White paper

Innovative directions for automation in cyber security operations

Authors
Frank Fransen, Bart Gijsen, Richard Kerkdijk,
Reinder Wolthuis, Ruggero Montalto

April 2023

TNO innovation for life

Acknowledgement

This whitepaper was written as part of the Program Cybersecurity Noord-Nederland.

The Program Cybersecurity Noord-Nederland has received funding from the RSP of the Province of Groningen and the municipality of Groningen.

Abstract

Despite heavy investments in their cyber defenses, most organizations are unable to keep pace with the ongoing evolution of threats and attack methods. Present day practices and solutions simply do not suffice to deal with the persistence and sophistication of professional threat actors. TNO believes that the gap between defenders and attackers can only be bridged through a fundamental game changer and that automation might hold the key towards evening the odds. This whitepaper describes the vision about automation in cyber security operations that TNO has developed to this end, as well as key innovations in the fields of automated security reasoning and automated response that TNO is currently invested in.

Innovative directions for automation in cyber security operations

Contents

1 Introduction

As cyber-attacks became more sophisticated and their disruptive effects (both on business and society) increased, organizations with a strong dependency on ICT (particularly ICT applied in vital processes) have gradually elevated their cyber defenses. Strategies typically included an increased focus on security monitoring and incident response capabilities, often through the establishment of dedicated Security Operations Centers (SOCs) and Computer Security Incident Response Teams (CSIRTs).

To further strengthen their resilience to cyber-attacks, many organizations have subsequently complemented this with Cyber Threat Intelligence (CTI) and threat hunting practices. While this evolution has arguably increased defensive capabilities, threat actors have also been stepping up their game and have consistently managed to come out ahead.

This observation was already well expressed in ENISA's Threat Landscape Report of 2017 [ENISA]:

"The cyber security community is still far from striking the balance between defenders and attackers."

More recently, a global survey into the state of cyber security resilience [Accenture] revealed that a vast majority of security leaders agrees with the following statement:

"Staying ahead of attackers is a constant battle and the cost is unsustainable."

What these statements illustrate is that there is still a rather substantial imbalance between attackers and defenders. A principal cause lies in the given that defensive practices tend to rely heavily on human effort and expertise, while advanced attacks are often automated to great extent. This has rather visible effects in operational practice, where attacks can be launched in a matter of seconds, while discovery and containment might take days, weeks or even months. An obvious driver for pursuing automation in cyber security operations is therefore to accelerate the *speed of response*, not only to reduce the direct damage that incidents might induce, but also to discourage attackers from further attempts and strengthen confidence in ICT infrastructures. Here, we note that the need to accelerate response is not limited to (cyber) attacks, but also extends to the remediation of newly discovered vulnerabilities in equipment and software.

Apart from speed, the desire to automate cyber security operations is also driven by resource related factors. Automation solutions can for instance *relieve security specialists* from (what could be) routine (repetitive) tasks and allow them to focus more emphatically on complex activities such as threat hunting or developing new use cases for security monitoring solutions. Since pressure on FTEs is not uncommon among security operations teams, freeing up resources might be an essential requisite for facilitating such advanced practices. A shift towards more complex duties will likely also have a motivating effect on security specialists and allow organizations to *retain talent* for a longer period of time, which is appealing considering the structural shortage of qualified cyber security staff [GAP], [Frost].

The need to automate cyber security operations is widely recognized and has already led to an array of commercial and technological developments. On the product side, the advent of playbook[1] driven security automation through so-called Security Orchestration, Automation and Response (SOAR) solutions has been quite prominent. Meanwhile, advances have also been made in standardizing languages and formats that can facilitate security automation. An example of the latter is the OpenC2 language maintained by OASIS [OpenC2], which enables machine-to-machine control of an organization's tools and applications for cyber security. To bridge (or at least minimize) the gap between attackers and defenders, however, TNO believes that the concept of automating (cyber) security operations needs to be taken significantly further. This whitepaper presents TNO's vision on next generation automated security and some of the key innovations it pursues in this field.

1 The term playbook refers to a computer program that specifies a workflow of relatively simple actions. Security oriented playbooks are executed by SOAR tools.

2 Pursuing automation in the SOC and CSIRT space

While a great variety of cyber defenses might be automated to more or less extent, particular potential seems to lie in the automation of SOC and CSIRT operations. In part, this is likely due to the fact that human driven SOC and CSIRT practices have already evolved at great pace in recent years. Figure 1 depicts the typical context of SOC and CSIRT teams and positions the automation solutions that TNO and its partners are currently pursuing for such environments.

Note that Figure 1 reflects the setup of a (large and mature) organization that maintains in-house SOC and CSIRT provisions to protect a self-managed technical infrastructure. The automated security solutions and innovations presented in this paper, however, are equally applicable to situations in which operational security duties have been outsourced to a Managed Security Service Provider (MSSP) or where (some of) the technical infrastructure is maintained by third parties.

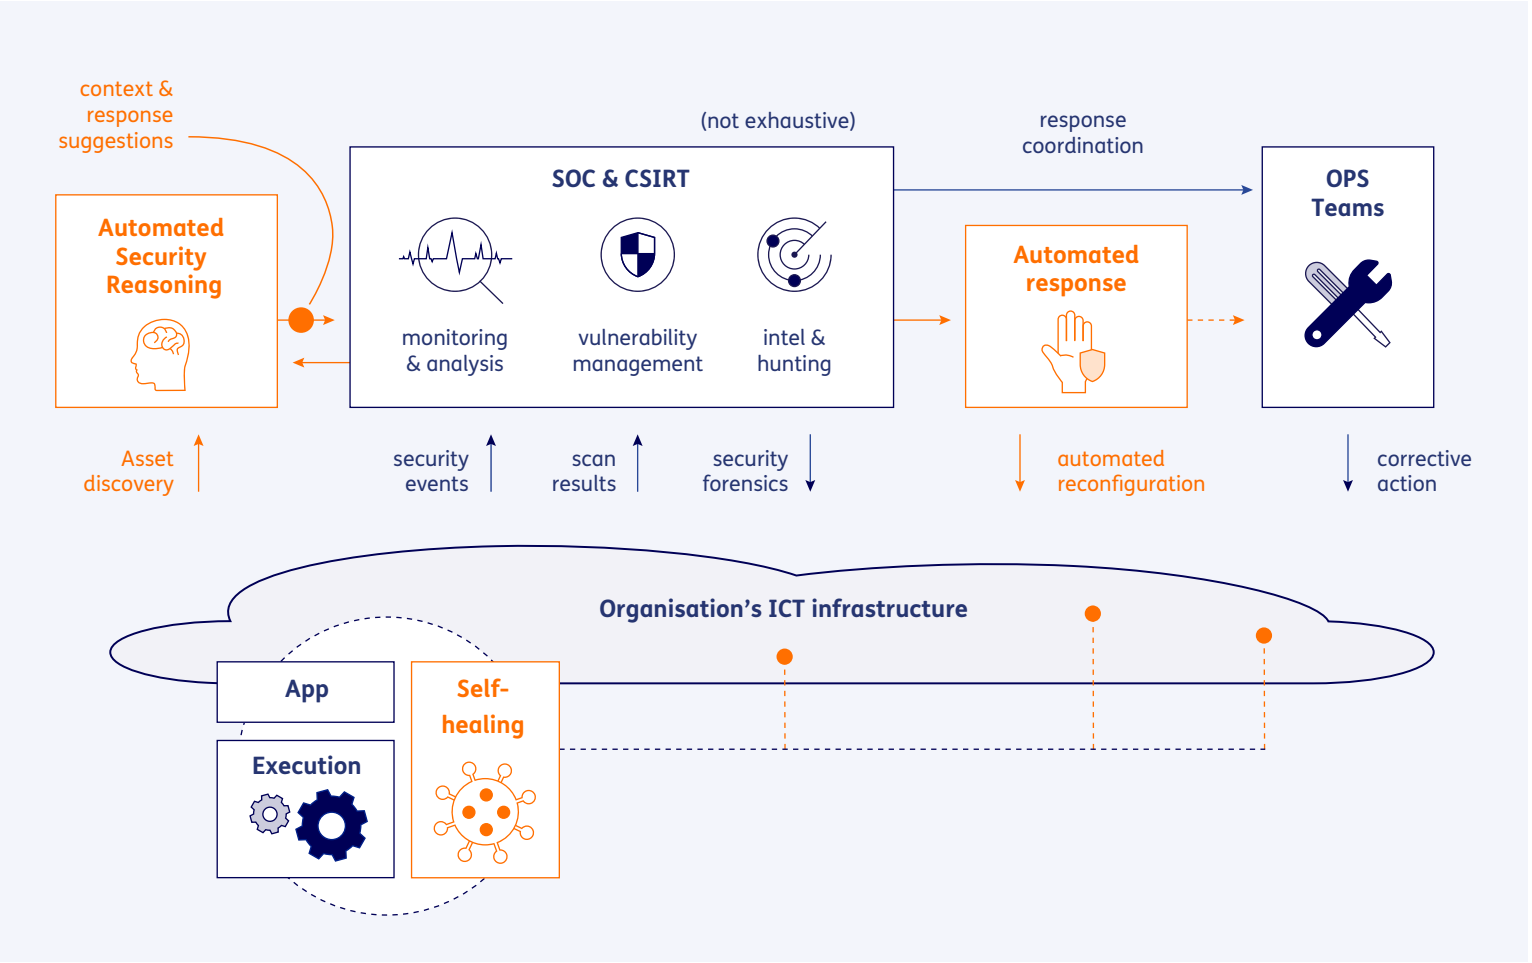

Figure 1: Vision on automation in the SOC and CSIRT space.

In present day ICT infrastructures (depicted in blue in Figure 1), security analysts residing in the SOC assess security relevant events and follow up on any observation that might represent an actual security incident. Such events are typically put forward by an array of security monitoring solutions[2]. Upon establishing the characteristics of the incident at hand, the analyst decides on an appropriate response strategy and subsequently initiates an actual mitigation in the ICT infrastructure. The task of executing corrective actions typically resides with operations teams that maintain the organization's systems and applications.

To request such actions, SOC analysts often employ the ticketing system that their technology unit is accustomed to. In particular cases, an incident might need to be escalated to the organization's CSIRT for further handling. This typically takes place if the nature or severity of an incident exceeds the SOC's mandate. CSIRT teams are usually comprised of more experienced cyber security specialists and better equipped to conduct thorough (forensic) investigations on the affected assets.

The need for automation is generally recognized among (mature) SOC and CSIRT teams, most prominently in the field of incident response where the aforementioned SOAR solutions are increasingly gaining traction.

Automated security innovation

Moving forward, TNO believes that significant advancements can be achieved by pursuing the following directions:

- **Automated security reasoning.** Under this header, TNO is particularly invested in the concept of *real-time dynamic risk assessment*. When employed to this end, security reasoning technology can provide SOC and CSIRT specialists with a contextual understanding of threats and events and allow them to take (more) informed decisions on subsequent response strategies. The basis for such reasoning lies in elaborate models of an organization's ICT infrastructure and the tactics and techniques of adversaries that may be targeting it.
- **Automated response.** Here, TNO pursues the concept of *automated playbook generation and execution*, to fully automate response strategies for newly emerging threats or attacks with no or only limited human intervention. In addition, TNO has been exploring *self-healing technology* that allows ICT assets to anticipate, withstand and recover from threats and attacks in a fully autonomous fashion. The latter concept was inspired by the human immune system and correspondingly operates directly on the technical assets themselves.

The remainder of this white paper will focus on these innovative directions and the results that have been achieved to date. Across all these initiatives, it is stressed that TNO does not envisage automation technology for cyber security operations to replace the human analyst altogether. Rather, automation solutions can help such analysts become more effective, either by relieving them from (what could be) repetitive tasks or by supplying them with (contextual or threat related) insights that human experts could likely not develop at the same pace.

2 Mature setups typically include a Security Information and Event Management (SIEM) solution, an Intrusion Detection System (IDS) and an Endpoint Detection and Response (EDR) solution.

3 Automated Security Reasoning

Devising an appropriate response to (newly discovered) threats and attacks involves a great degree of analytics. Whenever new threat insights are gathered, for instance, organizations need to determine if and how their particular technical infrastructure might be affected.

Similarly, when monitoring systems (e.g. SIEM, IDS or EDR solutions) raise a new security event, organizations will want to understand the larger sequence of attack steps to which their infrastructure is being subjected and the (critical) assets that might be at risk.

Establishing such contextual understanding and extrapolating this into an effective course of action not only requires knowledge of (cyber) threats and attack techniques, but also of the organization's business and technology landscape.

Human analysts typically need some time to piece everything together and might need to consult other specialists in their organization (e.g. engineers and operators of the ICT infrastructure under consideration) to complete their expert appraisal. The workload (and thus the lead time) involved in such analysis is often substantial and will likely increase further as ICT infrastructures become more complex and diverse (e.g. due to the rapid deployment of new technologies such as cloud and IoT).

Security analysts can already avail of various tools and solutions to assist them in comprehending the nature and impact of particular (security related) events. Contemporary Threat Intelligence Platforms (TIPs), for instance, ingest and correlate threat information from a variety of sources to gather a composite view on specific threats or threat actors, while asset registration tools such as an organization's CMDB (Configuration Management Database) capture the status and ownership of IT assets and the business processes that rely on them.

The process of reasoning towards a diagnosis and determining an appropriate way forward, however, is still largely a human expert task. TNO believes that this process can be supported more directly by leveraging technologies for modelling infrastructure as well as cyber adversary behavior. The envisaged setup is explained further in the following section.

3.1 Real-time dynamic risk assessment

At its core, the concept of real-time dynamic risk assessment aims to supply security analysts with two elementary forms of insight:

- **Situational awareness.** This refers to the automated analysis and extrapolation of security relevant events to provide security analysts with an appropriate situational understanding, e.g. concerning attack paths that may result from newly discovered threats or vulnerabilities or assets that are likely to be affected by an ongoing attack.
- **Option awareness.** This concerns the automated generation of potential Courses of Action (CoAs), each accompanied by an appraisal of expected (mitigating) effects and possible business trade-offs, that an analyst could consider in response to a threat or event.

These combined insights can contribute to (more) informed decisions on the most appropriate response to specific threats or incidents. Correspondingly, the concept of real-time dynamic risk assessment is often positioned as a form of *security decision support*. Figure 2 presents the conceptual technical environment through which the desired insights can be produced. In essence this environment encompasses the following functional elements:

- **Infrastructure modelling.** Foundational capabilities to compile machine readable expressions of ICT assets (component I in Figure 2) and the larger ICT infrastructure in which they reside (component II). To facilitate a meaningful security analysis, resulting models must be greatly detailed (to the level of operating systems, applications and system configurations) and include a wide variety of (cyber) security relevant properties (e.g. vulnerabilities, patch levels, access control policies and firewall configurations).
- **Adversary modelling.** Modelling of attacks (component III in Figure 2) that can be executed on individual technical assets (e.g. specific routers or server types) and ultimately of adversaries that employ such attacks in a particular combination or sequence (component IV). To ensure realistic simulations, it is desirable that such adversaries (and their associated attack techniques) can be positioned at various points in the modelled infrastructure (representing different 'footholds' from which the adversary might launch its initial activities).
- **Attack Defense Graph (ADG) analysis.** The ability to predict how attacks might propagate across the modelled ICT infrastructure and to which extent they are impeded by security controls that the organization has already deployed (component V in Figure 2). To express the likelihood that a particular attack will in fact be successful, the module could for instance calculate a so-called Time-To-Compromise (TTC).
- **Course of Action (CoA) assessment.** The ability to generate a set of viable Courses of Action (e.g. deployment or reconfiguration of specific security controls), compare their respective (threat mitigating) effects (e.g. in terms of TTC reduction) and determine relevant business trade-offs (which are implicitly included in component VI in Figure 2). The latter refers to the (potential) impact of the threat or attack itself as opposed to the impact that a Course of Action might have on particular business processes.

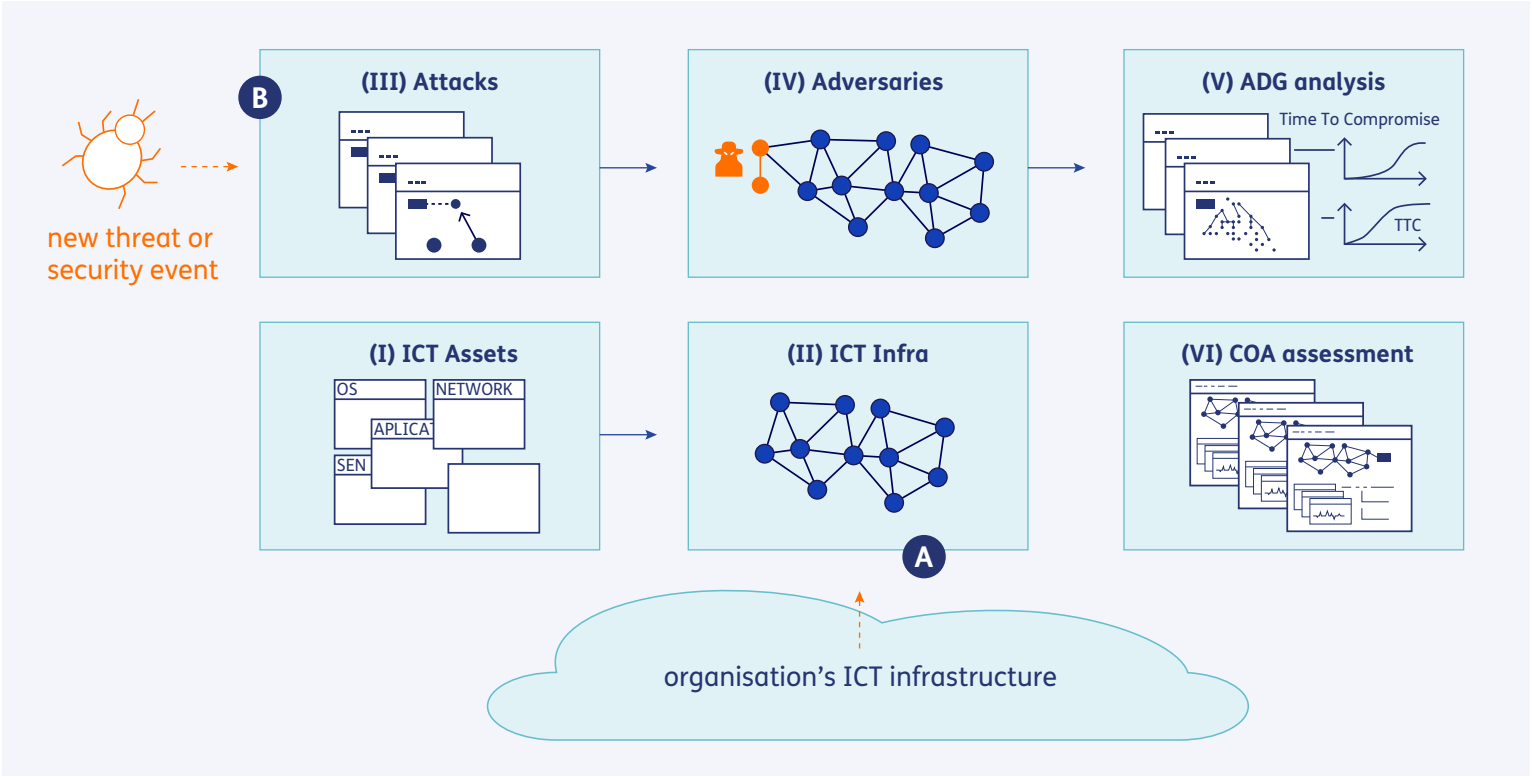

Figure 2: Core building blocks of automated security reasoning environment.

To ensure that the security reasoning environment supplies the analyst with realistic appraisals and recommendations, it must continuously be fed with up-to-date information on (newly discovered) threats and events and (changes in) the organization's ICT infrastructure.

In view of this, the setup depicted in Figure 2 also requires the following distinct capabilities:

A. **Infrastructure discovery.** The ability to identify assets, configurations and system structures within the organization's operational ICT infrastructure, e.g. by leveraging network scanners, asset discovery tools, vulnerability scanners and CMDBs, and processing the outcome into an infrastructure model that accurately depicts the present operational reality (a 'digital twin' or close approximation thereof). Considering the dynamics of typical ICT infrastructures, this process will need to take place on a continuous basis. Once a base representation is in place, however, the scanning and discovery effort can largely be focused on relevant changes (e.g. configuration adjustments or newly deployed technical assets).

B. **Threat and event ingestion.** The ability to update attack and adversary models based on newly received threat information (intelligence) and to match security events that were raised by the organization's security monitoring systems (SIEM, IDS, EDR, etc.) with attack steps for which models are already in place. The former allows ADG simulations to represent the most current threat insights, whereas the latter provides an understanding of the larger sequence of attack steps that an observed event might form part of.

The philosophy that underlies the concept of real-time dynamic risk assessment is that security analysts should be offered appropriate support and guidance whenever the security state of an operational ICT infrastructure needs to be reassessed. In practice, this can be triggered by a variety of events and situations. Newly discovered vulnerabilities or alerts raised by the organization's security monitoring systems are obvious examples, but changes in the ICT infrastructure itself (e.g. a newly discovered technical asset) might also open up new attack paths and thus require further appraisal. Developing and optimizing a security reasoning environment that facilitates this properly is the focus of several innovation projects that TNO is or has been involved in. A prominent example is the pan-European SOCCRATES project [SOCCRATES].

Automated security reasoning in the SOCCRATES platform

SOCCRATES (SOC & CSIRT Response to Attacks & Threats based on attack defense graphs Evaluation Systems) is an EU funded project under the Horizon 2020 program [SOCCRATES]. The project is coordinated by TNO and its partners include knowledge institutes, security solution providers and end user organizations. The core objective of the project is to develop and trial a next level automation platform that enhances the effectiveness of SOC and CSIRT operations.

This platform revolves around 5 use cases that represent situations in which an organization needs to reassess its security state and determine if and how it should respond in order to protect its interests:

- **UC1** Response on detection of ongoing attack.
- **UC2** Response on newly received Cyber Threat Intelligence.
- **UC3** Response on discovery of new vulnerable assets.
- **UC4** Response on discovery of system configuration change.
- **UC5** Response on deployment of new systems in infrastructure.

Each use case triggers a particular workflow in the SOCCRATES platform that is governed by the central orchestration and integration engine. Security reasoning technology plays an instrumental role in all cases and forms the basis for any Course of Action (CoA) that the platform initiates. The foundation for such reasoning is established by a dedicated *Infrastructure Modelling Component* (IMC) that collects network and asset data from a variety of sources (a.o. vulnerability scanners, DHCP servers, Netflow enabled devices and CMDBs) and aggregates this in a centralized graph database.

The resulting infrastructure model is fed to a separate *ADG Analyser* that performs probabilistic attack simulations, plots the most probable attack paths to high value assets (example shown in Figure 3) and calculates a corresponding Time-To-Compromise (TTC).

The ADG Analyser is based on Foreseeti's SecuriCAD solution [SecuriCAD] and employs the Meta Attack Language [MAL] to capture infrastructure assets and access (attack) techniques. SecuriCAD also has the ability to suggest mitigations. Correspondingly, it also serves as the platform's *Course of Action (CoA) Generator*. These Courses of Action are applied and evaluated in additional ADG simulations to determine their effect on the expected TTC. They are also fed to the platform's *Business Impact Analyzer* (BIA) component to determine their financial and operational impact[3]. TTC reduction values and business impact appraisals are both presented to the security analyst in order to facilitate decisions on actual CoA deployment.

3 For an explanation of the underlying business logic, we refer to the elaborate project documentation that is available via [SOCCRATES].

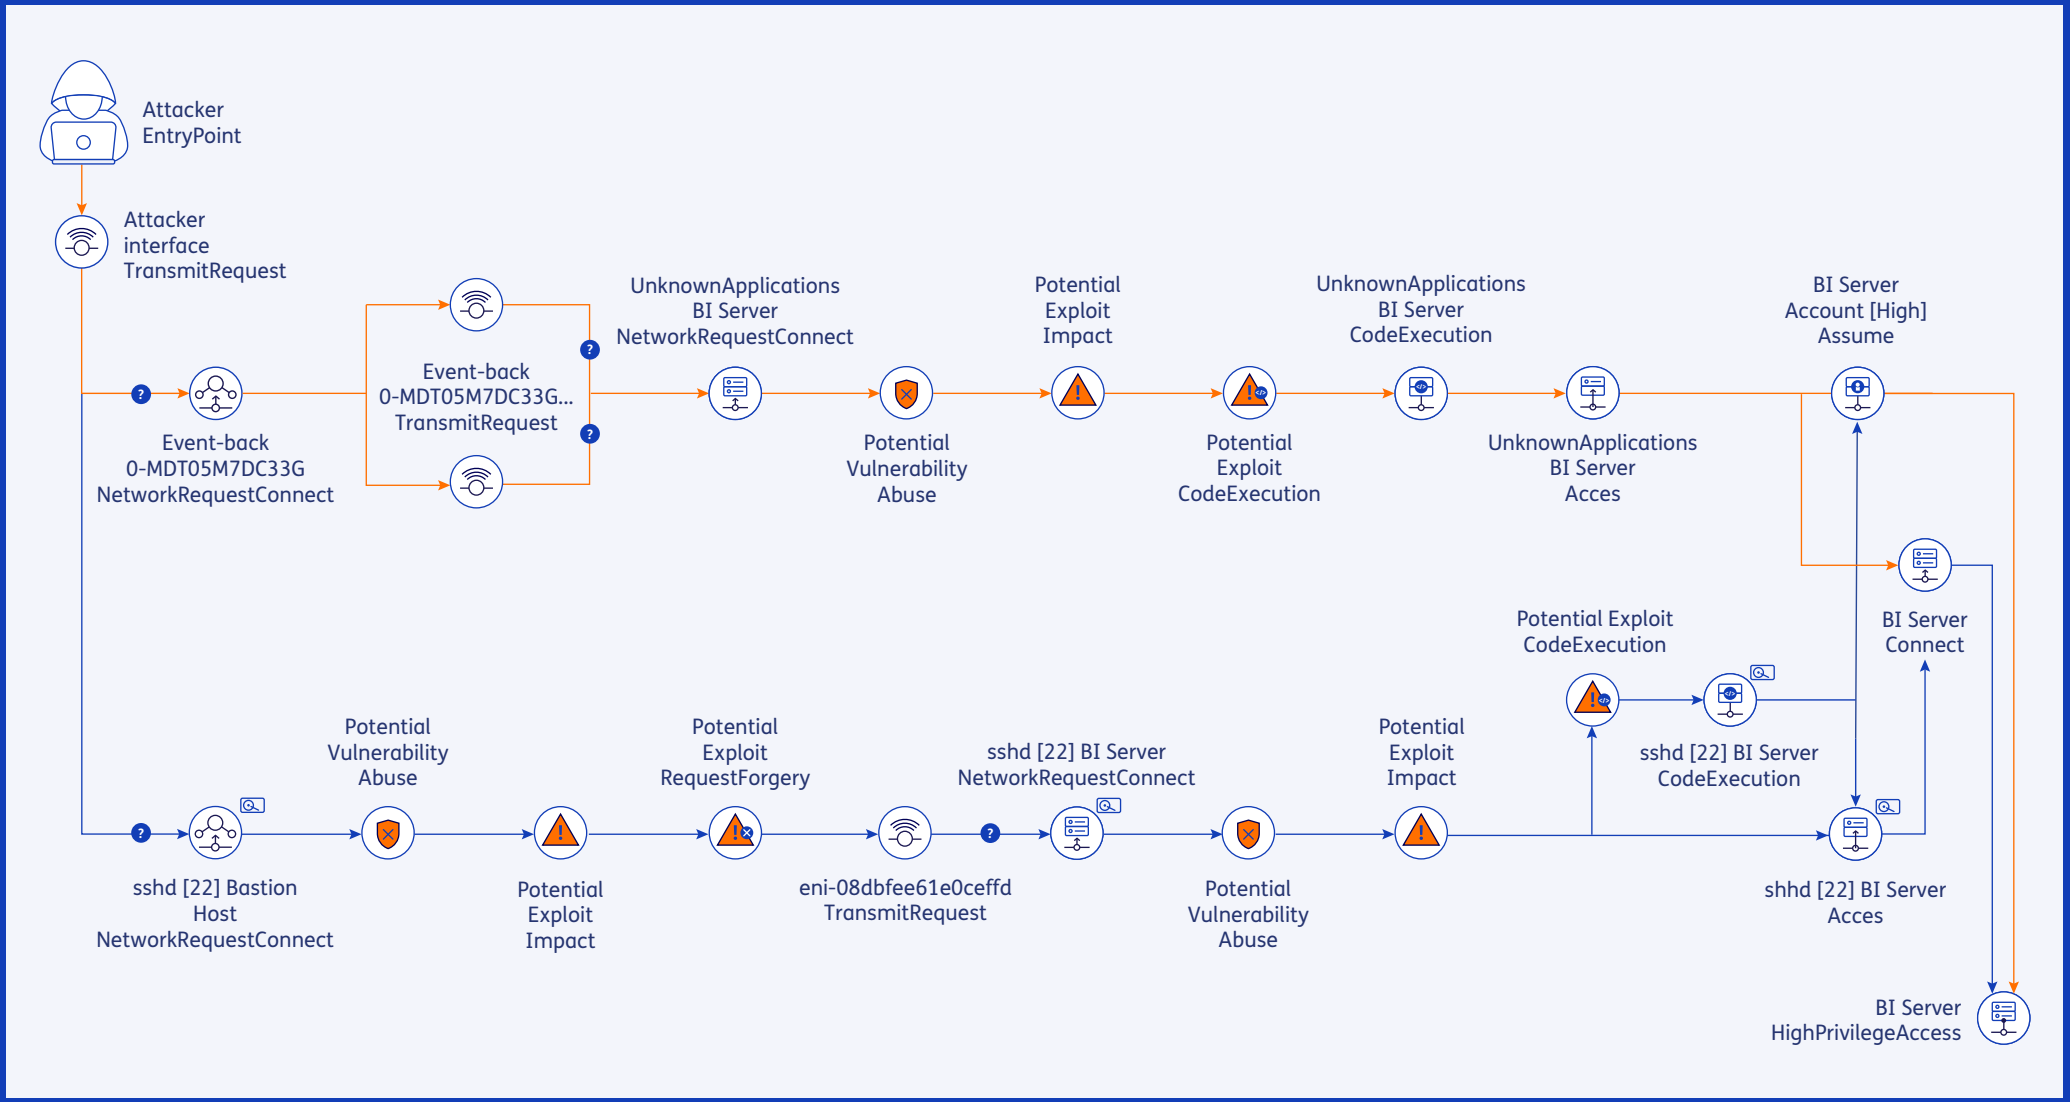

Figure 3: Attack graph that shows how simulated attack steps propagate through infrastructure.

4 Automated response

An organization's response to a (suspected) security incident traditionally relied heavily on human effort and expertise. While large and mature organizations usually maintained standardized processes for incident management, the tools to support them were typically limited to fairly generic workflow and ticketing solutions. In practice, security analysts addressing something as common as a (suspected) malware infection would find themselves manually querying asset databases, gathering or verifying intelligence and 'cutting and pasting' information from one system to the next in order to complete their appraisal of the incident at hand. Actual mitigation would often involve a ticket to the responsible technology units, e.g. requesting to blacklist certain URLs and domains or block particular IP addresses, and if the mitigation required a substantial change or reconfiguration in the organization's ICT infrastructure, this would have to be onboarded in regular maintenance and change procedures (e.g. based on ITIL or a similar framework).

On the whole, therefore, the process of resolving a security incident could easily become cumbersome and time-consuming.

To optimize their daily operations, many SOC and CSIRT teams established automation in specific process steps through self-built scripts and customized tools. As the need for automating security operations became more evident, however, manufacturers of security solutions introduced the concept of orchestrated security automation. This laid the foundation for so-called Security Orchestration, Automation and Response (SOAR) solutions, which Gartner defines as 'technologies that enable organizations to take inputs from a variety of sources (mostly from Security Information and Event Management – SIEM - systems) and apply workflows aligned to processes and procedures' [Gartner]. SOAR tools essentially facilitate the automation of predefined incident response workflows that are captured in so-called 'security playbooks'.

Their adoption is on the rise and under the supervision of OASIS, efforts are also ongoing to standardize playbooks that implement course of actions (CoA) for security operations [CACAO]. While CoA standardization is under development, we note that present SOAR installations mostly focus on automated information collection, thus relieving security analysts of many manual queries and lookups. Automated execution of actual responsive actions, while technically feasible, is typically positioned as a later step, not least because it would require a reconsideration of IT maintenance procedures and the mandate of SOC and CSIRT teams therein.

In parallel to the rise of security playbooks, there are developments to add automated response functions to endpoint and network detection solutions. Traditionally, such detection tools are used for identifying anomalies and alerting cyber experts once an anomaly is identified. Although this is a good practice, it fails to counter attacks that require a fast response.

To circumvent this drawback of traditional endpoint and network detection tools, vendors have started to introduce relatively simple and small impact response functions in their products. For example, once a network detection tool detects anomalous traffic over a certain connection, it responds immediately by temporarily suspending the traffic flow over the connection. For Gartner, this trend has been the reason to rename the cyber security market that was formerly known as 'network traffic analysis' and refer to it as 'network detection and response' (NDR) since midst of 2020[4]. Although EDR and NDR are interesting examples of automated response, we will not elaborate on these but focus our attention on more generally applicable security orchestration technology and the emerging concept of self-healing.

4 https://www.gartner.com/en/documents/3986225/market-guide-for-network-detection-and-response.

4.1 Automated Playbook Generation

The current state of the art in playbook driven security orchestration and automation has a limitation. Security playbooks need to be designed and maintained by security analysts. In view of this, TNO has been exploring the concept of *automated playbook generation*. This not only relieves maintenance effort on the part of human analysts, but also enables automated resolution of attacks and incidents that were not previously encountered.

In order to automate this task, playbook generation is part of an integrated security control-loop referred to as MAPE (Monitoring, Analysis, Plan, Execute). MAPE was introduced as a generic concept to implement autonomous systems [IBM] and has been widely adopted in the field of automated security. A MAPE-loop can be composed from existing security tools, such as monitoring & detection, cyber reasoning and system configuration tools. The integration of each of these tools into an automated MAPE-loop can be realized by a set of playbooks that are consecutively executed by a SOAR tool.

Some of these playbooks implement relatively small, deterministic functions, such as (re-)formatting data that is exchanged between tools or updating the risk level of the to-be-protected-ICT-environment based on received threat intelligence or security monitoring alerts. However, pre-programming of playbooks is not straightforward in all steps of the MAPE loop. In particular, the step between a suggested CoA and its execution would require an insurmountable number of pre-programmed playbooks.

In fact, pre-programming this step would require the knowledge of any possible CoA that may be needed in the future. Moreover, those playbooks would need to be re-programmed any time an infrastructure component is replaced.

A better solution than pre-programming the execution of any (sequence of) possible CoAs is to dynamically compose a playbook from sub-playbook templates. The reduced complexity of the sub-playbook template enables them to be pre-programmed. The challenge that remains is to develop a playbook generator that can parse any suggested CoA and generate the according composition of sub-playbook templates and fill in the specific details (e.g. IP addresses, host names, credentials).

In 2018, TNO started to explore the options for dynamic playbook generation and built two prototypes for two different scenarios. Figure 4 shows a schematic overview of the automated playbook generation proof of concept.

Looking ahead even further, beyond the automated playbook generation PoCs, automated response technology should evolve into self-protecting and self-healing technical infrastructures. Those infrastructures should autonomously anticipate, withstand and recover from emerging threats and ongoing attacks.

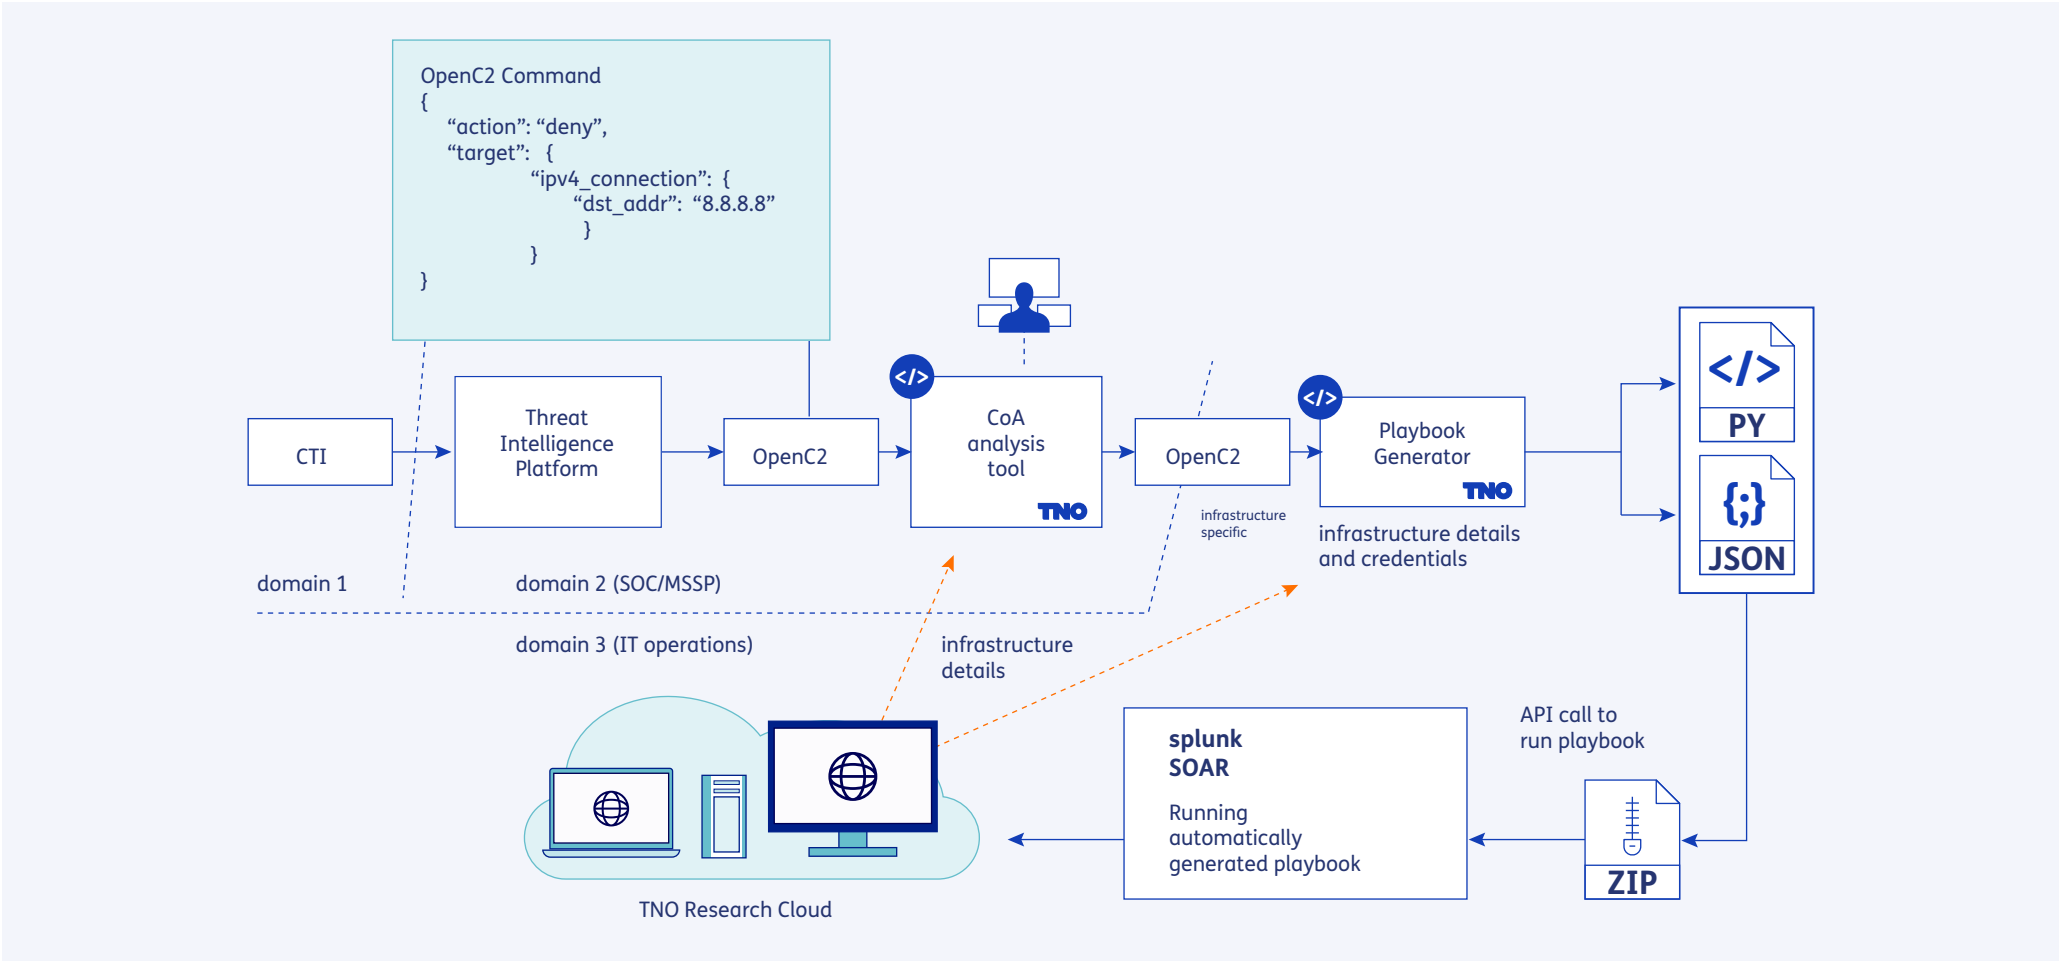

Figure 4: Automated playbook generation extended with OpenC2.

Playbook generation Proof of Concept (PoC)

In the first playbook generation PoC, the 'to-be-protected-ICT-environment' is installed in the TNO Research Cloud, that consists of a cluster of virtualized servers and network switches that provides computation, storage and networking services. The use case of this PoC is to scan the ICT infrastructure for vulnerable assets and automatically resolve the vulnerabilities found.

The ICT infrastructure is described by a JSON description of the servers, the installed software and the network connections. The SecuriCAD tool is programmed to load the JSON model and calculate the most vulnerable attack paths. These attack path calculations are based on attack techniques and vulnerability data that is made available to SecuriCAD and by identifying which vulnerable components are present in the infrastructure model.

Based on the vulnerability scores and attacker profile(s) that are relevant for the ICT infrastructure, SecuriCAD calculates the most critical attack paths. The level of vulnerability of an attack path is expressed in an expected 'time-to-compromise'. Once the attack paths are calculated, SecuriCAD identifies which defensive measures (e.g., patching some of the software components, running a virus scanner or setting a firewall rule) are most effective to increase the time-to-compromise of the most critical attack paths. The proposed measures are sent to the playbook generation software that parses the SecuriCAD output and creates a file indicating the suggested measure(s) for the machine(s) that should be patched or reconfigured. The file contains new playbook code that calls other (sub-) playbooks for the execution of machine specific functions.

The generated playbook is then uploaded into the Splunk Phantom SOAR tool and executed on the ICT infrastructure. This prototype can be extended and made as complex and precise as needed to patch the infrastructure in the best conceivable way. It is possible to rerun SecuriCAD on an updated version of the infrastructure model which includes the modification made by the tool.

A second iteration of the automated playbook generation prototype was made in the Cyber Security Noord-Nederland project, by implementing the OpenC2 interface standard. Figure 4 shows a schematic overview of the extended automated playbook generation prototype.

In the second PoC the playbook generation prototype was extended by implementing an OpenC2 interface between the TIP and the CoA analysis tool. These OpenC2 messages include sequences of (action, target)- instructions that specify which action (e.g., deny access) needs to be performed for which target (e.g., a specific IP address). The SOC analyst receives the messages and for each of them he can select which actuator (e.g., a firewall, specific host) should perform the described actions. Then the prototype software implements the commands (e.g. create file, deny IP, or deny URL). Based on the choice of the analyst, a specific Phantom playbook is generated and executed on the selected ICT component.

4.2 Self-healing for Cyber Security (SH4CS)

Complementary to the mentioned security orchestration and automation technology is an emerging technology that is referred to as self-healing for cyber security[5]. The term 'self-healing' was first coined in 2001 in the article titled 'The dawning of the autonomic computing era' [IBM]. The drivers for development of the autonomic computing concept were similar to those for automated security. Autonomic computing is aimed at overcoming a wide range of operational problems due to the ever increasing complexity of IT systems and their labor-intensive maintenance.

Autonomic computing is comprised of eight characteristics of which self-healing and self-protection are most notable in the context of cyber security. Self-healing and self-protective systems are defined as systems that are able to recover from a failed component and to prevent and recover from unauthorized access. In this paper, we refer to self-healing for cyber security (SH4CS) as the broader definition of self-* characteristics that are applied to autonomously secure cyber systems (i.e. the systems initiate and execute security actions themselves). This technology complements security actions by human experts and enable them to focus on more complex actions.

The self-healing for cyber security concept is inspired by defense patterns of the human immune system. So far, basic self-healing functions are implemented by modern container (management) platforms Docker, Kubernetes and Rancher. These platforms are particularly useful for providing high flexibility with respect to application (re-)deployment and centralized monitoring, scheduling and scaling of individual containers[6]. Built on those features each platform also provides out-of-the-box self-healing features. These features are still basic in the sense that they periodically execute if-then-rules, where the if- and the then-parts are executed by one platform.

SH4CS extends these basic functions to more advanced functions that can combine arbitrary if clauses (e.g. monitoring events from a SIEM and/or from a security monitoring system) with a then-action that can be executed by any container platform. By combining if clauses the concept can also be applied in a risk adaptive manner. The SH4CS software[7] that TNO developed recently demonstrates how the self-healing concept can be extended from basic if-then-rules to decentralized MAPE-loops.

5 https://www.abnamro.com/nl/nieuws/abn-amro-eerste-afnemer-van-innovatieve-zelfhelende-cybersecurity-software.
6 https://kubernetes.io/docs/concepts/workloads/pods/.
7 https://github.com/TNO/self-healing-4-cyber-security.

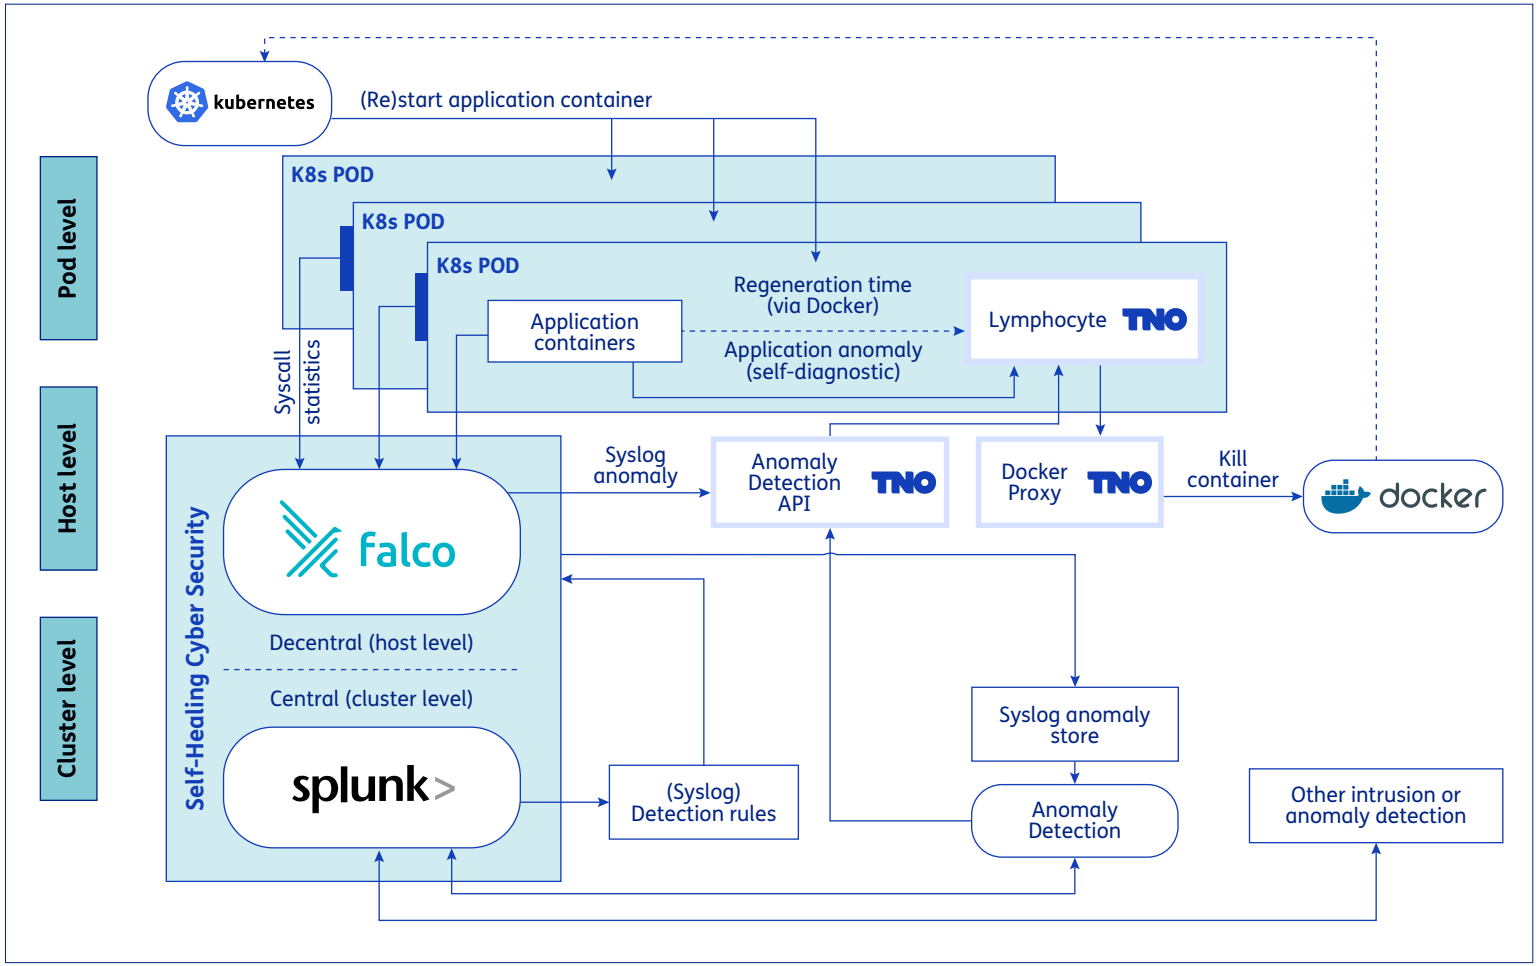

Figure 5: Adaptive SH4CS components and interaction loops.

Self-healing for cyber security Proof of Concept

Kubernetes (K8s) supports the function to automatically detect a failing application container and restart it. In the SH4CS PoC this is used to implement a periodic container regeneration concept (similar to how human body cells regenerate themselves), once it is complemented with software that checks the remaining time-to-regeneration for the container and instructs Docker to kill the container upon expiry.

Based on this capability to regenerate containers, the SH4CS software can also influence the moment at which regeneration will take place. This can be timer based, but SH4CS can also regenerate the container once a certain cyber security event is detected. In a proof of concept, the Falco security monitoring tool for Kubernetes was used to detect and alert specific events (Falco does so by monitoring all system calls that are made by the application containers and generating an alert once a malicious syslog is detected). By letting the detection of a cyber security event from a specific container trigger its regeneration, the SH4CS concept becomes adaptive to cyber security events.

Figure 5 shows the SH4CS software components developed by TNO (Lymphocyte, Anomaly detection API and Docker Proxy) and the MAPE interaction loops between those software components and Kubernetes, Docker and Falco platform software. It also shows that the concept is complementary to (and not instead of) centralized automated response concepts, such as a Splunk SOAR solution.

The translation of multiple inputs (i.e. residual time-to-regeneration and cyber security triggers) to multiple possible responses (e.g. immediate container kill or reduction of residual life-time, or any other action that Docker can perform) requires security reasoning logic. As opposed to the playbook generation concept, the prototyped SH4CS logic is pre-programmed and decentralized; each application container is accompanied by a SH4CS container, or Lymphocyte, that executes the logic. The decentralized implementation coincides with the decentralized operation of the human immune system, which is essential for a fast response and scalable defensive actions.

5 Take-aways

To summarize TNO's vision and work in the field of automating cyber security operations, we conclude this whitepaper with the following take aways:

- There is an evident need to pursue heavy automation in cyber security operations, not only to bridge the significant gap between attackers and defenders but also to relieve pressure on scarce security resources in a challenging labor market.
- While a great variety of cyber defenses might be automated to more or less extent, TNO believes that there is particular potential in the automation of SOC and CSIRT operations. Promising directions to this end include the further development of automated security reasoning and automated response technologies.
- Automation in cyber security will not happen overnight. Apart from the need to mature automation technologies further, organizations will need to develop trust in software solutions that reconfigure technical assets and security controls without human intervention. Many will also need to rethink their operations and maintenance procedures. In most cases, therefore, a gradual transition towards increasingly automated cyber security operations seems likely.
- While automation technology will play an instrumental role in the further evolution of cyber security, TNO does not envisage that it will replace the human analyst altogether. Rather, the vision is that automation solutions can help such analysts become more effective, either by relieving them from (what could be) repetitive tasks or by supplying them with (contextual or threat related) insights that human experts could likely not develop at the same pace.

As a final note, the authors emphasize that the need to automate cyber security operations is widely recognized and has led to an array of innovation initiatives across the globe. This whitepaper addressed some of the innovations that TNO is heavily invested in, but readers are encouraged to also explore the ongoing work of solution providers and other knowledge institutes in his field.

6 References

[Accenture] K. Bissell, J. Fox, R.M. LaSalle and P. Dal Cin, How aligning security and the business creates cyber resilience, Accenture, November 2021, https://www.accenture.com/_acnmedia/PDF-165/Accenture-State-Of-Cybersecurity-2021.pdf#zoom=40.

[CACAO] Collaborative Automated Course of Action Operations (CACAO) for Cyber Security TC, OASIS, https://www.oasis-open.org/committees/tc_home.php?wg_abbrev=cacao.

[ENISA] Threat Landscape Report 2017, ENISA, Final Version 1.0, January 2018, https://www.enisa.europa.eu/publications/enisa-threat-landscape-report-2017.

[Frost] 2017 Global Information Security Workforce Study Benchmarking Workforce Capacity and Response to Cyber Risk, Frost & Sullivan, https://iamcybersafe.org/wp-content/uploads/2017/06/Europe-GISWS-Report.pdf.

[GAP] Cyber security Talent: The BIG GAP in Cyber Protection, Capgemini Digital Transformation Institute, February 2018, https://www.capgemini.com/resources/cybersecurity-talent-gap/.

[Gartner] Security Orchestration, Automation and Response Solutions Reviews and Ratings, Gartner, https://www.gartner.com/reviews/market/security-orchestration-automation-and-response-solutions.

[IBM] The dawning of the autonomic computing era, IBM Systems Journal, vol. 42, no. 1, pp. 5-18, 2003.

[MAL] Foreseeti AB, Meta Attack Language - The open source platform for creation of cyber threat modelling systems, https://mal-lang.org/.

[OpenC2] Open Command and Control (OpenC2) TC, OASIS, https://www.oasis-open.org/committees/tc_home.php?wg_abbrev=openc2.

[SecuriCAD] SecuriCAD for automated threat modelling and attack simulations, Foreseeti AB, https://foreseeti.com/.

[SOCCRATES] EU H2020 SOCCRATES project, https://www.soccrates.eu/.

Authors

Frank Fransen
Senior cyber security scientist

frank.fransen@tno.nl

Bart Gijsen
Senior cyber security consultant

bart.gijsen@tno.nl

Richard Kerkdijk
Senior cyber security consultant

richard.kerkdijk@tno.nl

Reinder Wolthuis
Senior cyber security consultant

reinder.wolthuis@tno.nl

Ruggero Montalto
Project Manager

ruggero.montalto@tno.nl

TNO innovation for life

Context

This whitepaper was written as part of the Program *Cybersecurity Noord-Nederland.*

The Program *Cybersecurity Noord-Nederland* has received funding from the RSP of the Province of Groningen and the municipality of Groningen.

PROGRAMMA CYBERSECURITY NOORD-NEDERLAND

tno.nl/automated-security