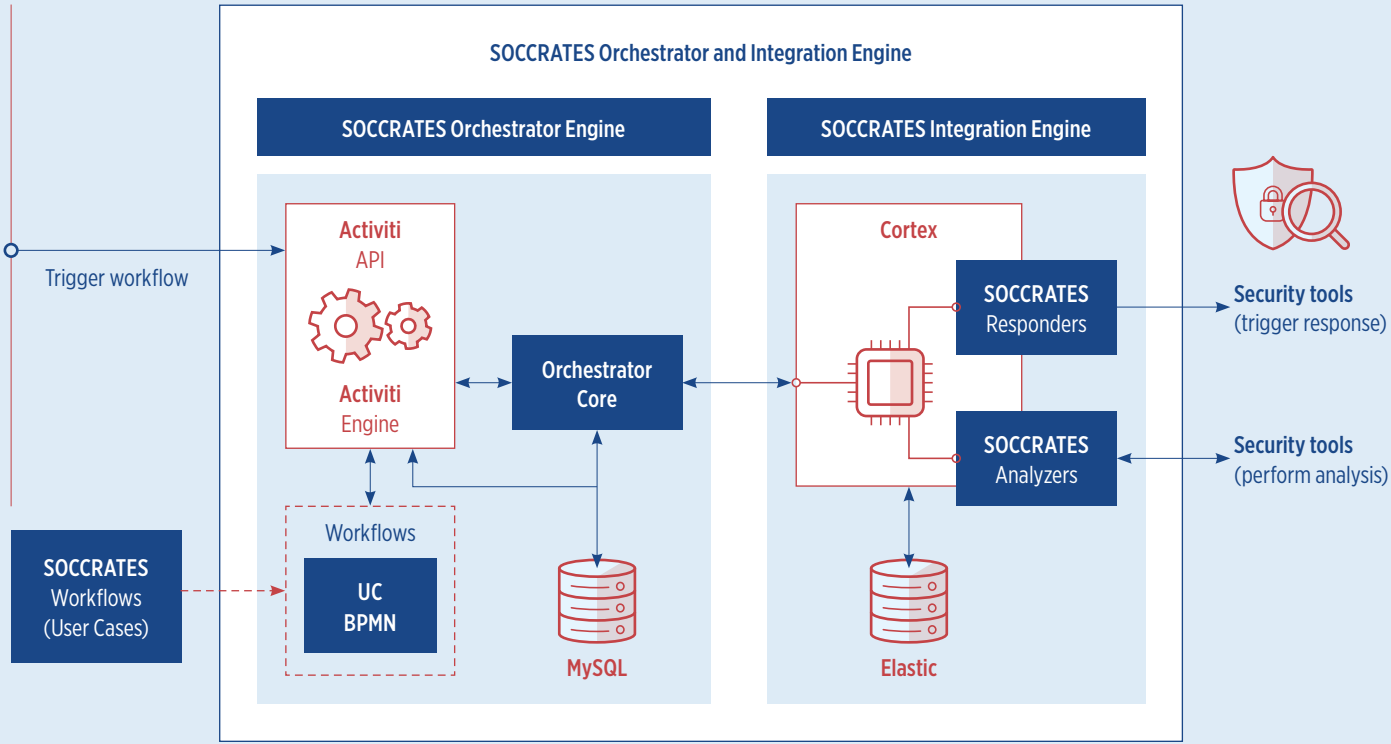
SOCCRATES Orchestrator and Integration Engine
SOCCRATES Orchestrator Engine
SOCCRATES Integration Engine
Trigger workflow
Activiti
API
Activiti
Engine
Orchestrator
Core
Workflows
UC
BPMN
MySQL
Cortex
SOCCRATES
Responders
SOCCRATES
Analyzers
Elastic
SOCCRATES
Workflows
(User Cases)
Security tools
(trigger response)
Security tools
(perform analysis)

Research roadmap

Based on the identified future SOC/CSIRT needs and the work in the SOCCRATES project, the following research challenges are identified:

Orchestration and workflows: SOCCRATES has developed its own Orchestration an Integration Engine out of two open source components. Additional research and development is needed to further improve and extend the capabilities. This includes updating the workflows to make them valid for a wider variety of potentially target SOC/CSIRT infrastructure models (hybrid, cloud based, virtualized), enable better tuning of the workflow to a specific environment, improve visualisation, and extend the integration with other open-source tools commonly used in SOC/CSIRT environments. Furthermore, research is needed into simultaneous triggering of events that are related to the same security event or incident. For the long-term it is interesting to do research on the integration of AI into security orchestration that can learn from the decision-making process and make suggestions for future handling of security events.

Automated Response: In this document, we highlighted the need for speed in incident response when attacks are becoming more automated. To reduce the mean time to detect, respond and recover, security teams have to increase the level of automation in incident response. This includes automated response by automatically reconfiguring security functions, network systems, etcetera. In many organisations and certain domains, automated response will not be accepted easily. To understand the concerns and provide automated response capabilities that can safely be introduced, more research is needed. For instance, to address questions like: What automated response actions are acceptable? What are the conditions under which automated response actions are allowed? How to integrate a human in the loop? How to integrate automated response into the IT service management (e.g. ITIL)? Furthermore, research and development work is needed in integration of security tools and standardisation of interfaces, protocols and data formats to enable automated response, e.g. openC2 from OASIS and Interface to Network Security Functions (i2nsf) from IETF.

Automated Playbook Generation (for CoA): as described above, for the translation from abstract recommended CoAs to a specific ICT environment, automated playbook generation is needed. Research is needed on how to automatically create playbooks from e.g. a library of modular templates.

Automated playbook generation methods may also be useful for more dynamic analysis and response planning of security events. By having the ability to build playbook steps on the fly, an AI agent may be able to dynamically create and modify a playbook for the security analysis of so called unknown-unknown attacks. In general, research is needed in how security automation and orchestration can become more flexible and dynamic, and less pre-defined.

Human-Machine interaction in SOC and CSIRT operations: human-machine interaction in security operations will become an important topic for research in the coming years. Research questions include: How to best interact with the human SOC and CSIRT analysist to enable efficient and effective security decision making? How will the task of SOC and CSIRT analysist change over time? What will the role be of a SOC and CSIRT analysist be in human-assisted tool-centred analysis? How to ensure that SOC and CSIRT analysist can still perform when automated security tools cannot be used (relying too much on automation)? Last but not least, how to train the human SOC and CSIRT analysist to work with the highly automated security tools?

Data Standardization: Investigate the feasibility of the usage of standardized formats (such as STIX, SCAP and ontology work by OCA and OpenDXL) between the different components integrated in the SOCCRATES orchestrator is one of the research topics that could be included in the roadmap to be explored. Currently, each SOCCRATES component has its own data format for requests and responses, which are used internally by the orchestrator to share information throughout the different components involved in the tasks included in the workflows defined. This normalization in the data formats could help to ease future extensions of the workflows and pave the way for the integration of the platform with other security products or tools used in SOC and CSIRT environments.

OVERALL PERSPECTIVE AND RECOMMENDATIONS

Overall perspective

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. Although SOCCRATES has had its main focus on the technical (IT) side of this transformation, we also gathered valuable technical and non-technical insights from the preparation and deployment of the SOCCRATES main pilot and the design and implementation of the SOCCRATES platform itself.

Role of the human analyst

The transformation will trigger a major change in the man-machine interaction, as we have learned from the pilot preparation. SOC and CSIRTS need to transform from the human centric organisations that they currently are (with scattered use of automated tools where available) to highly automated environments that seamlessly support the SOC/CSIRT workflows with integrated and balanced automated capabilities and functions, such as provided by the SOCCRATES platform. The human will remain an important factor in SOC and CSIRTs, but the activities of the human analyst in this highly automated environment will change considerably. Instead of creating its own situational awareness by collecting data from very many sources and correlating this data manually, the analyst will be supported by excellent situational awareness and extensive option awareness created by the automation platform. This will require a different skill set from the human analysts, they will have to act and take decisions on a different aggregation level and might need less knowledge and understanding of technical infrastructural details. To trigger the optimal mitigation actions, analysts will be able to choose from a set of elaborate Course of Action options that are presented to them, including all potential consequences on both

technical and business level for each of those options. We think the job profile of the analyst will change quite drastically and this change will trigger the need for new educational programs. We think that more research, involving Human Factor experts, is needed to study the interaction between human analysts and highly automated environments, to ensure that this combination provides the most effective results. This should cover two directions; it is not only the question how to optimally prepare the human analyst for his task, but also a question on how to design the automation in such a way that it optimally fits and serves human capabilities. A human factor related aspect to take into account also is that the human analyst needs to be equipped to take decisions under pressure and in crisis situations. This is in principle not different compared to the current situation in SOC and CSIRTs, but the other aggregation level of presented information and necessary decisions could be of influence on this human behaviour under pressure. An there of course is the important continuity question; how can we ensure that the SOC and CSIRT analysist can still perform his or her duties when automated security tools cannot be used due to e.g. power outage, bugs or cyber-attacks on the tooling itself?

In conclusion, we may not forget the human analyst operating in the automated SOC and CSIRT environment, and make sure that training material is developed and provided to the SOC & CSIRT analysts to enable them to work efficiently and effectively within such highly automated SOC and CSIRT environment.

Transformation in a non-stop environment

While transforming, the show must go on; attackers will not pause their attacks to enable transformation of SOCs and CSIRTs, so the defensive shields must stay up. This introduces huge challenges in implementing the automation functions into the SOC/CSIRT environment and ensuring that the analysts are up to their tasks and at the same time avoid interruptions in the SOC/CSIRT activities. On top of that, automation functions will have to be

integrated into already functioning SOC/CSIRT environments. It is impossible to start from scratch (green field), because a lot of tools are already in use that cannot be completely phased out all at once. So there needs to be a migration period in which new automation functions have to interact with existing (legacy) tooling; this will cause probably in the beginning sub-optimal profits of the newly added automation functions.

Proactivity

From the five SOCCRATES use cases, there are actually is only one use case on incident response and four use cases that focus on avoiding incidents: changes in infrastructure (use case 4 and 5), discovery of vulnerabilities (use case 3) and reception of Cyber Threat Intelligence (use case 2). So the SOCCRATES vision is that more attention should be dedicated in preventing security incidents to occur by investing more effort in removing the entry points that are available to attackers to enter into the network. The logical reason for this is that incident response takes a lot of effort and time (and consequently, money) and also has unpredictable impact on business processes. Security incidents cannot be completely prevented, so mature capabilities to detect and respond to security incidents must be available, but it is always better to prevent an incident to happen. This of course starts with things such as good security by design and well-defined and implemented security controls. But the idea is that this is insufficient considering the continuously changing threat landscape, infrastructure, business processes and attacker's capabilities. So we need to take a pro-active approach and continuously monitor these changes and solve any issue that these changes might trigger as fast as possible before an attacker has the chance to exploit new attack opportunities. We think that his approach can lead to more stable and robust infrastructure and business processes and less time and effort spent on incident response.

Information overload

There are very many relevant information sources that feed the SOC and CSIRT, both internal and external. Those information sources have many different formats and contain information at business, process and technical level. It is very important to abstract exactly that information that is relevant to a SOC and CSIRT and also it is very important that the information that is needed

is available at the time it is actually needed. Taking into account the huge continuous information feeds, the support of AI and Machine Learning techniques to select, compress and interrelate the information is key.

Interaction between SOC/CSIRT and operations departments

In current SOC/CSIRT operations the SOC and CSIRT analysts need to interact with operational staff that manages the infrastructure to implement necessary changes in the infrastructure, to avoid incidents or to mitigate incidents. The operational staff is responsible for implementing the changes and bears responsibility to ensure that the infrastructure maintains its functionality. The interaction between SOC and CSIRT analysts and staff of operational departments currently mainly goes through the use of email, phone calls and ticketing systems. But we envisage more automated executed Curse of Action, directly triggered from the automation (SOCCRATES) platform and interacting with the infrastructure itself, bypassing the human SOC/CSIRT analyst and the operational staff. This automation is made possible because of the introduction of infrastructure that can be managed in an automated way, with the use of e.g. Network Function Virtualization, Software Defined Networking and Cloud platforms. Of course, this automation also comes with some challenges:

- Automation will decrease the time that is needed to implement changes in the infrastructure, but also limits the control that there is on those changes. The automated functions could have design flaws or bad implementation choices that on its turn could trigger unforeseen behaviour of the infrastructure
- It could be unclear who exactly bears the responsibility for the implemented changes when they are triggered by automated functions. Is it the SOC/CSIRT analyst, the implementor of the automation platform? Just like with the introduction of self-driving cars, thought needs to put into these issues to come up with solutions that satisfy all those involved and that can also be evaluated and traced back in a later stage.

Move to devops

In general, the IT world is moving towards more agile development methodologies, in which new software and infrastructure is developed in much shorter cycles

than in traditional (waterfall) development. Also Devops methodologies become more popular, in which an integrated team is responsible for both development and operations of software or infrastructure. Supported by Continuous Integration and Continuous Development (CI/CD) pipelines, new releases can in principle be put into production each few minutes. This of course also presents challenges to the situational awareness of the SOC/CSIRT, which could change any minute. It could also be of influence on the option awareness, which suggests (automated) implemented changes for mitigation of incidents or proactive response. The timing of the response needs to be aligned in one way or another with the short cycles of the devops environment.

Interaction between SOC/CSIRT and business departments

Traditionally, it has proven to be difficult to explain to business people what the potential impact on their business could be of vulnerabilities that are found in the infrastructure or of security incidents that are detected. We think that this is one of the major improvements of the SOCCRATES platform, that provides the information to explain the effect of vulnerabilities or attack on an infrastructural asset in business terminology (e.g. expected impact, Return of Response). This will really bridge the gap between SOC/CSIRT analyst and business responsible and will improve decision making with better and understandable arguments. However, a prerequisite to do this effectively is that the business processes are structurally recorded, including some properties that are needed to assess impact of vulnerabilities and incidents, such as criticality of a business process. This is an area that needs to be improved if we want to profit optimally of automation as provided by the SOCCRATES platform. Also here, we need to solve the responsibility issue. If responses are highly automated, how much influence does the business responsible still have in deciding to e.g. switch off business processes? Will this be limited to defining requirements of the automated functions or does the business responsible still need to be consulted before changes are implemented that will have major impact on business processes? This is also an issue for further research.

Interaction between different SOC and CSIRTs

As explained on page 11, SOCs and CSIRTs come in various formats and on various levels. Currently, there already is exchange of threat information on different levels: technical (e.g. IoC's), but also on a more tactical level (e.g. Tactics, Techniques and procedures). We expect the level of information exchange to grow rapidly. This will happen between SOC/CISRTS of the same level, but also on SOC/CSIRT and CSIRT on different levels. The growth will be triggered by the fact that collaboration against cyber threats is more effective than individual defence, but also triggered by regulatory bodies. The latter want to have insight in the threats that are relevant to e.g. their national critical infrastructures but also want to ensure that their critical infrastructures have sufficient information and means to protect themselves. With the increased automation of SOC/CSIRTs environment, it could also be beneficial to connect those automated environments in real-time, e.g. connecting the SOCs of all the different entities in a supply chain or connecting all of the SOCs in a particular sector. This could have huge benefits, such as expanding the situational and option awareness. But on the other hand, this also comes with challenges, such as privacy, responsibility and even commercial related issues. We think this also is a topic for further research.

IT SOC/OT SOC Convergence

The focus of SOCCRATES is on SOC/CSIRT in IT infrastructure. But a huge amount of infrastructure is built up from Operation Technology (OT) components, such as electricity networks, drinking water supply, chemical industry or generic industrial environments. OT components are intelligent devices for sensing, monitoring, control, automation and communications. They typically have a combination of a physical component (e.g. a valve, a switch) and an IT component and typically have a very long life time. These different devices communicate with each other using a mixture of proprietary and standard protocols which in many cases lack support for security mechanisms such as proper encryption. We see an increase of 'OT SOCs' that monitor security in such environments. Those are mostly specifically built for their OT task and separated from IT SOCs. But the OT world is moving more to the IT world, using more open and standardized protocols (in stead of proprietary closed

protocols) and adding more IT functionality to OT components. We therefore expect that in OT environments, the OT SOC and IT SOC will start to merge into one 'converged IT/OT SOC' which of course will also be highly automated. But we expect that the specific requirements of OT environments will also induce specific requirements to the automation environment. This typically also is an area for further research.

Standardization

To enable automation and also to enable collaboration between automated environments, it is key to have a high degree of (industry) standardisation. The need for standardization not only is on technical level, but also on process level, and business level. Standardization facilitates collaboration because different SOC/CSIRTs can easily connect, makes integration between different tools easier, but also provides direction to new product development and makes education much more effective. The cyber security industry, and ICT industry in general, should intensify standardisation activities required to enable automation in security operations.

EU sovereignty

We consider SOC and CSIRT environments to be a key element in building up resilience against cyber-attacks. But at the same time this specifically makes them vulnerable for attacks, especially supply chain attacks. Considering the geopolitical situation in the world, it would be very prudent to ensure that development of the essential tooling that is needed by EU companies for automated SOC and CSIRT environments is designed and manufactured in the EU. Although this of course does not provide 100% guarantee that (supply chain) attacks will not impact the automation tooling for SOC and CSIRT, it at least provides some level of confidence and control on the products that are applied.

Recommendations

We think that SOCCRATES has made a major contribution to the automation of SOC and CISRT environments. But a lot remains to be done. This is why we recommend to further develop and evaluate the SOCCRATES platform, to ensure that even more of the innovative con-

cepts and ideas will end up in commercial products. Of course, the SOCCRATES partners will all continue with the SOCCRATES developments in their own context, but we think that more collaborative research efforts are needed to increase the maturity level of the developed technology.

We recommend to invest in the research topics that are described in the research roadmaps in previous chapters. We strongly believe that this is the direction that we need to take to increase the maturity level of automation in SOC and CSIRT environments.

As mentioned in this paper, SOCCRATES focuses on the technical automation platform for SOC and CSIRTs of large organizations and MSSPs. But in further research we should take into account the broader picture (e.g. not only technical focus, see also the 'overall perspective in the previous section).

REFERENCES AND ENDNOTES

References

C. Islam, M. B. (2019). A multi-vocal review of security orchestration. ACM Computing Surveys (CSUR), 1-45.

Claudio Neiva, C. L. (2020). Market Guide for Security Orchestration, Automation and Response Solutions. Gartner.

Cybersecurity Insider. (2021). 2021 SOAR Report - Swimlane.

Edgars Taurins. (2020). HOW TO SETUP UP CSIRT AND SOC. European Union Agency for Cybersecurity (ENISA). doi:10.2824/056764

ENISA. (2018). Information Sharing and Analysis Centres (ISACs) - Cooperative models. Athens: European Union Agency For Network and Information Security.

FIRST. (2019). Computer Security Incident Response Team (CSIRT) Services Framework, version 2.1. Forum of Incident Response and Security Teams, Inc. (FIRST.Org). Retrieved from https://www.first.org/standards/frameworks/csirts/csirt_services_framework_v2.1

Johnson Kinyua, L. A. (2020). AI/ML in Security Orchestration, Automation and Response: Future Research Directions. Intell. Autom. Soft Comput.

Martin Husák, M. L. (2021). System for Continuous Collection of Contextual Information for Network Security Management and Incident Handling. The 16th International Conference on Availability, Reliability and Security, (pp. 1-8).

N. Brownlee, a. E. (1998, June). RFC 2350 - Expectations for Computer Security Incident Response. IETF. Retrieved from https://datatracker.ietf.org/doc/html/rfc2350

Pontus Johnson, R. L. (2018). A Meta Language for Threat Modeling and Attack Simulations. Proceedings of the 13th International Conference on Availability, Reliability and Security (ARES 2018) (pp. 1 - 8). New York: Association for Computing Machinery.

S. Jajodia, S. N. (2011). Cauldron Mission-centric Cyber Situational Awareness with Defense in Depth. 2011-MILCOM 2011 Military Communications Conference (pp. 1339-1344). IEEE.

S. Noel, E. H. (2016). CyGraph: graph-based analytics and visualization for cybersecurity. Handbook of Statistics, 35, 117-167.

Sotirios Katsikeas, P. J. (2020). An Attack Simulation Language for the IT Domain. International Workshop on Graphical Models for Security (pp. 67-86). Springer.

Tesink, S. (2020). D2.2 System Architecture & Interface Specifications (Initial). European Commission.

Trapero, R. (2020). D6.1 Initial version of the SOCCRATES Platform Orchestration, Reconfiguration and Front-end. European Commission.

Yan Shoshitaishvili, M. W. (2017). Rise of the hacrs: Augmenting autonomous cyber reasoning systems with human assistance. Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, (pp. 347-362).

Zimmerman, C. (2014). Ten Strategies of a World-Class Cybersecurity Operations Center. The MITRE Corporation.

Endnotes

1 Asset Management (ID.AM) is the first Category of cybersecurity activities in the NIST Cyber Security Framework (https://www.nist.gov/cyberframework).

2 CIS control 1 and 2 of the CIS Critical Security Controls for Effective Cyber Defense [1] are on Iinventory and control of enterprise assets, and of software assets.

3 https://www.microsoft.com/security/blog/2021/03/02/hafnium-targeting-exchange-servers/

4 https://www.cisa.gov/uscert/ncas/alerts/aa21-356a

5 For example, in Fortinet FortiGate products: https://ics-cert.kaspersky.com/publications/reports/2021/04/07/vulnerability-in-fortigate-vpn-servers-is-exploited-in-cring-ransomware-attacks/

6 https://www.eset.com/int/about/newsroom/press-releases/products/brute-force-attacks-targeting-remote-access-increased-during-the-covid-19-pandemic-eset-confirms/

7 https://www.shadowserver.org/what-we-do/network-reporting/

8 https://shodan.io

[9] https://censys.io
[10] https://www.bitsight.com/
[11] https://securityscorecard.com/
[12] https://www.ntia.gov/sbom
[13] https://www.gartner.com/en/documents/2487216/definition-threat-intelligence
[14] https://nsarchive.gwu.edu/document/17212-united-kingdom-government-threat-intelligence
[15] https://attack.mitre.org/
[16] http://detect-respond.blogspot.com/2013/03/the-pyramid-of-pain.html
[17] https://www.ecb.europa.eu/paym/cyber-resilience/tiber-eu/html/index.en.html
[18] https://github.com/center-for-threat-informed-defense/adversary_emulation_library
[19] https://oasis-open.github.io/cti-documentation/
[20] https://www.misp-project.org/
[21] https://attack.mitre.org/docs/ATTACK_Design_and_Philosophy_March_2020.pdf
[22] https://solutionsreview.com/security-information-event-management/the-8-best-soar-security-companies/, https://www.trustradius.com/security-orchestration-automation-and-response-soar, https://geekflare.com/es/best-soar-tools/, and https://www.gartner.com/reviews/market/security-orchestration-automation-and-response-solutions
[23] https://shuffler.io/
[24] https://www.oasis-open.org/committees/cacao
[25] https://www.oasis-open.org/committees/openc2

APPENDIX A

SOCCRATES PROJECT

SOCCRATES (SOC & CSIRT Response to Attacks & Threats based on attack defence graphs Evaluation Systems) is an EU funded project under the Horizon2020 programme that has the following main challenge:

How can SOC and CSIRT operations effectively improve their capability in detecting and managing response to complex cyber-attacks and emerging threats, in complex and continuously evolving ICT infrastructures while there is a shortage of qualified cybersecurity talent?

The main objective of SOCCRATES is to develop and implement a security automation and decision support platform ('the SOCCRATES platform') that will significantly improve an organisation's capability (usually implemented by a SOC and/or CSIRT) to quickly and effectively detect and respond to new cyber threats and ongoing attacks.

The SOCCRATES platform (see Figure 4) consists of an orchestrating function and a set of innovative components for automated infrastructure modelling, attack detection, cyber threat intelligence utilization, threat trend prediction, and automated analysis using attack defence graphs and business impact modelling to aid human analysis and decision making on response actions, and enable the execution of defensive actions at machine-speed.

Figure 4 **The SOCCRATES platform**

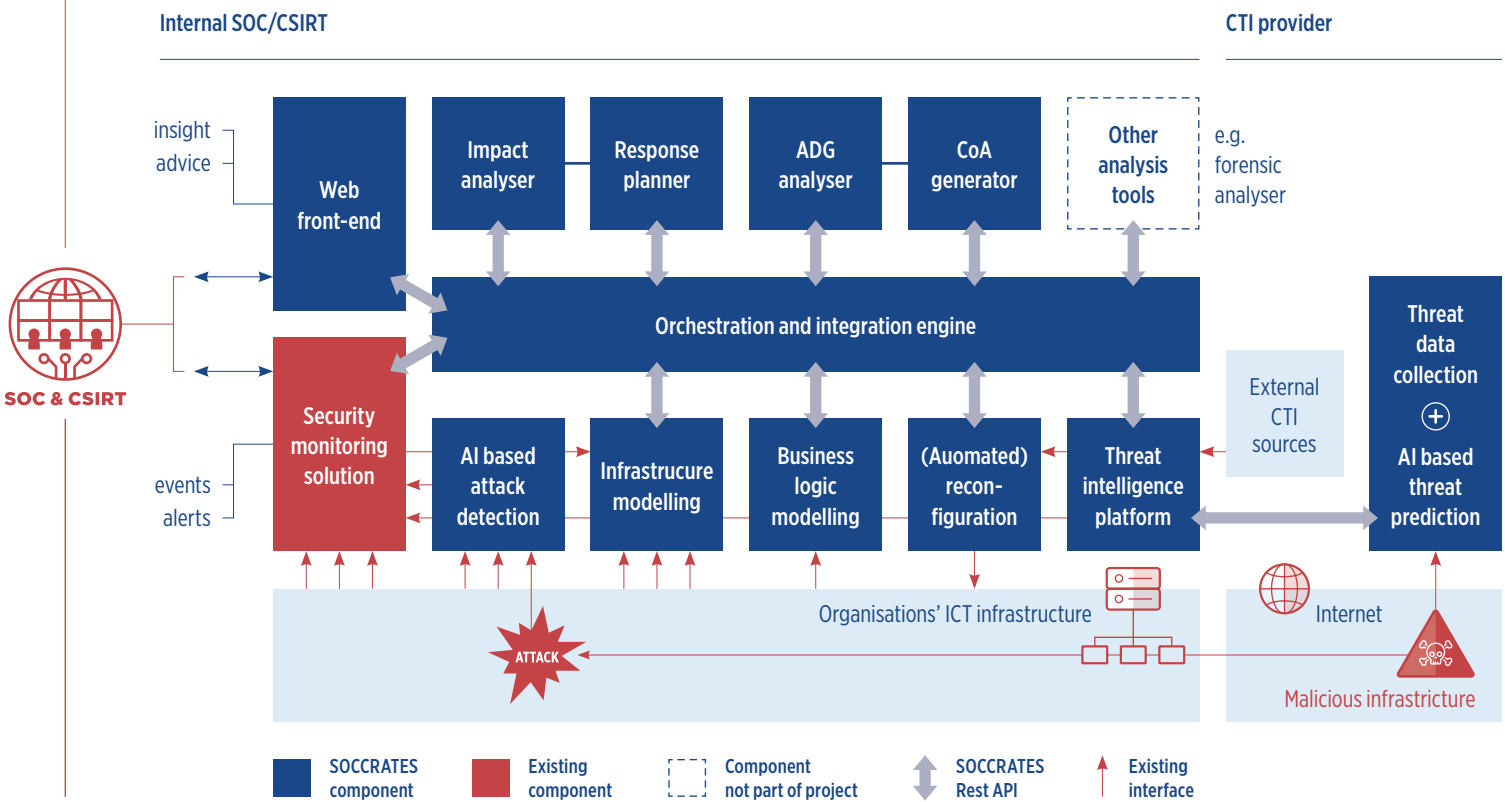

SOCCRATES has the following concrete project objectives:

- Deliver the SOCCRATES platform consisting of an orchestration function and a unique integration of innovative background solutions that seamlessly work together.
- Show that the SOCCRATES platform can improve SOC operations by evaluating the SOCCRATES platform in two diverse real-life pilot environments.
- Examine and illustrate the benefits of automation for selected SOC activities to help manage the cyber security skills gap in organizations.
- Prepare for successful exploitation by the SOCCRATES partners of the individual innovated components and the integrated SOCCRATES platform in commercial products that are offered to the market and are available for the European (business) community.

Please visit **www.soccrates.eu** for more information on the SOCCRATES project.

ABBREVIATIONS

This glossary serves as inventory of abbreviations used in the document.

Acronym	Description
ACT	semi-Automated Cyber Threat intelligence
ADG	Attack Defence Graph
AI	Artificial Intelligence
API	Application Programming Interface
APT	Advance Persistent Threat
AV	AntiVirus
BPMN	Business Process Model and Notation
CC	Command and Control
CERT	Computer Emergency Response Team
CMDB	Configuration Management Database
CSIRT	Computer Security Incident Response Team
CoA	Course of Action
CTI	Cyber Threat Intelligence
DGA	Domain Generated Algorithm
MTTD	Mean Time To Detection
DNS	Domain Name System
ICT	Information and Communication Technology
IDS	Intrusion Detection System
IMC	Infrastructure Modelling Component
IoC	Indicators of Compromise
IP	Internet Protocol
IPS	Intrusion Prevention System
IRM	Incident Response and Management
ITIL	Information Technology Infrastructure Library
LAN	Local Area Network
Mn	Infrastructure Model (at time n)
MSSP	Managed Security Service Provider
OT	Operational Technology
OS	Operating System
RORI	Return on Response Investment
SDN	Software Defined Network
SIEM	Security information and event management
SOAR	Security Orchestration, Automation and Response
SOC	Security Operation Centre
SOCCRATES	SOC & CSIRT Response to Attacks & Threats based on attack defence graph Evaluation Systems
SSL	Secure Sockets Layer
TIP	Threat Intelligence platform
TTC	Time To Compromise
UC	Use Case
VLAN	Virtual LAN

CONTRIBUTING EXPERTS

SOCCRATES partner	Experts
TNO	Marije Berding, Gabriela Bodea, Irina Chiscop, Federico Falconieri, Frank Fransen, Richard Kerkdijk, Sebastiaan Tesink, Reinder Wolthuis, Lenny Zilverberg
Mnemonic AS	Fredrik Borg, Siri Bromander, Martin Eian, Dan Rosenzweig Eig, Kristian Haga, Phanida Meesin, Mona Elisabeth Østvang, Ola Flølo Ringdalen, Fredrik Morten Sætran, Philip Christian Scheel, Geir Skjøtskift, Tom Spangebu
AIT Austrian Institute of technology GMBH	Agron Bajraktari, Ewa Piatkowska, Paul Smith, Francesca Soro
Kungliga Tekniska högskolan	Ashish Dwivedi, Mathias Ekstedt, Preetam Mukherjee, Giuseppe Nebbione, Wojciech Wideł
Foreseeti AB	Per Eliasson, Joakim Nydren, Erik Ringdahl
Institut Mines-Télécom - Télécom SudParis	Gregory Blanc, Hervé Debar, Christophe Kiennert, Marwan Lazrag, Paul Henri Mignot
Stichting the Shadowserver foundation Europe	Piotr Kijewski, David Watson
WithSecure	Ville Alkkiomäki, Marko Komssi, Seamus Moloney, Andrew Patel
Atos Spain S.A.	Maria Angeles Bermudez Garcia, Rodrigo Diaz Rodríguez, Alejandro García Bedoya, Susana González Zarzosa, Alberto Miranda García, Ruben Trapero Burgos, Jesús Villalobos Nieto
Vattenfall IT services Poland Sp. z.o.o.	Pawel Babski, Michal Celinski-Myslaw, Małgorzata Greń, Rafal Kondracki, Maciej Kosz, Piotr Szombierski

COLOPHON

© 2022 Partners of the SOCCRATES project

Editors Frank Fransen and Reinder Wolthuis
Design Studio Oostrum in cooperation with C10 Ontwerp
Print Impressed print & sign

Copyright
All rights reserved. No part of this document may be reproduced and/or published in any form by print, photoprint, microfilm, website or any other means without previous written permission.

Contact
reinder.wolthuis@tno.nl
www.soccrates.eu

Acknowledgement
This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement no. 833481

To receive a printed copy of the SOCCRATES vision paper, please send an email including your email address and postal address to: **info@soccrates.eu**.

SOCCRATES

SOC & CSIRT Response to Attacks & Threats based on attack defence graphs Evaluation Systems

www.soccrates.eu