SOCCRATES

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

D7.2

Pilot Application and Evaluations

Deliverable type:	Report
Contributing work packages:	WP7 Pilot Application and Evaluations
Due date of deliverable:	31/10/2022
Submission date:	31/10/2022
Dissemination level:	PU

Responsible organisation:	mnemonic
Editor:	Siri Bromander
Revision:	1.0

Abstract	This deliverable presents the results from pilots in the SOCCRATES project: the pilot implementation, alterations to the initial plan, measurement results and ethical and privacy considerations for the pilot implementations. The goal of the pilots was to validate the SOCCRATES platform in a realistic environment.
Keywords:	SOCCRATES, Pilots, Evaluation results

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement no. 833481
Call H2020-SU-ICT-2018 • Innovation Action • Start date: September 1st, 2019

Classification level: Public

Management summary

The SOCCRATES project developed and implemented a security automation and decision support platform for security operations centers (SOCs) and computer security incident response teams (CSIRT). In order to ensure that the platform is fit for purpose, the project carried out three pilots to validate the platform in realistic environments. The plan for pilot implementation and execution is described in a different deliverable: D7.1.

This deliverable describes the pilot implementations for all three pilots and the measurement results. The lessons learned from pilot implementations and execution are described in a different deliverable: D7.3 Best practice guide.

There were alterations to the implementations and execution of pilot 2 and 3 due to covid and unforeseen implementation isses on the pilot sites. However, all pilots were effective and together created foundations for KPI measurements.

The main part of this deliverable is detailed descriptions of the pilot implementations and measurement results on each KPI.

The reader is assumed to be familiar with D7.1.

Contributor(s)	Siri Bromander (mnemonic) Martin Eian (mnemonic) Philip Scheel (mnemonic) Richard Kerkdijk (TNO) Frank Fransen (TNO) Piotr Kijewski (Shadowserver) Irina Chiscop (TNO) Maciej Kosz (Vattenfall) Gabriela Bodea (TNO) Paul Smith (AIT)

Reviewer(s)	Mathias Ekstedt (KTH)

Security Assessment	**See deliverables table for security assessment requirements**
Approval Date	27/10/2022
Remarks	No comments

TABLE OF CONTENTS

1 INTRODUCTION – RATIONALE OF THIS DOCUMENT 6
1.1 The SOCCRATES project 6
1.2 This deliverable 7
1.3 Structure of this deliverable 7
2 USE CASES 8
3 PILOT IMPLEMENTATIONS 9
3.1 First pilot 9
3.1.1 Description of pilot site 9
3.1.2 Data sources 10
3.1.3 Hardware setup 10
3.1.4 Software setup 10
3.1.5 Experiences altering plan for second pilot 10
3.2 Second pilot 11
3.2.1 Mnemonic pilot site 11
3.2.2 Vattenfall pilot site 12
3.2.3 Shadowserver pilot site 14
3.3 Demonstrator 15
4 PILOT EVALUATION 18
5 EVALUATION OF KPIS 22
5.1 KPI1 - Prototype successfully generates a correct model of the pilot infrastructure (all UCs) 22
5.2 KPI2 - Prototype successfully improves detection capability (incl. detecting attacks in encrypted traffic that was not detected otherwise) (UC1) 23
5.3 KPI3 - Prototype successfully reduces total time of detection and response (UC1) 28
5.3.1 Time to detection 28
5.3.2 Time to response 29
5.4 KPI4 - Prototype successfully reduces total time from collection, analysis and determining pro-active mitigation new CTI (UC2) 31
5.5 KPI5 - Prototype successfully analyses the vulnerability of an asset in the context of the whole infrastructure, that could cause an incident otherwise (UC3) 33
5.6 KPI6 - Prototype reduces the time spent on analyses of the severity of vulnerable assets (UC3) 35
5.7 KPI7 - Prototype proposes at least one course of action that is more efficient/appropriate than what would otherwise have been proposed by a human analyst using templates (UC4 & UC5) 36
5.8 KPI8 - Prototype reduces the time spent on analysis and determining a course of action (UC4 & UC5) 36
5.9 KPI9 - Prototype successfully detects 25 new DGA seed from known malware family 37
5.10 KPI10 - Prototype successfully detects 5 new, previously unknown malware DGA 37
5.11 KPI11 - Prototype successfully clusters 30 malware families 38

5.12 KPI12 - Prototype successfully generate 20 new detection signatures/rules 39
5.13 KPI13 - New API successfully used by 10 partners 39
6 ETHICS AND PRIVACY 40
7 CONCLUSION 41
8 REFERENCES 42
9 ABBREVIATIONS 43
ANNEX A - METRICS FOR PILOT PERFORMANCE EVALUATION 45

1 Introduction – Rationale of this document

This section introduces the SOCCRATES project and defines the goals of this deliverable.

1.1 The 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.

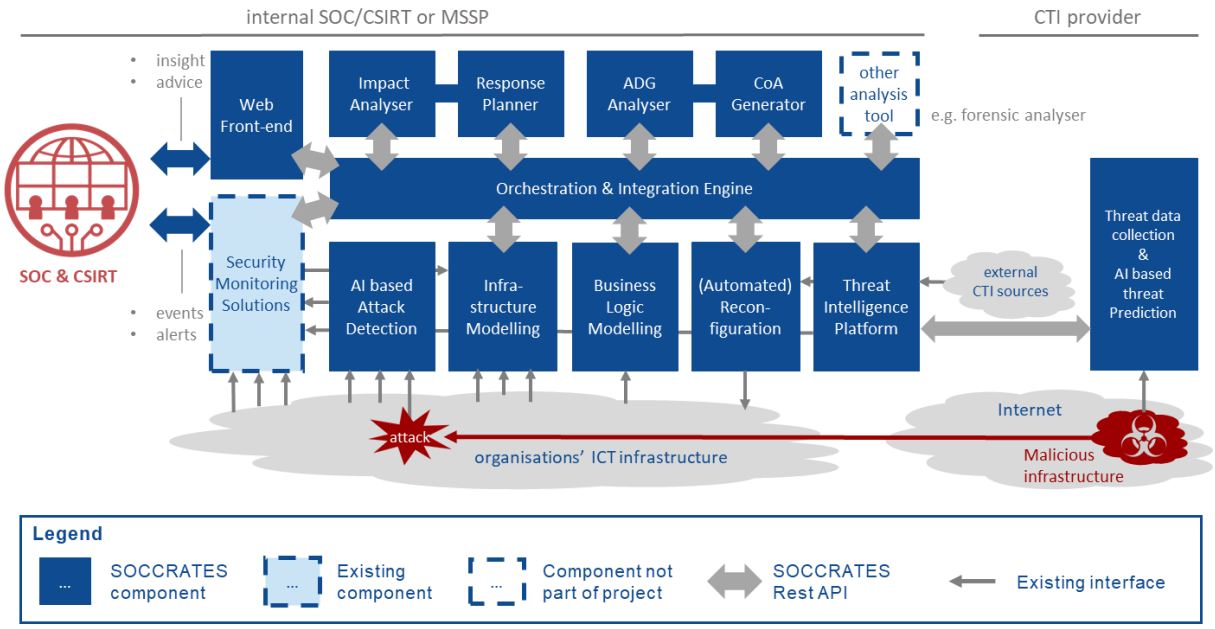

Figure 1-1 – The SOCCRATES platform

The SOCCRATES platform (see **Figure 1-1**) 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.

SOCCRATES has the following concrete project objectives:

1. Deliver the SOCCRATES platform consisting of an orchestration function and a unique integration of innovative background solutions that seamlessly work together.
2. Show that the SOCCRATES platform can improve SOC operations by evaluating the SOCCRATES platform in two diverse real-life pilot environments.
3. Examine and illustrate the benefits of automation for selected SOC activities to help manage the cyber security skills gap in organizations.
4. 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.

1.2 This deliverable

This deliverable describes the pilot execution and evaluation results of the SOCCRATES project. SOCCRATES planned to perform three pilots. The goal of the first pilot was to validate a number of SOCCRATES components in a realistic environment. The implementation and execution is described in this deliverable. The goal of the second pilot was to validate the complete SOCCRATES platform in a realistic environment. The implementation and execution is described in this deliverable. The goal of the third (demo) pilot was scoped to cover the full SOCCRATES platform except the parts covered in the first pilot in order to showcase the functionality and results of the SOCCRATES platform to stakeholders and external parties.

The KPI measurements described in deliverable D7.1 were not completed in pilot 2 alone due to implementation issues, but in total covered by pilot 2 and pilot 3 (demo). The results from measurements are described in this deliverable. Lessons learned and best practices from execution of pilots are described in D7.3.

1.3 Structure of this deliverable

In section 2 we describe the use cases which SOCCRATES has aimed to solve and served as the foundation for the pilot execution. Section 3 describes the pilot implementation specifics and Section 4 evaluates the pilots. Section 5 details the results from measurements of the defined KPIs. Section 6 describes relevant etchics and privacy consideration and we conclude in section 7.

2 Use Cases

The platform is structured to solve the chosen SOCCRATES use cases and the pilots are set up to test these. The use cases are described in detail in SOCCRATES Deliverable 2.1 and details can be found there. For the convenience of the reader, a short description of them are included in the following.

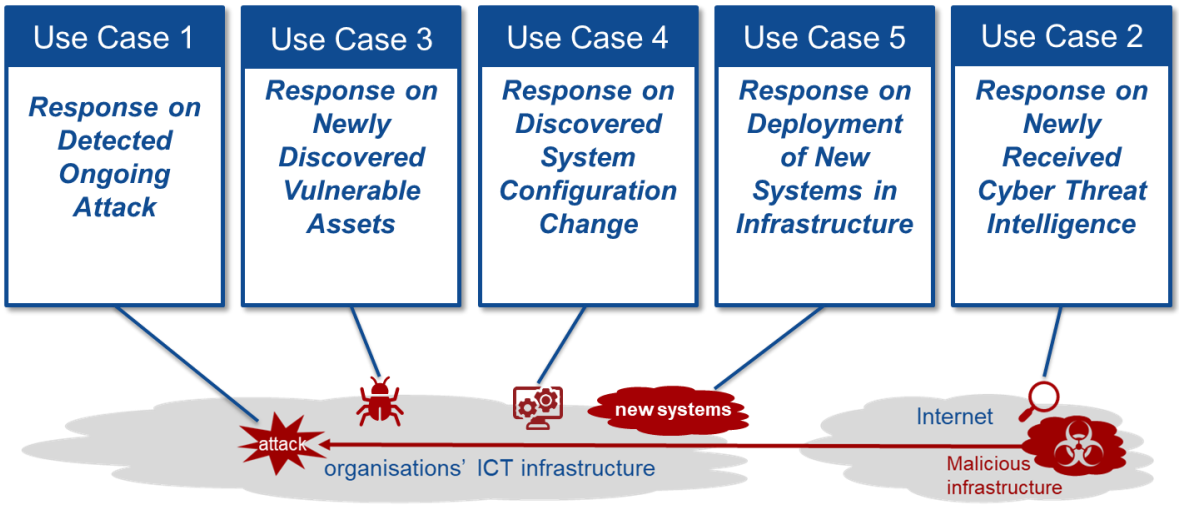

Figure 2-1 The SOCCRATES Use Cases.

- ***Use Case 1: Response on Detected Ongoing Attack***
  Detect ongoing attacks and automatically analyse the attack, automatically determine the best response, and initiate deployment of the selected response.
- ***Use Case 2: Response on Newly Received Cyber Threat Intelligence***
  Continuously collect new threat information, automatically analyse the potential business impact and determine best options for proactive mitigation.
- ***Use Case 3: Response on Newly Discovered Vulnerable Assets***
  Automatically detect vulnerabilities on assets in the ICT infrastructure, assess if they enable new attack paths, determine and initiate mitigation actions.
- ***Use Case 4: Response on Discovered System Configuration Change***
  Automatically detect configuration changes on assets in the ICT infrastructure, assess if they enable new attack paths and determine if action is needed.
- ***Use Case 5: Response on Deployment of New Systems in Infrastructure***
  Automatically detect introduction of new systems to the ICT infrastructure. Automatically assess the new situation and determine if security measures are needed.

3 Pilot implementations

The planned pilot setups are described in D7.1 [D7.1]. The following details the information given in D7.1, and describe the deviations from the planned setup. The lessons learned from the pilot executions are described in Deliverable 7.3 [D7.3].

3.1 First pilot

The goal of the first pilot was to deploy and test the AI based Attack Detection (AAD) components of the SOCCRATES project. These components would provide analysis of network, DNS and Syslog data to detect malicious activity. It would also serve as a test of the deployment strategy for the full-scale pilot.

Design-wise, the goal was to provide enough data from a semi-realistic environment to the analysis engines in the AAD component, and in this way avoid having to synthesize data.

3.1.1 Description of pilot site

The pilot site was an already existing network used for demonstrations and tool testing in mnemonic. While it does not contain any real end users, it simulates a miniature Small Business network with common components such as Windows clients and active directory services, along with detection tools for malicious software and common network analysis tools.
The most relevant pre-existing components used by the Pilot was the suricata open source IDS sensor, which is used to generate netflow, DNS and network traffic detection. Additional virtual machines were added for use during the pilot. In order of implementation:

3.1.1.1 A virtual machine with elasticsearch and kibana

This component was installed for log retention, distribution and debugging purposes.

3.1.1.2 A virtual machine with docker

This component was host to all AAD components, and allowed us to validate processes surrounding the future use of docker swarm as a means to deploy the components of SOCCRATES.

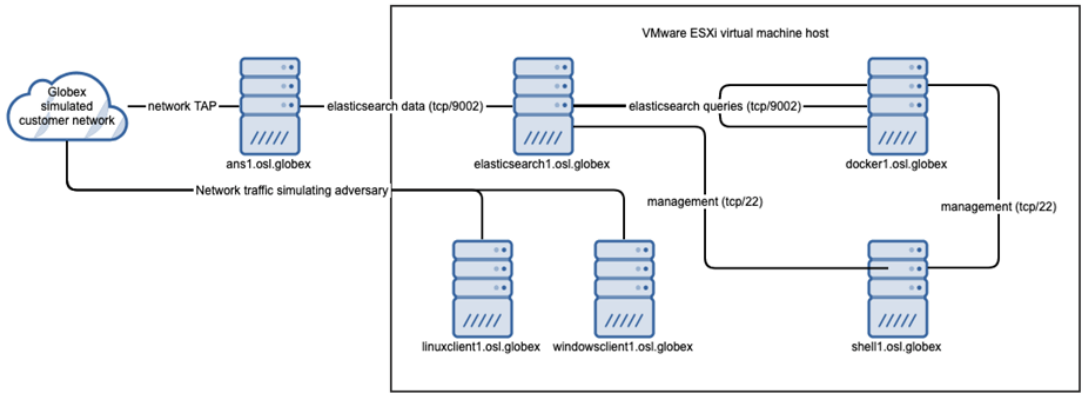

Figure 3-1 The mnemonic First Pilot environment.

Classification level: Public

3.1.2 Data sources

The data sources used for the initial Pilot was DNS, Network traffic (netflow) and Syslog data for a single affected host. The DNS and Netflow data was provided by a Suricata IDS sensor. This IDS can produce json-formatted summaries of several protocols, where netflow and DNS data is included. Since Suricata IDS is an open source product, and has a standard json format, the field names used by Suricata in these logs were also used as the standard for SOCCRATES input data.
Syslog data was provided by a standard CentOS machine (linuxclient1.osl.globex), with standard logging settings. This machine was central in any simulated malicious activity attempted during the pilot.

3.1.3 Hardware setup

All components mentioned under “Description of pilot site” were hosted on a single VMware ESXi server.

3.1.4 Software setup

An Elastisearch instance was provided by mnemonic for use as a data source by all software under testing. The Elasticsearch instance was used as a common instance of data for all consumers in the AAD component of the platform. The goal of keeping data in Elsticsearch was to deduplicate the data required by the platform components, and have a common pipeline for data. This also allowed for development of common interfaces for interacting with the platform data (i.e. deduplication of code.
A docker host was installed, running on a CentOS virtual machine in our environment. Two clients were created to simulate attacker and victim in the environment, without negatively affecting other parts of the shared infrastructure. We also developed a simple script to simulate user activity on the linux host, which ensured that there would be “background noise”/normal traffic for the attacker to hide in. Without such activity, any traffic to/from the victim or attacker host would be considered an anomaly. Software updates were provided using a github repository.

3.1.5 Experiences altering plan for second pilot

Based on the experience with deploying the AAD component in the first pilot, the SOCCRATES project identified that more emphasis had to be placed on unit testing, integration testing and deployment testing. Moreover, only after successful integration and deployment testing the software of the SOCCRATES Platform was release for deployment in the pilot environents. The results of the testing and verification have been described in deliverable D6.3 [D6.3]. The testing strategy of the SOCCRATES project consisted of unit testing of each of the SOCCRATES components by the individual work package that developed the component, and integration and deployment testing by WP6. In particularly, WP6 setup a testing environment and process for:

- *Platform Integration testing* – Testing to verify that components can be integrated as specified in D2.3 [D2.3], and to verify that the SOCCRATES platform can execute the SOCCRATES Use Cases (run a whole workflow successfully in Activiti).
- *Platform deployment testing* – Preparation of all software for deployment, and testing to verify that the SOCCRATES Platform can be deployed as a whole. As the SOCCRATES components are developed by different partners, this testing activity focuses on testing the combined deployment of all components in a single staging environment. This included containerizing all software components and setting up a GitLab CI/CD pipeline for deployment in the TNO Research Cloud.

Classification level: Public

Furthermore, a network traffic generator NEtwork Data producer (Ned) was developed to test the consumption by components like the AAD and the IMC of large amounts of network data.

3.2 Second pilot

The second pilot was implemented on three pilot sites: mnemonic (MSSP), Vattenfall and Shadowserver. The mnemonic and Vattenfall pilot sites were to install the full SOCCRATES platform, and the Shadowserver pilot was implemented to provide an API for querying and collecting relevant data. The plan was to perform all of the KPI measurements during the second pilot.

3.2.1 Mnemonic pilot site

The second pilot utilized mnemonic's production environment, in an attempt to give a realistic site with a sufficient amount of traffic to challenge the SOCCRATES components. mnemonic is a technology-heavy company, 90+% technical personell MSSP company of around 250 employees, most with either one or two client machies. The monitoring covers the Norwegian offices, where the majority of employees and infrastructure is located. Since we are heavy on automation and virtualized machines, this comes out to between 10k and 20k assets (clients + servers) in our environment. The environment is described in Figure 3-2.

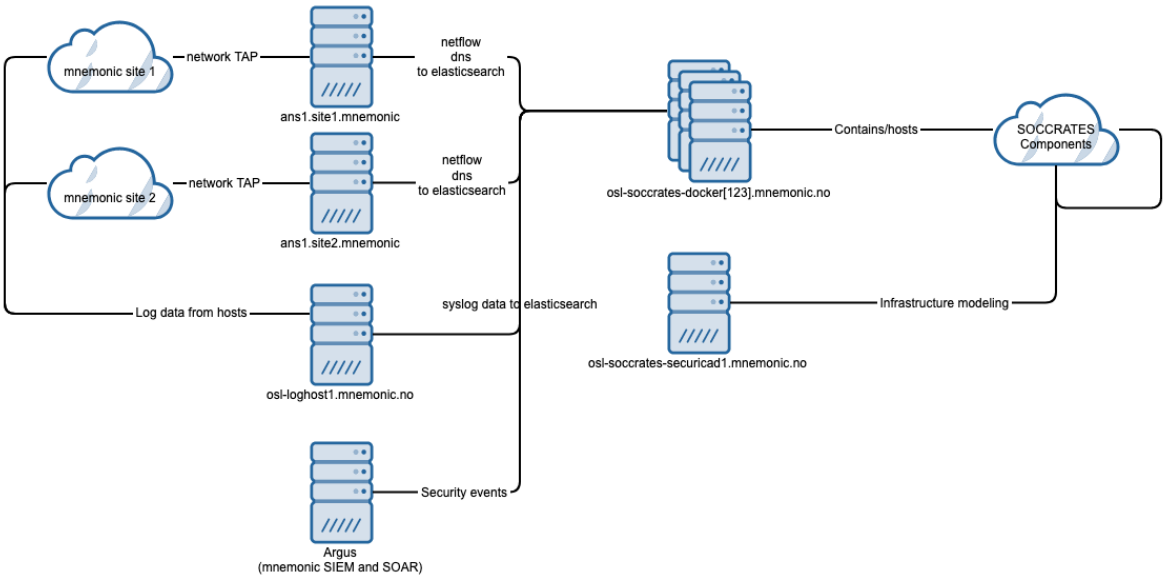

Figure 3-2 The mnemonic Second Pilot environment.

3.2.1.1 Description of pilot site

The goal of the second pilot was to deploy the full SOCCRATES platform, and test the various KPIs as described in Deliverable 7.1 [D7.1]. Since the goal of the second pilot was to test the SOCCRATES platform in a real life environment, the second pilot was based on data from the production network of mnemonic.

3.2.1.2 Data sources

The data sources used for Pilot 2 was identical to that used in Pilot 1. Netflow data was provided by the suricata output for netflow data. DNS data was provided by the suricata output for DNS data.

Syslog data was provided from one of the centralized Syslog servers in our environment.
It is important to note that the network data source in this pilot site differs from that seen in other parts of the project. The data comes from centralized sensors, which are fed by taps at central switches in the infrastructure. This means we have broad coverage at key points between networks, but little coverarge far out in the networks.

3.2.1.3 Software Setup

Software applications were distributed as individual docker images, with configuration distributed as a git repository. Additionally, site-specific configuration was kept in its own sub-repository, which allowed for version-controlled configuration changes. Due to architectural reasons, the one exception was the SecuriCAD infrastructure modeler, which uses container technology itself. Due to this, the software was installed according to manufacturer's instructions on a dedicated machine. Management of the pilot applications was done by way of one of the native docker-managers, Portainer. This tool also allowed for central management and log investigations of the applications.

3.2.1.4 Hardware setup

In terms of hardware, the components of the pilot were run on 2 physical machines running VMware ESXi. These, in turn, hosted 4 virtual machines, 3 machines for docker and 1 machine for SecuriCAD. In terms of distribution, it was expected that a lot of the processing power would be used by specific docker-containers. As such, the more powerful physical machine host only hosted the primary docker host, while the second physical machine hosted two docker hosts and the SecuriCAD machine.

3.2.1.5 Analysts

The Pilot site provided 4 analysts for evaluation purposes, distributed between different career stages. The evaluators were, as follows:

- A Threat Analyst
- A senior analyst
- Two SOC analysts

3.2.1.6 Pentest

For evaluation purposes, the mnemonic internal pentesting team was engaged to simulate an attacker in our infrastructure during the evaluation-phase of the pilot. The team made a plan for simulating an advanced persistent threat, and would have executed realistic actions according to this plan, with some modifications to adjust for discoverability underway, if the AAD components should not discover the activity. This would have allowed for comparison between the service delivery based on IDS and regular analysts, and the SOCCRATES platform.

When the AAD components would later be unable to scale to the amount of data in our environment, the plan was made available to the analysts to simulate investigating the pentest.

3.2.2 Vattenfall pilot site

SOCCRATES platform deployed at Vattenfall was fed with data from selected parts of the production environment. Due to security reasons, infrastructure which was built for the SOCCRATES platform

needed to be heavily isolated from the rest of environment. One business application was chosen for in-depth business process and function modelling needed to demonstrate SOCCRATES capabilities.

3.2.2.1 Data sources

Following data sources were used with a filtering applied limiting data to relevant to SOCCRATES platform pilot only:

- Network flow telemetry covering both east-west and north-south traffic
  - NetFlow records exported from virtualization hypervisors (Data Center)
  - NetFlow-equivalent records exported by a sensor capturing raw network traffic from the SPAN port of a physical switch (Office network)
- DNS logs from internal DNS servers
- Vulnerability Scan data

To avoid data duplication and to standardize data format, network flow and other records provided to SOCCRATES platform were stored in Elasticsearch instance from where they are made available for other components.

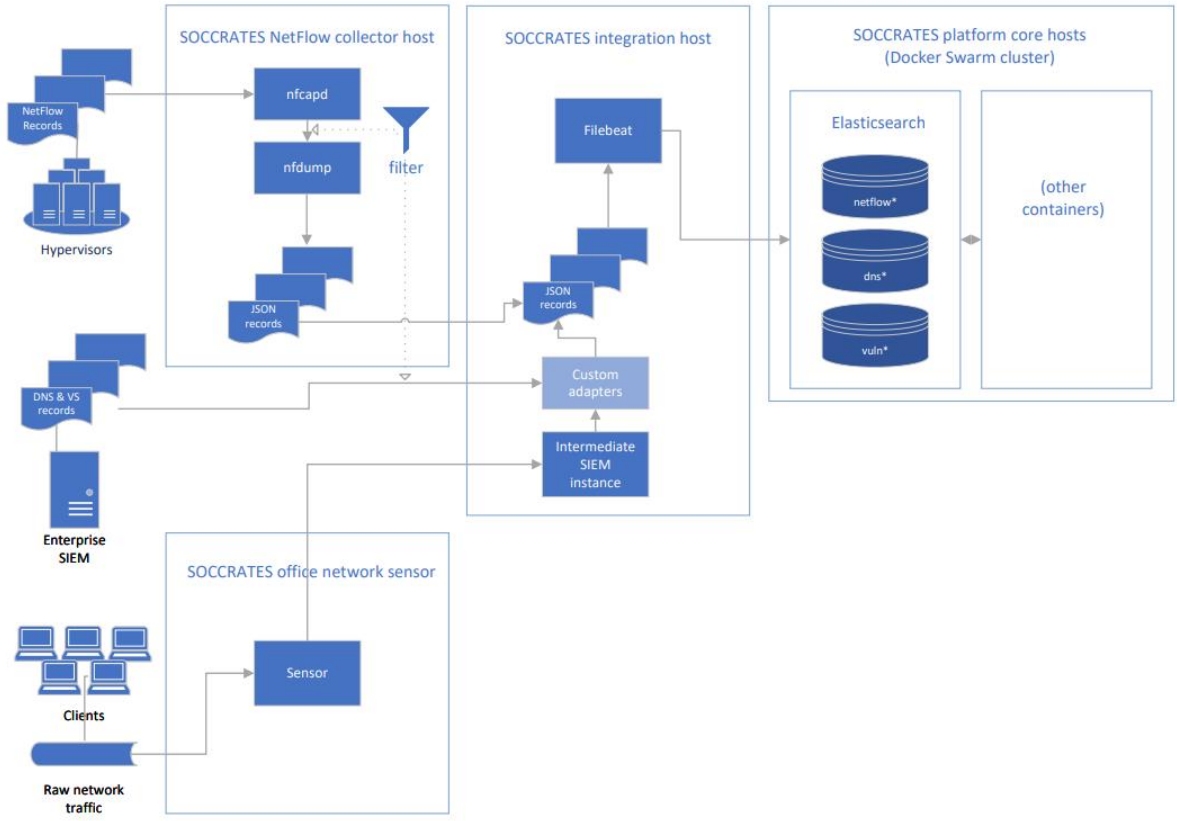

Figure 3-3 The Vattenfall Second Pilot environment.

3.2.2.2 Software Setup

Similarly to Mnemonic pilot site, software applications were distributed as individual docker images, with configuration distributed as a git repository and an exception of SecuriCAD infrastructure modeler software which was installed on a separate host.

Due to internal security requitements, it was not possible to download SOCCRATES software directly on the target hosts. Additional server was used to retrieve docker images and git repository, perform security checks and finally upload them to the restricted SOCCRATES host.
It was also necessary to introduce additional components needed for data ingestion, transformation and filtering.

3.2.2.3 Hardware setup

In terms of hardware, Table 3-1 gives a list of hosts that were used for SOCCRATES platform pilot at Vattenfall.

Table 3-1 The Vattenfall Second Pilot hosts.

Host	Type	Description / Purpose
SOCCRATES core host 1	Virtual	Docker host for running SOCCRATES components
SOCCRATES core host 2	Virtual	Docker host for running SOCCRATES components
SOCCRATES ADG host	Virtual	Server running non-dockerized componens: SecuriCAD software
Network flow collection host – Data Center	Virtual	Server receiving, processing and transforming NetFlow records
Network flow collection host – Office network	Physical	Sensor host capturing network traffic from SPAN port and creating NetFlow-equivalent records
Intermediate SIEM host	Virtual	Host providing intermediate filtering, conversion and protection layer between SOCCRATES platform and enterprise SIEM system
Software download host	Virtual	Server allowed to download files from external SOCCRATES repository

3.2.2.1 Analyst

The Pilot site provided 3 SOC analysts for evaluation purposes. The evaluators were, as follows:

- A Threat Operations analyst
- An Incident Response (CSIRT) analyst
- SOC analyst (Tier-1)

Due to pilot implementation obstacles and platform not achiving fully operational state, evaluation by analysts was limited to the SOCCRATES demonstrator.

3.2.3 Shadowserver pilot site

The Shadowserver Pilot site has been operating independently from the mnemonic and Vattenfall sites. Unlike the mnemonic and Vattenfall pilots it operates only one component - the Threat Data Collection and Threat prediction component that involves analysing Shadowserver sandbox domain traffic, with a specific focus on analysing DGA domains.

The DGA Detective tool is operating on a dedicated VM server. The DGA Detective runs with a gRPC API. Shadowserver uses a Redis queue to submit domain names for evaluation that are then made available as individual Redis keys using the provided domain name.

As mentioned in D4.3[D4.3] the Shadowserver pilot has been operating since mid 2021 and delivering concrete real-world results that have been shared operationally with the outside world, contributing to new DGA detections, sinkholing of threats and the sharing of infection information with the outside world through free Shadowserver daily feeds.

Aside from Shadowserver threat feeds, data from that component is shared with external entities and the other SOCCRATES pilot sites via a RESTful based API that has new specific SOCCRATES methods implemented. Access to the API was provided to TNO, Mnemonic and Vattenfall for development and integration testing.

As a result API improvements were made by adding features to filter out potential FP. Vattenfall also engaged in querying the API through the Middleman tool which identified potential areas for improvements in the tool as well as identified some potential FP domains which had to be whitelisted.

For the SOCCRATES pilots, the API can be used in two ways. The first is an enrichment service for the Threat Intelligence Platform. The second way is to download new suspect domains from the API and apply those as detectors on the network traffic obtained in the pilots (implemented by the Middleman tool in SOCCRATES).

Aside from access to project partners, API access to external partners was also provided. This was announced in an invitation sent by Shadowserver to National CSIRTs on September 15th to apply for an API key to access the SOCCRATES methods. As of September 29th, 15 National CSIRTs applied.

3.3 Demonstrator

The goal of the demonstrator was to demonstrate how the SOCCRATES platform components collaboratively create situational awareness and option awareness, so that the SOC analyst can take an informed decision on further action. In deliverable D6.4 [D6.4] the SOCCRATES demonstration environment has been described. In this section a summary is given.

The SOCCRATES demonstration environment consists of two parts. The SOCCRATES Platform and the Infrastructure Under Control, see Figure 3-4. For the Infrastructure Under Control it was not only necessary to define and implement the IT infrastructure, but also define a business story in line with the assets in the IT infrasructure. This was necessary to be able to make a business model for the Business Impact Assessment component. In particularly, we had to create a graph with dependencies between the technical assets and the business functions and processes. For the Demonstration we decided to adopt ArchiMate models of a fictitious insurance company ArchiSurance[1]. The Open Group defined the ArchiSurance as a Case Study for ArchiMate enterprise modeling language.

[1] https://github.com/archimate-models/archisurance

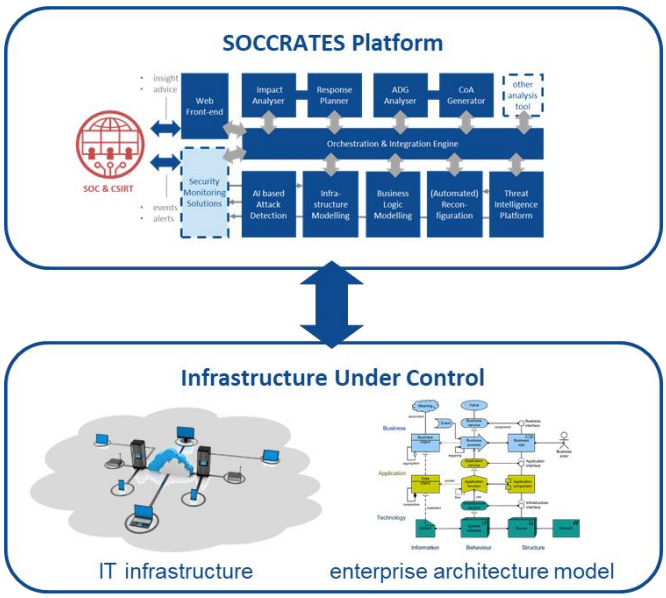

Figure 3-4: SOCCRATES demonstration environment

The infrastructure under control was designed and setup by WithSecure in AWS. In Figure 3-5 the IT architecture is depicted that was based on ArchiSurance Case Study. It consists of a couple of subnets with office PCs, servers, a firewall, and IT management systems (e.g. AD and Withsecure Elements).

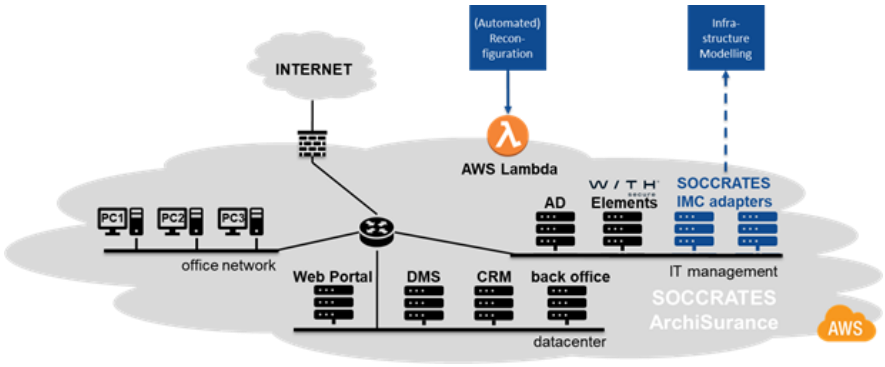

Figure 3-5: SOCCRATES ArchiSurance IT Architecture with SOCCRATES Platform integration

For integration with the SOCCRATES Platform the IMC RADAR Adapter has been deployed in the AWS environment to push data from WithSecure™ Elements Vulnerability Management to the IMC. Also the IMC AD Adapter has been deployed to push Active Directory data from the AD to the IMC. To demonstrate automated response to an ongoing attack an AWS Lambda interface was developed which can be used to request to (a) stop a VM, (b) isolate a host via EDR, and (c) add filtering rules to the firewall.
The SOCCRATES Platform has been deployed to operate on the Infrastructure Under Control. The deployment of the SOCCRATES Platform consists of the following:

- *ADG Analyser* deployed on AWS hosted by foreseeti,
- *TIP* and *IMC* deployed in AWS hosted by mnemonic, and
- the rest of *SOCCRATES Platform* deployed in Azure hosted by TNO.

Figure 3-6 is a conceptual overview of the the SOCCRATES Platform deployment.

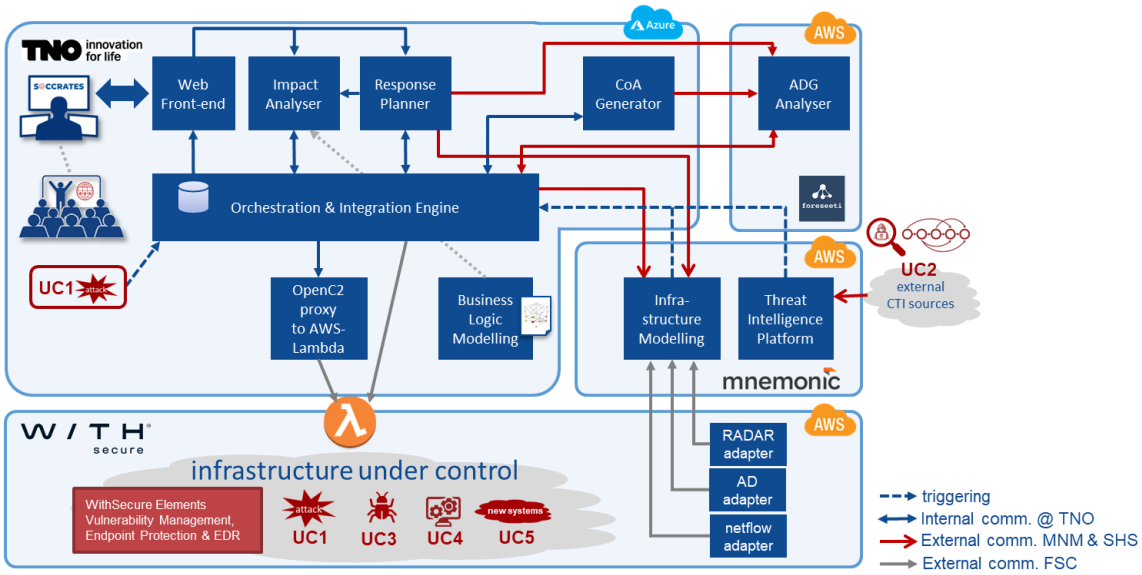

Figure 3-6 : Overview of the SOCCRATES Demonstration environment.

The demonstrator of the SOCCRATES Platform is focussed on demonstrating use cases 1 (all three scenario's), 2 and 3. For ease of demonstration a Control Panel has been added to the SOCCRATES Web Front End for triggering the demonstration use cases, see Figure 3-7. The SOCCRATES Web Front End also provides a generic capability to trigger workflows.

A detailed walk through of the SOCCRATES demonstration use cases is provided in deliverable D6.4 [D6.4].

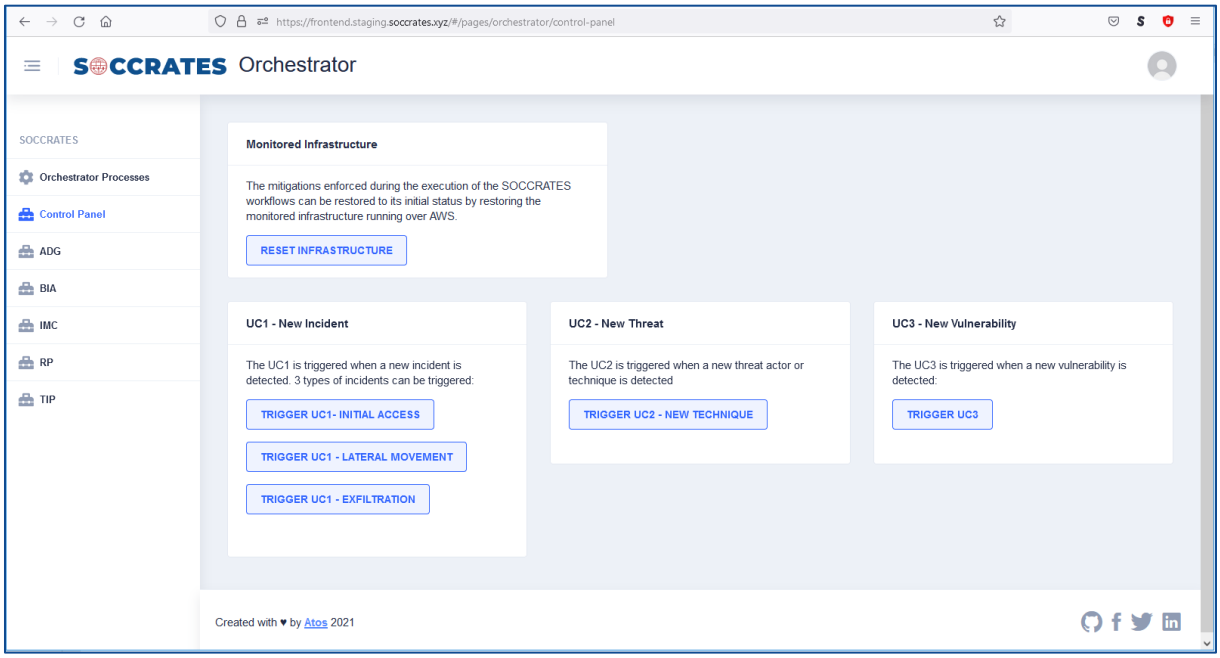

Figure 3-7: Screen shot of the SOCCRATES Front End for triggering the demonstration use cases.

4 Pilot evaluation

As explained in Deliverable 7.1 [D7.1], the second pilot stage was pivotal for obtaining an in-depth understanding of the platform's (potential) value and to measure performance on a predefined set of KPIs. The mnemonic and Vattenfall pilot sites are TRL 7-9[2] environments, however, while the ambition for the SOCCRATES platform was to achieve TRL 6[3] [H-2020]. For the project, getting theTRL 6 demonstrator environment running in full meant that the target of the original project proposal was achieved. To maximize the exposure of actual SOC analysts to the platform and its automation capabilities, however, the project sought ways to also evaluate it in higher TRL environments. To this end, a hybrid setup was established that went well beyond the original ambitions:

- At the mnenomic pilot site, the project was able to get the workflow for use case 1 (specifically the "initial access" scenario) running on actual operational infrastructure. For a limited period of time, this allowed actual analysts to conduct evaluations with the SOCCRATES platform on actual alerts that had been triaged in the mnenomic SOC. Effectively this means that selected platform capabilities were operated and trialled at TRL 7.
- The demonstrator environment (see section 3.3) was used to let analyst teams at both Vattenfall and mnemonic become acquainted with the platform in its entirety. For all of the project's use cases (1-5), therefore, both pilot teams evaluated the SOCCRATES platform in its full TLR 6 manifestation.

Table 4-1 gives an overview of which pilot sites were used for evaluation of the different components, use cases and KPIs.

The project held two dedicated workshops with the analyst teams to collect their impressions and experiences. Questions put forward to the analysts largely revolved around the project's formal KPIs (see [D7.1]). Chapter 5 provides a detailed account of the responses and how they were used to appraise KPI performance. At a higher abstraction level, evaluating the platform with actual SOC analysts also led the project to some key insights and conclusions:

1. **The platform's modelling and analysis capabilities offer SOC and CSIRT analysts a unique contextual understanding, especially for newly discovered vulnerabilities and changes in the infrastructure under control.**
   In present day setups, it is usually not trivial for SOC and CSIRT analysts to ascertain which assets in an organisation's infrastructure are affected by a newly discovered vulnerability and to what degree this vulnerability poses an actual, material threat to the organisation's business. This is obviously challenging for the multi-tenant SOC at mnemonic because it serves such a wide multitude of clients and infrastructures. The more focused internal SOC at Vattenfall, however, experiences similar issues and usually needs substantial time and effort to get a full situational picture as well. For changes in the organisation's infrastructure (use case 4/5), both teams indicate that

[2] Technology Readiness Level (TRL) 7: System prototype demonstration in operation environment
[3] TRL 6: Technology demonstrated in relevant environment

Table 4-1 Pilot evaluations.

	Pilot	First	Second			Demo	
	Pilot site	mnemonic	mnemonic	Vattenfall	Shadowserver	Demonstrator	AAD Demonstrator
SOCCRATES Components	AAD						
	ADG						
	AR						
	CoA						
	BIA						
	IMC						
	OIE						
	RP						
	TIP						
	Threat Prediction						
SOCCRATES Use Cases	UC1 Initial Access						
	UC1 Lateral Movement						
	UC1 Exfiltration						
	UC2 Vulnerability						
	UC2 Technique						
	UC3						
	UC4						
	UC5						
KPI Eval	KPI1						
	KPI2-3						
	KPI3-8						
	KPI9+						

Legend

Demonstrated

Demonstrated (partially)

the challenge is even greater because deployment of new systems or mutations in network configurations are typically not reported to the SOC and the IT asset inventories to which they do have access are rarely up to date. Any insight that the SOC does have in infrastructural changes results from its continuous vulnerability scanning effort, but this is typically indirect and incomplete. In view of all this, both analyst teams feel that the modelling and analysis capabilities incorporated in the SOCCRATES platform offer a unique contextual understanding that would greatly help them resolve issues such as the recent log4j vulnerability.

2. **Across all use cases considered, the platform's modelling and analysis capabilities will likely <u>enhance analyst decisions</u> on the most appropriate Course of Action (CoA).**
   Whenever a new threat or event is presented to them, experienced analysts will often form an idea of the underlying issue (e.g. "data exfiltration") and the most appropriate mitigation (e.g. "block selected outgoing traffic") fairly quickly. They typically rely on their initial suspicions because experience has taught them that these are usually correct. While the SOCCRATES platform

could arguably present a diagnosis or recommendation that they had not come up with themselves, the analysts doubt that this would happen often. They do, however, feel that the platform could enhance the quality of their decisions by substantiating their suspicions with material evidence. Thus, while they are not convinced that the platform's modelling and analysis capabilities would regularly surpass their own appraisal, the analysts see great value in the increased certainty that these capabilities introduce into the overall process.

3. **The platform's structured and predefined workflows offer value in themselves, not least because they force analysts to reflect on threats and events, their relevance for the organisation and the effects of actions that they choose to initiate.**
   Both analyst teams responded very positively to the SOCCRATES platform's structured and predefined workflows as enforced by the central orchestrator component. For one thing the feeling is that these workflows can accelerate the process of responding to a threat or event, although experienced analysts do not believe that it will save them considerable amounts of time (see also chapter 5). More importantly, however, both analyst teams see great value in the predetermined human decision points that each of the workflows incorporates. At such decision points, the analyst is essentially forced to take a step back and reflect on the issue at hand – what exactly does the threat or event under consideration signify, what actual harm could it do to the organization's infrastructure and business, what are the (potential) side effects of the mitigation that is about to be initiated, etc. Similar to the previous point, therefore, experienced analysts believe that platform's structured workflows will likely enhance the quality of their decisions.

4. **Establishing confidence and trust in the platform's output will not be trivial and any recommendation that it generates will need to come with a clear and visible substantiation.**
   Both analyst teams emphasized that actual adoption of the automation technologies incorporated in the SOCCRATES platform can only be successful under the premise of sufficient trust. Analysts need to be reassured that the platform did in fact perform a credible and reliable analysis before they will follow any of its recommendations. To that end, each CoA recommendation will need to be accompanied by a clear and easily accessible explanation. Such substantiation is of course available in the platform, but mostly resides in the individual functional components (ADG, Impact Analyser, Response Planner etc). Upfront, the project's vision was that analysts could (and would) access these individual components to gather whatever detailed insights they required. The analysts involved in the evaluation, however, emphasized that the core of this substantiation should be readily available in the platform's front-end whenever a recommendation is put forward or otherwise they would simply resort to doing the analysis themselves. This was an important take away from the evaluation workshops. Moving forward, a key challenge will be to design a user interface that offers sufficient substantiation (and thus confidence) whilst at the same time avoiding an unmanageable overload of information.

To place the above in a proper perspective, the following remarks are of importance:

- The evaluation workshops were conducted with fairly experienced SOC and CSIRT analysts. Some of their younger, less experienced team members would likely get even more value from the platform's capabilities, but this was not assessed in detail.
- The evaluation workshops did not cover the second phase pilot at the Shadowserver site (see section 3.2.3). This pilot was, however, executed in full and one of the project's major successes. Its results are reflected in the KPI evaluation (specifically KPIs 9-13) in the following chapter.

Finally we note that the project actually aspired to do even more elaborate trials at the operational pilot sites (i.e. at TRL 7 or above). To this end the team prepared an extensive set of metrics and measurement instruments. While these could not be employed within the boundaries of this work, they have been documented in annex A for future reference.

5 Evaluation of KPIs

The plan described in D7.1 [D7.1] was to evaluate KPIs 1-8 at the Vattenfall and mnemonic pilot sites. Due to the pilot site challenges described in previous sections, this plan had to be amended for KPIs 1-8. For each of these KPIs below, we describe the assessment methodology and the results.

5.1 KPI1 - Prototype successfully generates a correct model of the pilot infrastructure (all UCs)

KPI1 was assessed at the mnemonic pilot site, following the measurement methodology from D7.1. The assessment was carried out by Martin Eian and Fredrik Borg from mnemonic, with help from the mnemonic Infrastructure team to verify our findings.

The IMC at the mnemonic site was populated with data from two site specific adapters. The first adapter provided data about hosts, software, networks and vulnerabilities from the asset inventory in mnemonic's Argus platform. The origin of this data is primarily vulnerability scans. The second adapter provided the firewall rule sets from mnemonic's firewalls.

Our main focus was on the asset inventory, since we suspected that there might be incorrectnesses in that data set. We also did a random sampling of the firewall rule set in the IMC and compared this to the real infrastructure.

The IMC correctly represented the firewall rule set. For the asset inventory, we found several issues with the IMC data when compared to the real infrastructure:

- The Common Platform Enumeration (CPE) for software objects was too generic, it identified the type of software but not the version
- The IMC data model does not support Hot Standby Router Protocol (HSRP) addresses, this resulted in duplicate IP addresses for one host in the IMC
- Several hosts had an unknown host name
- Some hosts were identified by NetBIOS name rather than the fully qualified domain name
- Missing software objects (e.g. no sshd software associated with a login server running sshd)

After a thorough investigation, we concluded that these incorrectnesses were also present in the original data source that the IMC adapter fetched the data from. With the exception of the HSRP issue, we found no instances where the IMC itself or the adapters caused incorrectnesses. Furthermore, the IMC adapter was able to correctly associate vulnerabilities with hosts even if the software object was missing. This was solved using "Fact Chains", as described in the ACT Python API documentation[4].

Overall, the IMC provided a correct model of hosts, vulnerabilities, network topology and firewall rules. In order to improve the infrastructure model, we need to improve the data quality of the asset inventory. mnemonic is already running an internal project to address this. The IMC data model should be updated to be able to handle HSRP addresses and other load balancing solutions. Currently, such addresses may cause an incorrect infrastructure model.

[4] https://github.com/mnemonic-no/act-api-python

5.2 KPI2 - Prototype successfully improves detection capability (incl. detecting attacks in encrypted traffic that was not detected otherwise) (UC1)

The original measurement methodology should compare the output of the AAD component to a baseline of output of individual anomaly detection tools. Our evaluation has sought to go beyond this approach by examining the benefits of enriching the outputs from ADSs with a knowledge-based system – specifically an Evidential Network (EN) – in order to account for uncertainties that stem from detector performance and opportunities that arise from combining events from several detectors.

To evaluate the AAD, a demonstrator tool has been implemented that can be used to simulate several attack scenarios. These scenarios – summarized in Table 5-1 – describe situations in which different SOCCRATES anomaly detection systems (ADSs) generate events. A description of the AAD Demonstrator tool can be found in SOCCRATES Deliverable D6.4. In short, the tool can be used to input scenarios to an Evidential Network (EN) that is configured to provide inferences, including a belief score, regarding which adversary tactic is being described by events that are generated by detectors. The EN that was produced for this purpose is depicted in Figure 5-1.

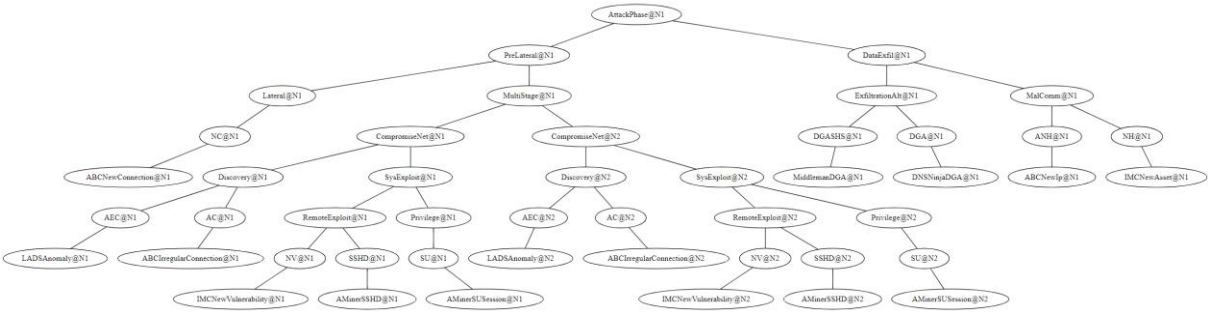

Figure 5-1 An overview of the Evidential Network that was developed for the SOCCRATES demonstration

It is important to note that the EN takes into account the performance of the detectors that are used in the SOCCRATES project. More specifically, False Positive (FP) and False Negative (FN) rates are used to calculate a Positive Predictive Value (PPV) using a method previously published by Frieberg et al.[5]. The FN and FP rates for the detectors are taken from SOCCRATES Deliverable D4.1. The PPV is intended to articulate the confidence one can have in the correctness of an alert from a detector – e.g. the FP rate of the detector is used, via the PPV, to derive a numerical confidence that an alert generated by a detector does represent a real anomaly.

The scenarios that are summarized in Table 5-1 describe situations in which SOCCRATES ADSs are generating events and the adversarial tactic they could be describing. The reason that the EN is providing inferences regarding the observed adversarial tactic is to support the selection of a suitable playbook or CoA, as part of a response to a detected attack (UC1). For each of the adversary tactics, different scenarios have been generated that provide more or less evidence of the tactic being achieved by an adversary. For example, the *Data Exfiltration – Lightest* scenario is using events only from DNS Ninja (which detects anomalous DNS behaviour) to infer that data exfiltration has been achieved by an adversary; meanwhile, the *Data Exfiltration – Strongest* scenario uses evidence from multiple detectors to infer the same objective. The intention is that as more evidence is provided by

[5] I. Friedberg, X. Hong, K. Mclaughlin, P. Smith and P. C. Miller, "Evidential Network Modeling for Cyber-Physical System State Inference," in IEEE Access, vol. 5, pp. 17149-17164, 2017, doi: 10.1109/ACCESS.2017.2718498.

Classification level: Public

the SOCCRATES ADSs, the greater the belief score is generated by the EN and, correspondingly, an increased confidence an analyst can have regarding the inferences.

Table 5-1 An overview of the scenarios that have been used to evaluate the AAD and KPI2

Scenario	Evidence (Events)	Rationale
Data Exfiltration		
Lightest	DNSNinjaDGA@N1	Unusual DNS behaviour is detected by DNS Ninja that could indicate data is being exfiltrated using the DNS protocol. No other evidence is present.
Light	MiddlemanDGA@N1; DNSNinjaDGA@N1	Unusual DNS behaviour could indicate an adversary is attempting to exfiltrate data.
Strong	*Light*; ABCNewIP@N1; IMCNewAsset@N1; AMinerSUSession@N1; AMinerSSHD@N1; IMCNewVulnerability@N1; ABCIrregularConnection@N1; LADSAnomaly@N1; ABCNewConnection@N1	In addition to unusual DNS behaviour, other anomalies provide supporting evidence regarding the presence of an adversary.
Strongest	*Light;* ABCNewIP@N1; IMCNewAsset@N1	Both unusual communication to an external host, indicated by the ABC Tool and the IMC providing evidence of a new host (asset), and unusual DNS behaviour provide the strongest evidence of potential data exfiltration.
Initial Access		
Light	IMCNewVulnerability@N2; AMinerSSHD@N2; AMinerSUSession@N2	Evidence of unusual behaviour that could indicate an attempt to compromise a host is found on one network – in this case Network Two. Additionally, a known vulnerability could provide supporting evidence of compromise. and vice versa.
Strong	*Light;* LADSAnomaly@N2; ABCIrregularConnection@N2	In addition to evidence of a host being compromised (see Initial Access – Light), evidence of discovery activities is present, which collectively increases confidence that a host has been compromised.
Lateral Movement		
Lightest	ABCNewConnection@N1	An anomaly is detected that could indicate the presence of communication between hosts that has not previously interacted; this could be an indicator of an attempt at lateral movement.
Light	LADSAnomaly@N1; ABCIrregularConnection@N1; LADSAnomaly@N2; ABCIrregularConnec-	Strong indications that a host has been compromised on one network – in this case N2 – and indicators of an adversary conducting discovery activities on a second

Scenario	Evidence (Events)	Rationale
	tion@N2; IMCNewVulnerability@N2; AMinerSSHD@N2; AMinerSUSession@N2	network (N1); however, there is no evidence of a second host being compromised.
Strong	*Light*; IMCNewVulnerability@N1; AMinerSSHD@N1; AMinerSUSession@N1	The same evidence for Lateral Movement - Light but additional evidence that indicates that a host in a second network (N1) has been compromised, as evidenced by anomalies being generated by AMiner and the IMC indicating the presence of a vulnerability.
Strongest	*Strong*; ABCNewConnection@N1	The same evidence as for Lateral Movement – Strong but additional indication from the ABC Tool providing evidence of communication between two hosts that normally do not communicate.
Exceptions		
Normal	No detectors generating events	No events have been generated, indicating that no adversarial behaviour is present.
All Detectors	All detectors generating events	All the detectors are generating events, indicating anomalies. It is assumed the adversary has achieved the most advanced tactic, i.e. they have managed to perform initial access, lateral movement, leading to the exfiltration of data.

One of challenges that the AAD is intended to address are FP-related alerts that are generated by ADSs. False positives have the potential to waste SOC analyst's time and reduce confidence in the detector's ability to correctly identify anomalies and, therefore, attacks. To address this issue in SOCCRATES, the PPV of a detector is calculated based on its detection performance – the FP and FN rate. The PPV is used by the EN when deriving the belief score regarding a hypothesis about what adversary tactic has been achieved. The idea is that unreliable detectors that generate larger quantities of FPs are trusted less by the EN when making inferences about their root cause (adversary tactic). An analyst can use the belief score to make a judgment regarding whether to take action – a relatively low belief in a hypothesis may attract no response; whereas a higher belief may be used as a basis to initiate a response activity (e.g. initiate a playbook associated with data exfiltration).

The results that are presented in Table 5-2 aim to illustrate the effect of trustworthiness in a detector and how the EN uses it to determine a belief score in an inference. In this example, we consider the detection performance of the DNS Ninja ADS, which detects abnormal use of the DNS protocol. Based on the detection performance that is described in SOCCRATES Deliverable D4.1, the PPV for DNS Ninja is 0.95 – in other words, it has a relatively low FP rate. Using the EN presented in Figure 5-1, when DNS Ninja generates an alert, a hypothesis is generated that a data exfiltration attack has occurred with a belief score of 0.513. This score takes into account both the detectors performance and expert knowledge that captures whether unusual DNS behaviour is representative of data exfil-

tration – it could also indicate C2 communications, for example. In practice, more evidence is required to make this assertion. However, what is noteworthy is that as the PPV of DNS Ninja decreases (i.e. the detector is deemed less reliable, because it generates more FPs), the belief in the overall hypothesis decreases, also. This arrangement could be used in several ways by a SOC analyst: one simple approach would be to assign an arbitrary threshold, based on empirical analysis or analyst experience, wherein any belief in a hypothesis from a (set of) detectors that is below the threshold is not actioned by the SOC. In this case, for example, if an unreliable version of DNS Ninja (with a PPV of 0.7 or lower) was generating alerts that resulted in a hypothesis below a threshold of 0.5, those inferences from the AAD would be disregarded. Meanwhile, inferences with a higher belief score could be subject to further investigation and execution of a corresponding playbook.

Table 5-2 The effect of DNS Ninja's detection performance and corresponding PPV on the belief scores regarding the presence of data exfiltration

PPV of DNS Ninja	Data Exfiltration Belief
0.5	0.4026
0.7	0.4578
0.9	0.513
0.95	0.5267
0.99	0.5378

Another way to improve detection performance is to use multiple sources of evidence from different detectors. This approach can additionally have the benefit of detecting adversarial behaviour that uses encrypted traffic, e.g. by using evidence that is provided by Host-based Intrusion Detection Systems (HIDSs) and flow-based Network-based Intrusion Detection Systems (NIDSs). The SOCCRATES AAD makes use of a selection of these technologies to a) identify malicious behaviour that is using encrypted communication; and b) improve detection performance.

One way to exfiltrate data and conduct C2 communications is using encrypted protocols, such as HTTPS[6]. Whilst it is not possible to inspect the details of this communication, because it is encrypted, other network flow-based characteristics could be used to identify potentially malicious behaviour[7]. In the SOCCRATES project, the ABC Tool and L-ADS can perform flow-based anomaly detection. This ability has been demonstrated for data exfiltration over DNS traffic and can be used to provide supporting evidence of this activity. Of course, data that is exfiltrated over DNS is not encrypted but here we use it as a stand-in for exfiltration over HTTPS, for example. Consider a scenario in which neither DNS Ninja nor The Middleman have been able to detect unusual behaviour associated with DNS traffic, because it is "encrypted". Using the EN that is shown in Figure 5-1, evidence from the ABC Tool and the IMC could be used to infer that exfiltration is underway using alternative data sources – namely, flow data and asset information. In this situation, when the ABC Tool and IMC indicate anomalies, the EN is able to make the inferences that are shown in Table 5-3. Because of the configuration of the EN, the alternative evidence that is provided by the ABC Tool and the IMC, it can correctly infer the presence of data exfiltration. This example is somewhat contrived and artificial, but it

[6] See the MITRE ATT&CK technique for details: https://attack.mitre.org/techniques/T1048/002/

[7] Rui Dai, Chuan Gao, Bo Lang, Lixia Yang, Hongyu Liu, and Shaojie Chen. 2020. SSL Malicious Traffic Detection Based On Multi-view Features. In Proceedings of the 2019 the 9th International Conference on Communication and Network Security (ICCNS 2019).

serves to highlight the capability of the AAD, using a combination of detectors, to identify malicious behaviour that uses encrypted traffic.

Table 5-3 The inferences (belief scores) made by the EN when there is an absence of unusual DNS behaviour but abnormal flow behaviour – detected by ABC Tool – and the IMC has identified the presence of a new IP address.

Initial Access	Lateral Movement	Data Exfiltration	Normal
0.1424	0.1135	**0.6049**	0.1391

Another way to improve detection performance of ADSs is to combine events from different detectors. The results shown in Table 5-4 demonstrate how the AAD can be used to provide inferences about adversarial tactics with higher confidence, given that more detectors generate events. Each row in the table represents one of the scenarios that are summarized in Table 5-2, wherein there are scenarios in which fewer detectors provide support for a hypothesis (those labelled as "light") and others that provide more evidence (those labelled as "strong"). It can be seen that as more detectors generate an event, the belief in a hypothesis increases. For example, for the lightest data exfiltration scenario, in which DNS Ninja is generating an event, the EN generates a belief score of 0.513; meanwhile, when other detectors generate events, the belief can reach 0.8556 for the strongest case. This ability is demonstrated very well for the lateral movement tactic, in which a single detector results in an incorrect hypothesis (initial access) but, as more detectors that describe lateral movement generate events, the EN is able to correctly identify the adversary tactic with increased confidence (belief score).

Table 5-4 The belief scores for the scenarios that are described in Table 5-2; as more evidence is provided the belief score correspondingly increases. Green cells indicate a correct hypothesis, whereas a red cell indicates otherwise.

Scenario	Initial Access	Lateral Movement	Data Exfiltration	Normal
Data Exfiltration				
Lightest	0.1756	0.138	**0.513**	0.1734
Light	0.1463	0.1159	**0.5949**	0.1429
Strong	0.1024	0.1134	**0.7596**	0.0246
Strongest	0.0518	0.0472	**0.8556**	0.0454
Initial Access				
Light	**0.4032**	0.264	0.1914	0.1413
Strong	**0.4654**	0.2881	0.163	0.0835
Lateral Movement				
Lightest	**0.3657**	0.2881	0.1992	0.147
Light	0.2743	**0.4833**	0.1818	0.0606
Strong	0.223	**0.5371**	0.1864	0.0535
Strongest	0.1798	**0.5873**	0.1895	0.0434
Exceptions				
Normal	0.2695	0.205	0.2555	**0.2699**
All Detectors	0.0537	0.1601	**0.7656**	0.0206

To conclude, we argue that the AAD is able to improve detection capabilities in several ways: i) it is able to account for unreliable detectors with high FP and FN rates by incorporating this information into the beliefs it generates about the presence of adversarial tactics – this should assist analysts

when prioritizing alerts from different detectors; ii) the AAD has the potential to detect malicious behaviour that uses encrypted traffic by using information from multiple data sources to infer the presence of adversarial behaviour; and iii) the AAD is able to capitalize on multiple detectors to increase confidence in the hypotheses it generates about the nature of the adversary tactic that has been achieved by an attacker.

5.3 KPI3 - Prototype successfully reduces total time of detection and response (UC1)

5.3.1 Time to detection

We propose that the belief scores that are generated by the AAD (see KPI 2) can be used in combination with the inference that garners the highest score to determine a relative Time To Detection (TTD). In general, the EN provides a higher belief score when more reliable evidence (events) is provided by detectors. Arguably, a higher belief score from the EN should enable analysts to more rapidly confirm the achievement of a given adversarial tactic by an attacker. Conversely, lower belief scores require more intervention from an analyst to determine whether a tactic (e.g. data exfiltration) has been achieved by an attacker. These interventions include determining whether events generated by an AD are false positives, e.g. by inspecting the network flow records (or full packet captures) that generated an event. This can require not insignificant time from an analyst and delays the establishment of situation awareness – with respect to the adversary tactic. Therefore, we propose to associate the TTD with the belief scores that are associated with inferences, which are generated by the AAD. As indicated in Table 5-2, this promotes faster detection when an ADS is deemed more reliable because of its detection performance, resulting in higher belief scores from more reliable detectors.

As discussed earlier, we associate the TTD with the highest belief score for an inference that is generated by the AAD. In this context, we suggest the TTD is the time to verify by an analyst the achievement of an adversary tactic (goal) by an attacker (i.e. initial access, lateral movement, or data exfiltration) in order to engage a tactic-specific CoA. Moreover, we argue the two values (TTD and belief) are inversely correlated – the greater the belief, the faster the TTD. Intuitively, this proposal is supported by our analysis, because more evidence is required from reliable detectors in order to generate inferences with a higher belief score (see Table 5-2 and Table 5-4 that show this relationship)

For the purpose of illustration, we can define arbitrary thresholds, wherein we define relative TTDs that are related to ranges of belief scores from the AAD for the 'Lateral Movement' tactic[8], as shown in Table 5-5. We arbitrarily define three TTDs – i) *slower*, which indicates that significant analyst time and effort (e.g. performing extensive packet and log file analysis, based on FPs generated by detectors) is required to determine the presence of an adversary achieving a given tactic; ii) *fast*, wherein the evidence that is provided by the AAD from the detectors is used with some confidence to more rapidly determine (or confirm) an adversary tactic has been achieved; and ii) *faster*, wherein significant evidence is provided by the AAD along with a strong belief in the achievement of an adversary tactic – the intention is that an analyst can have significant confidence in an assertion about the situation and has sufficient evidence to rapidly confirm it, provided by the AAD.

[8] In this case, we have chosen to use the Lateral Movement tactic, as it provides a useful illustrative example, given the results of our experiments. A similar set of thresholds could be defined for the other tactics.

Table 5-5 Ranges of belief score and the relative Time to Detection (TTD) of an the 'Lateral Movement' tactic

Belief Score Range	Qualitative TTD
>= 0 and < 4	Slower
>= 4 and < 5	Fast
>= 5	Faster

Given these thresholds, the relative TTDs for the scenarios that are defined in Table 5-4 for the 'Lateral Movement' tactic are shown in Table 5-6. For the scenario in which the "lightest" evidence is provided, specifically just one detector indicates the presence of lateral movement (see Table 5-1), the relative TTD is deemed to be *slower.* In other words, an analyst will have to conduct more investigation to confirm this inference by the AAD or, alternatively, dismiss the inference as not credible for further investigation – a potentially legitimate action. Thereafter, as better evidence of the presence of lateral movement being achieved by the attacker is provided by the AAD, resulting in improved belief scores, the relative TTD decreases (becomes faster).

Table 5-6 An example relationship between the belief scores generated by the AAD and a relative TTD; the relationship between these values is defined in Table 5-5

Scenario	Belief Score	Relative TTD
Lateral Movement		
Lightest	0.3657	Slower
Light	0.4833	Fast
Strong	0.5371	Faster
Strongest	0.5873	Faster

Interestingly, for the 'Data Exfiltration' inferences, the 'Strong' scenario incorporates more events by the ADSs than for the 'Strongest' scenario (see Table 5-1). This is because experts have deemed that events that are generated by the detectors in the 'Strongest' scenario – the ABC Tool and IMC – are more reliable indicators of the presence of data exfiltration than the more numerous events that are generated by the other detectors.

In this analysis, we have argued there is a relationship between the inferences that are generated by the AAD, and associated belief scores, and the TTD required by analyst. The inferences that are generated by the AAD can be used to more rapidly affirm a hypothesis regarding the existence of a tactic being achieved by an adversary or, in order to reduce analysts' workloads, dismiss weak indicators until sufficient evidence is present in order to warrant a more concerted response from the SOC.

5.3.2 Time to response

The analysts from mnemonic and Vattenfall were shown a demo of Use Case 1, where the SOCCRATES platform was used to respond to a Lateral Movement scenario. The analysts were shown the OIE interface, the ADG Analyser interface and the BIA interface, and how these would be used during analysis. The analysts were then asked to respond to a series of statements. Figure 5-2 and Figure 5-3 show responses to a statement.

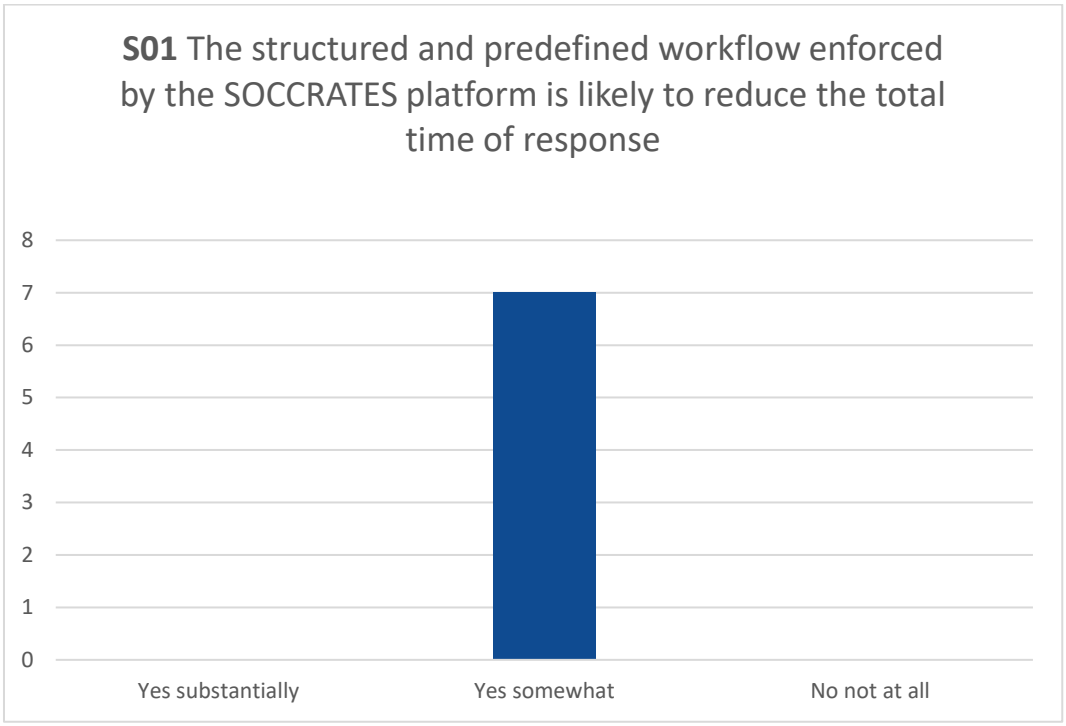

Figure 5-2 Response to statement S01

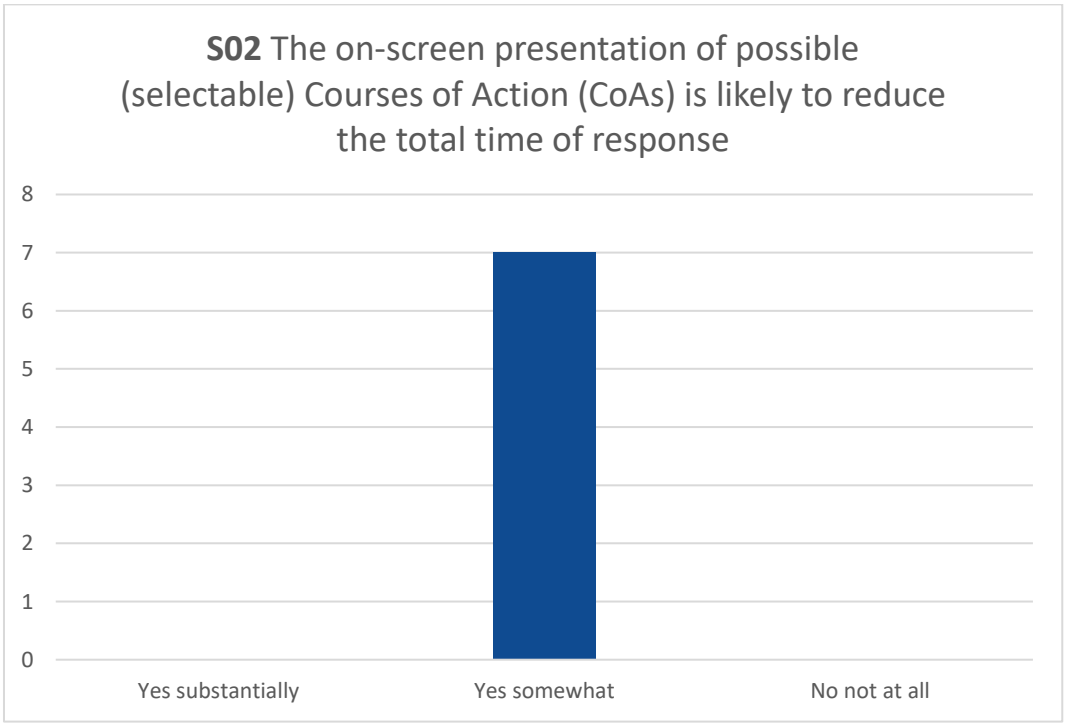

Figure 5-3 Response to statesment S02

All of the analysts responded “Yes somewhat”. Their assessment was that the SOCCRATES platform would help reduce the time to response, but not substantially.
Even though KPI3 is strictly about time to response, we also asked the analysts to respond to a statement about the quality of response. The response to the statement is shown in Figure 5-4.

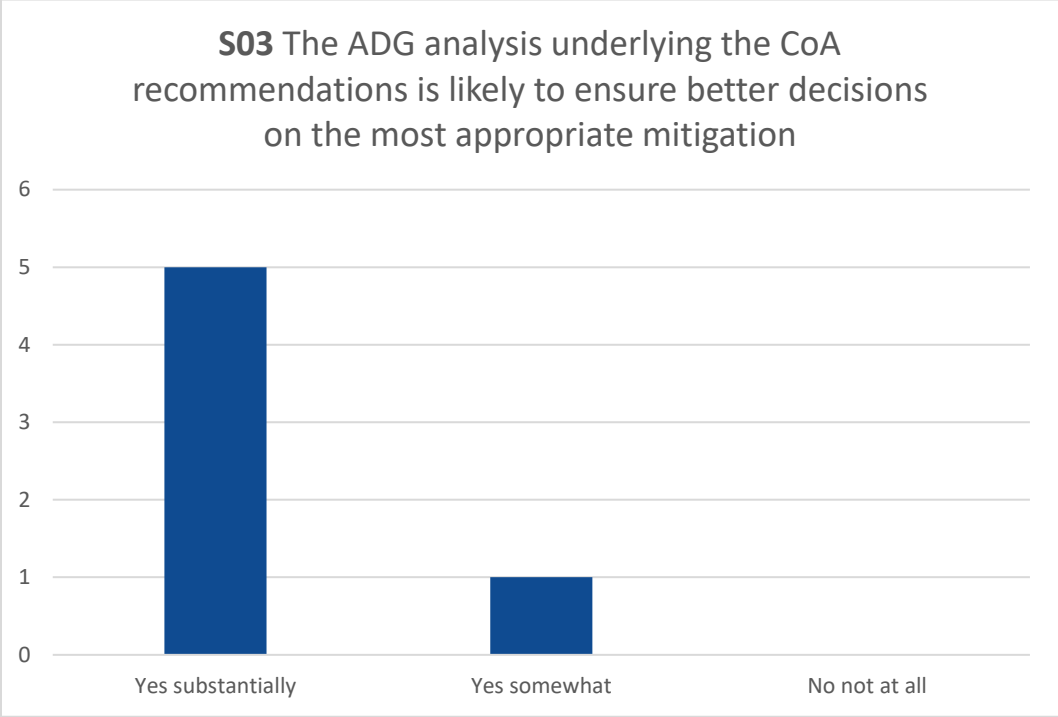

Figure 5-4 Response to statement S03

This result, along with comments from the analysts, show that their assessment is that the main contribution of SOCCRATES is to make better decisions, not to make faster decisions.

5.4 KPI4 - Prototype successfully reduces total time from collection, analysis and determining pro-active mitigation new CTI (UC2)

The analysts from mnemonic and Vattenfall were shown a demo of Use Case 2, where the SOCCRATES platform was used to respond to new CTI: The threat actor APT32 is now known to use the technique T1558 “Steal or Forge Kerberos Tickets”. The analysts were shown the OIE interface, the ADG Analyser interface and the BIA interface, and how these would be used during analysis. The analysts were then asked to respond to a series of statements, see Figure 5-5 and Figure 5-6.

Classification level: Public

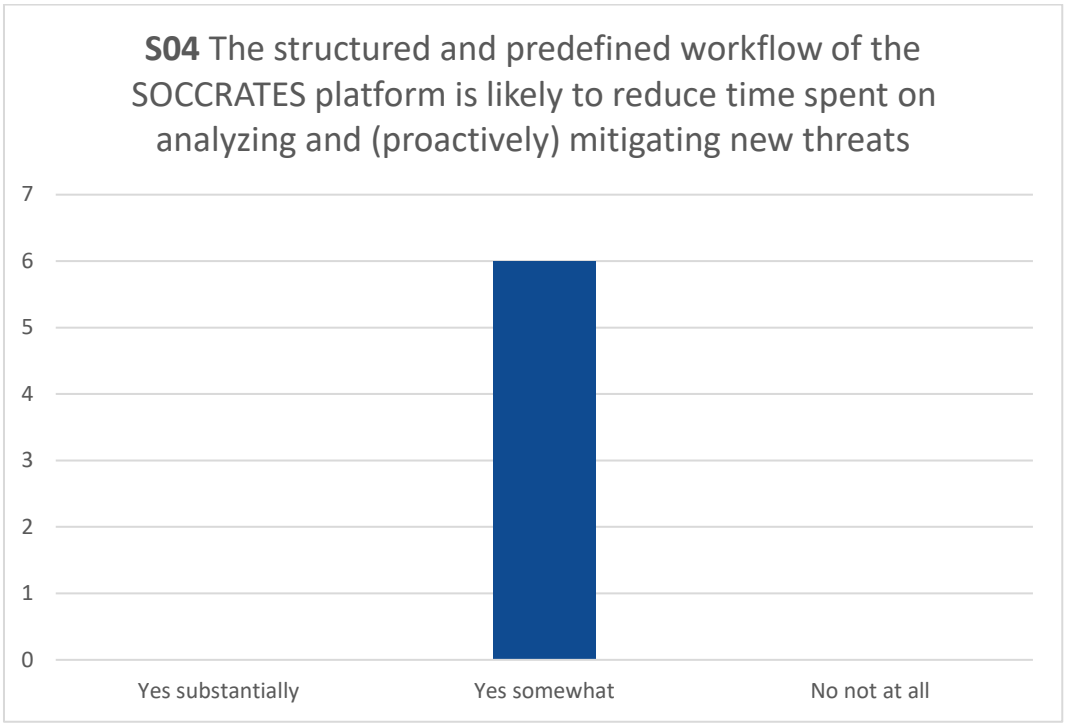

Figure 5-5 Response to statement S04

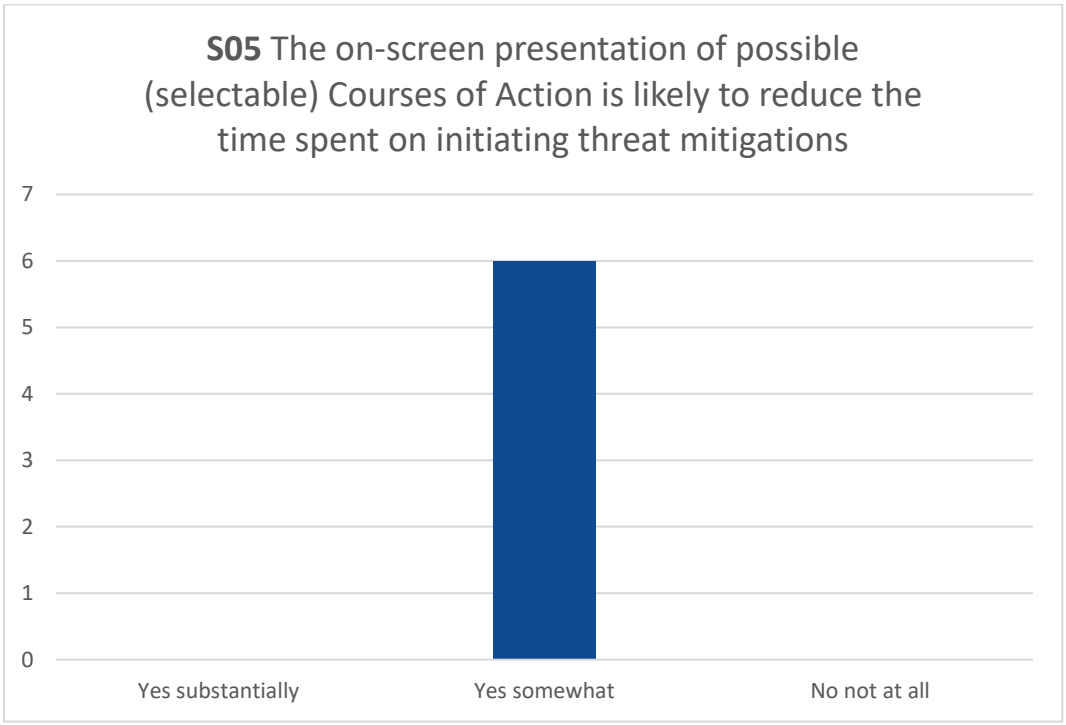

Figure 5-6 Response to statement S05

As for the previous KPI, the response from the analysts was an unanimous "Yes somewhat". Their assessment was that the SOCCRATES platform would help reduce the time to response, but not substantially.

Once again, we also asked the analysts to respond to a statement about the quality of response, see Figure 5-7.

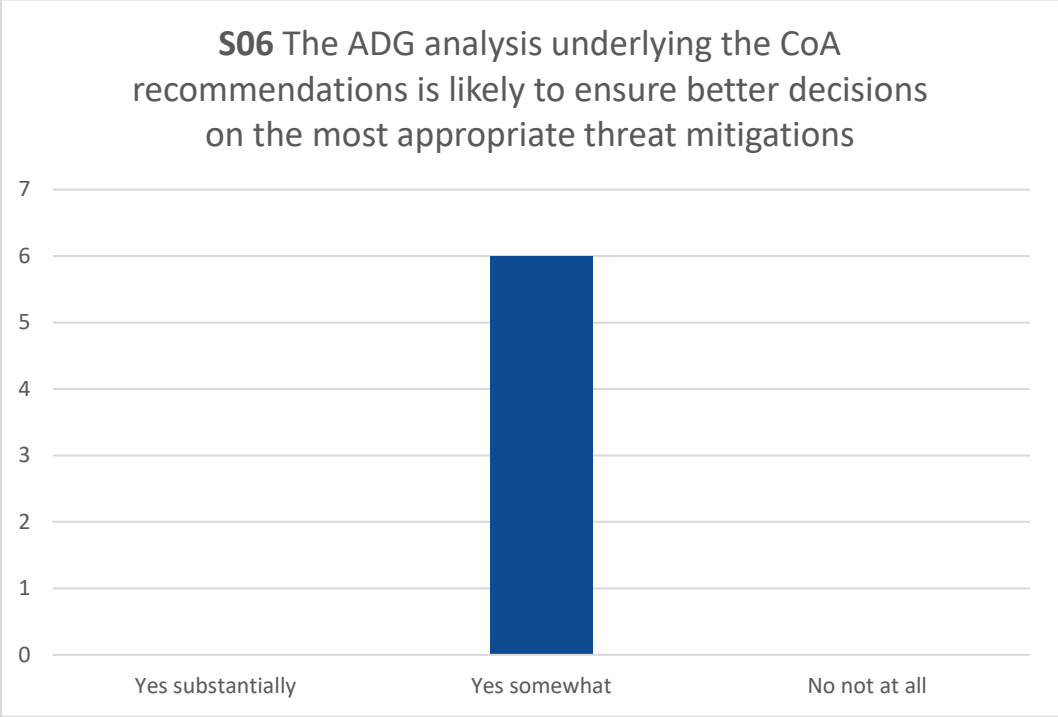

Figure 5-7 Response to statement S06

Compared to KPI3, the scores for this statement are substantially lower. This indicates that the analysts see more value in the SOCCRATES platform's support for incident response than its support for threat mitigation.

5.5 KPI5 - Prototype successfully analyses the vulnerability of an asset in the context of the whole infrastructure, that could cause an incident otherwise (UC3)

The analysts from mnemonic and Vattenfall were shown a demo of Use Case 3, where the SOCCRATES platform was used to respond to a new vulnerable asset: The web portal in the demonstrator environment is now vulnerable to CVE-2022-1292. The analysts were shown the OIE interface, the ADG Analyser interface and the BIA interface, and how these would be used during analysis. The analysts were then asked to respond to a series of statements. The response to the statements are shown in Figure 5-8 and Figure 5-9.

Classification level: Public

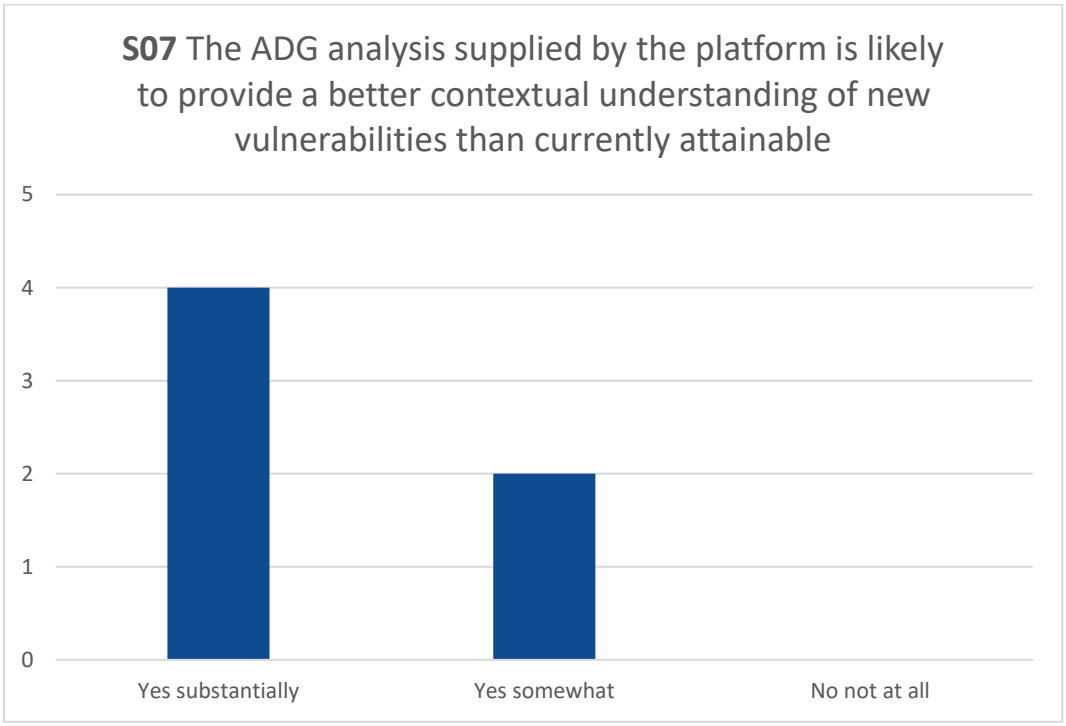

Figure 5-8 Response to statement S07

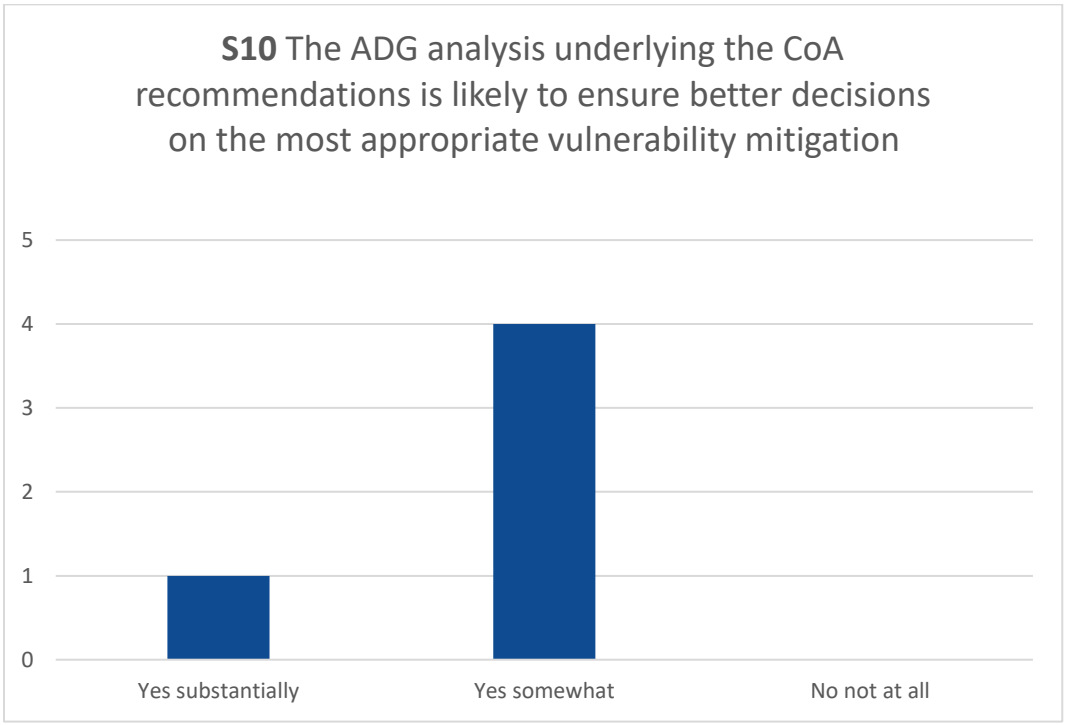

Figure 5-9 Response to statement S10

The responses from the analysts indicate that the SOCCRATES platform provides substantially better contextual understanding of new vulnerabilities, but that they are not fully convinced that this improved understanding will substantially improve their decisions.

5.6 KPI6 - Prototype reduces the time spent on analyses of the severity of vulnerable assets (UC3)

The analysts from mnemonic and Vattenfall were shown a demo of Use Case 3, where the SOCCRATES platform was used to respond to a new vulnerable asset: The web portal in the demonstrator environment is now vulnerable to CVE-2022-1292. The analysts were shown the OIE interface, the ADG Analyser interface and the BIA interface, and how these would be used during analysis. The analysts were then asked to respond to a series of statements, see Figure 5-10 and Figure 5-11.

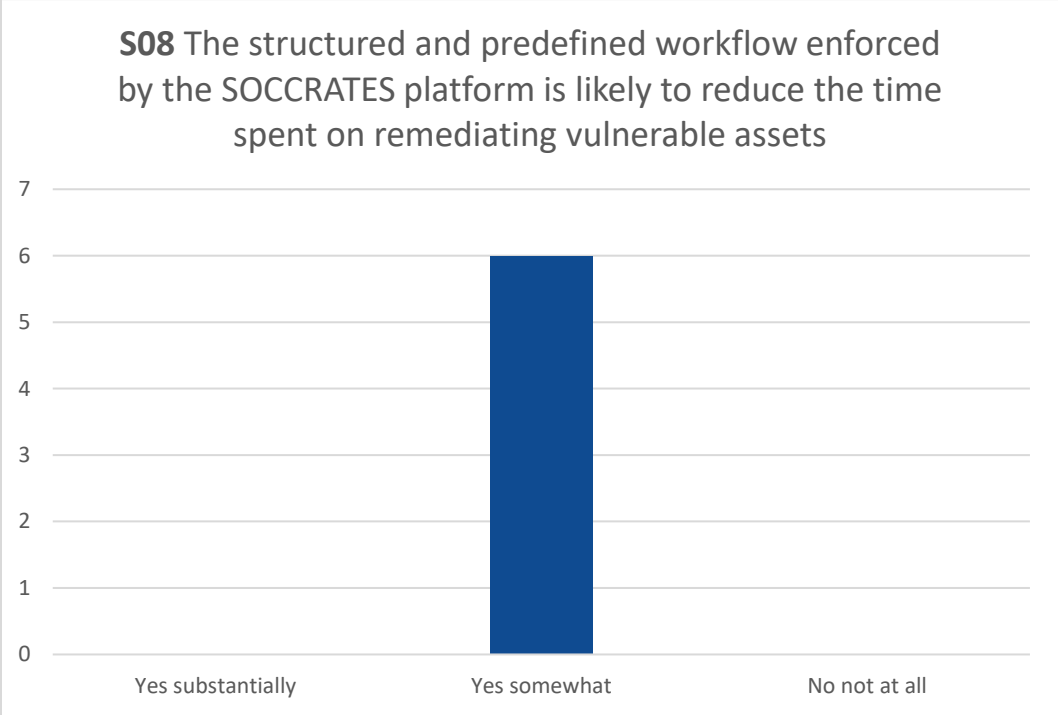

Figure 5-10 Response to statement S08

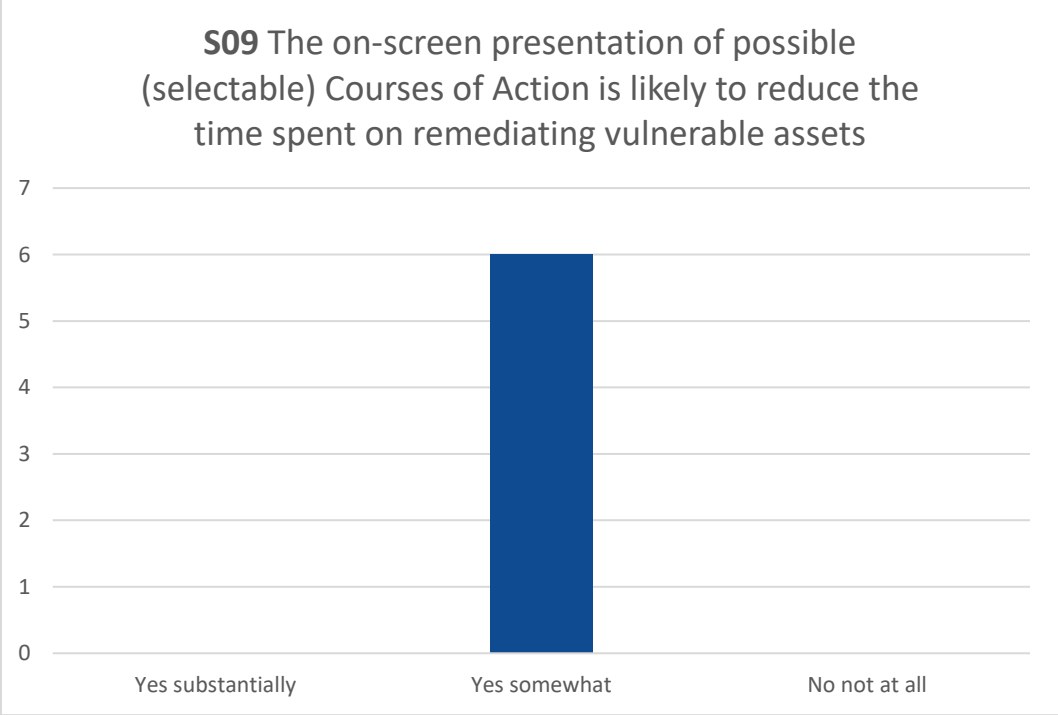

Figure 5-11 Response to statement S09

Once again, the assessment from the analysts was that the SOCCRATES platform would reduce the time spent on remediation somewhat, but not substantially.

5.7 KPI7 - Prototype proposes at least one course of action that is more efficient/appropriate than what would otherwise have been proposed by a human analyst using templates (UC4 & UC5)

Use Case 5 was not directly demonstrated to the analysts due to time constraints. The analysts from mnemonic and Vattenfall were told that the workflow for Use Case 4 is the same as for Use Case 3, but in this case the infrastructure change was a change in the firewall configuration. The analysis performed and shown by the OIE, ADG and BIA are equivalent to the analysis shown during Use Case 3. The analysts were then asked to respond to a statement, with the results shown in the figure below:

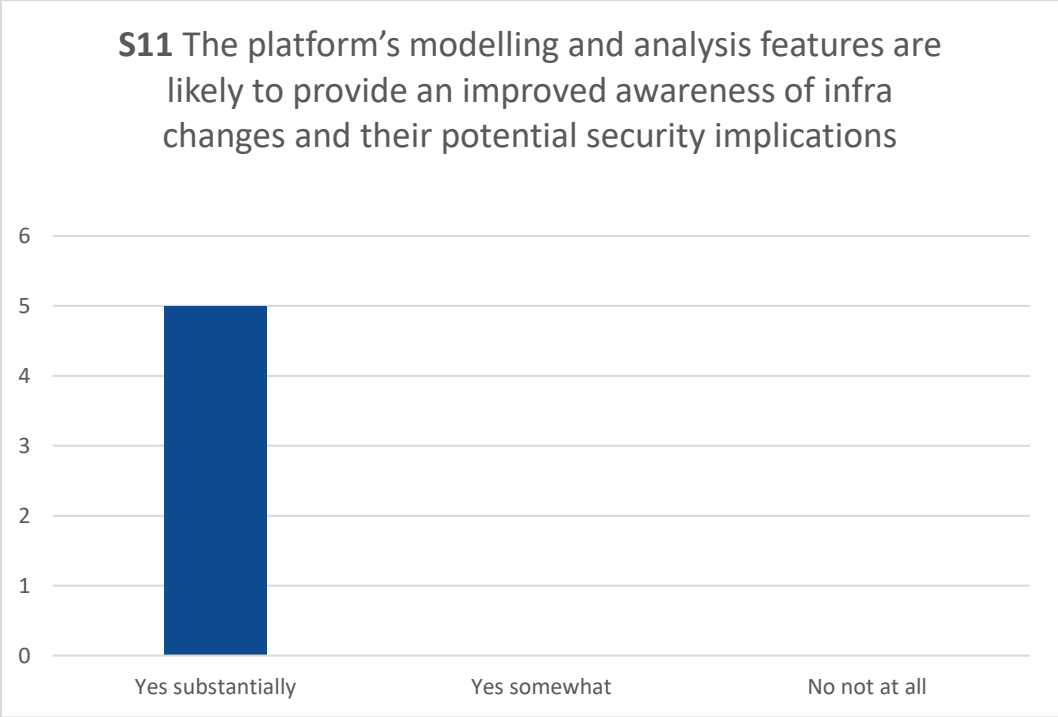

Figure 5-12 Response to statement S11

The analysts unanimously responded that the SOCCRATES platform would substantially improve the awareness of infrastrucrute changes and their potential security implications.

5.8 KPI8 - Prototype reduces the time spent on analysis and determining a course of action (UC4 & UC5)

We had prepared statements for assessing the time spent on Use Cases 4 and 5, but due to time constraints we were not able to get responses from the analysts for these statements. Therefore, we were unable to evaluate KPI8. Given the results for all the other time-related KPIs it is unlikely that the assessment would be that the time spent on analysis would be substantially reduced.

5.9 KPI9 - Prototype successfully detects 25 new DGA seed from known malware family

As written in D7.1, this KPI was heavily threat landscape dependent. We were unable to observe a known malware family employing 25 new seeds during the Shadowserver pilot evaluation. In a few cases, we were able to observe known families with unlabelled seeds, such as for Mydoom or Phorpiex.

The effectiveness of the prototype algorithms will be therefore measured on historical data collected by Shadowserver or demonstrated that having some of the data from a few seeds is enough to discover other seeds that were not already known by the classifier. To demonstrate this, we selected the TinyBanker (Tinba) banking malware, for which until now 144 seeds are known. The binary classifier used in the DGA detective has been trained on Tinba domain names produced with five of these seeds (tinba_dga_70146851, tinba_dga_a4932e33, tinba_dga_5f7a4794, tinba_dga_3ffd4605, 'tinba_dga_24c56758 as described in the DGArchive). We evaluated the prototype on all the domain names from the archive which are generated by the remaining 139 seeds. which have not been seen during the training phase. Figure 5-13 shows the distribution of accuracy scores. On average, the model achieves 99% accuracy, and there exists 22 seeds whose associated domain names are perfectly classified as DGAs. This result shows the generalization capability of the model, and demonstrated the achievement of the KPI.

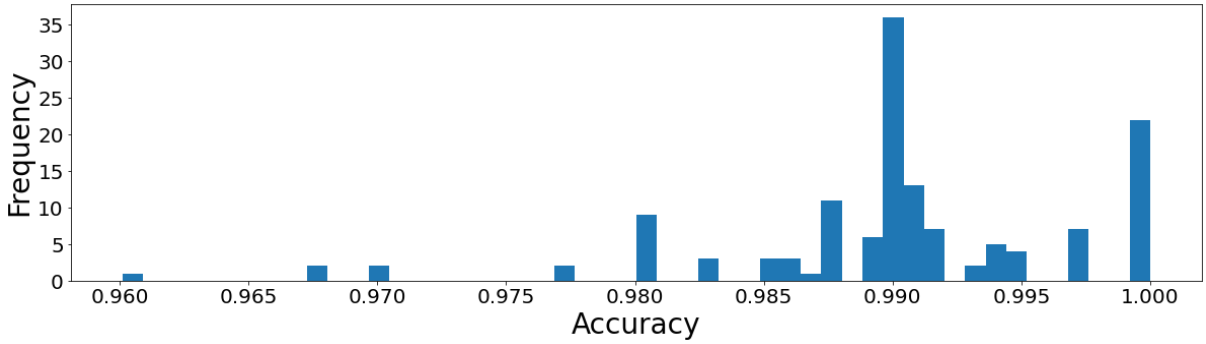

Figure 5-13 Distribution of accuracy scores.

5.10 KPI10 - Prototype successfully detects 5 new, previously unknown malware DGA

The DGA Detective tool running in the Shadowserver backend as part of the Shadowserver pilot successfully detected a number of previously unknown (not in the FKIE DGArchive at time of detection) new DGA algorithms. This includes a Phorpiex DGA, a M0yv DGA (both added to the DGArchive) and potential new DGAs for Softpulse, Razy, Nitol and Neksminer. A number of other DGA-like candidates were also discovered. This KPI was therefore reached based on in-the-wild malware spread observations.

The effectiveness of the prototype algorithms can also be measured on historical data collected by Shadowserver or evaluated versus DGAs that are not known to the classifier (not in the training dataset). To illustrate this, we evaluated the binary classifier included in the DGA Detective tool on the

following malware families (none of which were represented in the training data): FluBot, Sunburst, Phorpiex, M0yv, and Sharkbot. The results are shown in Table 5-7 The prototype results. The prototype is able to identify domains produced by all these families, which demonstrates the achievement of the KPI.

Table 5-7 The prototype results.

Family	Accuracy	Support (# samples evaluated)
Flubot	0.993	10000
Sunburst	1.000	5000
Phorpiex	0.846	1431
M0yv	0.630	888
Sharkbot	1.000	100

5.11 KPI11 - Prototype successfully clusters 30 malware families

This KPI refers to the ability of the prototype to identify at least 30 different malware families, a task initially intended to be performed with the aid of clustering methods. Despite some efforts of applying string grouping to domain names, presented in Deliverable 4.2 [D4.2], it was found that such methods were not delivering sufficiently accurate results, and so classification methods were studied instead. As described in Deliverable 4.3 [D4.3], DGA detective includes two deep learning models that perform malware family classification, trained on 80 and 52 malware families, respectively. Whilst the performance of both is affected by the class imbalance in the training set, the 52-class model achieves an F1-score above 80% for 30 malware families, and above 90% on 22 malware families. The specific results are shown in Figure 5-14. Although no traditional clustering methods were implemented in the project, the prototype identifies 30 malware families with sufficient precision, meaning that this KPI was achieved. It is worthwhile to mention that now that the DGA dataset is richer than at the beginning of the project, re-training the above mentioned model would likely achieve higher precision of even more malware families.

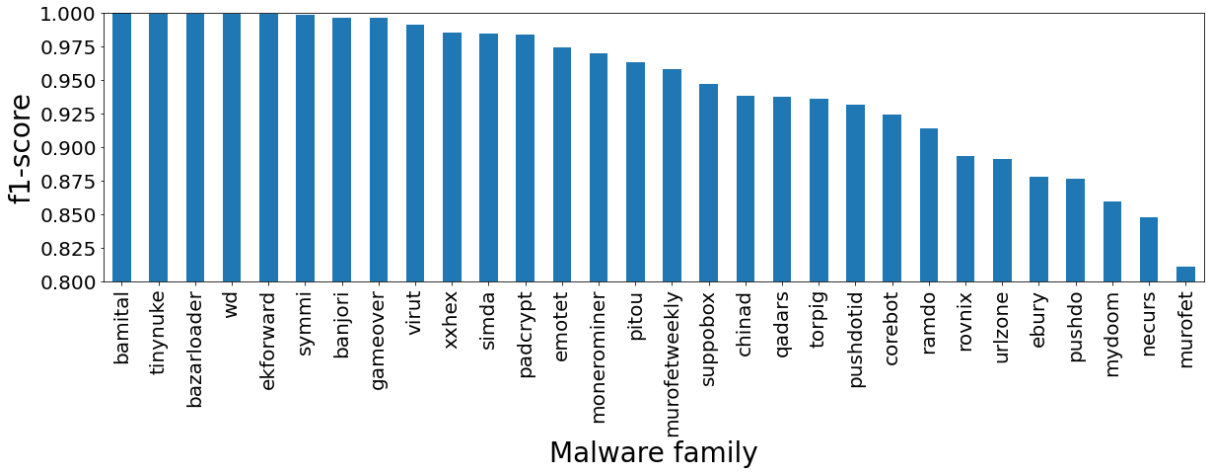

Figure 5-14 Performance results.

Classification level: Public

5.12 KPI12 - Prototype successfully generate 20 new detection signatures/rules

All domains that were classified by the TNO DGA Detective were also coded into snort rules to be used as potential signatures for external users or the other SOCCRATES (Mnemonic, Vattenfall) pilot platforms. Results were made available through the SOCCRATES API methods, specifically, the research/tno-export method. In practice, after filtering out potential false positives, this method typically generates hundreds of potential signatures daily. Thus this KPI has been easily achieved.

5.13 KPI13 - New API successfully used by 10 partners

Throughout the project duration, the Shadowserver SOCCRATES API has been used by Shadowserver, TNO, Mnemonic and Vattenfall.

On September 15th 2022, Shadowserver opened up the API to external, trusted partners. Specifically, the first group of users that was invited to make use of the API were the National CSIRTs from around the world. As of October 31st 2022 partners that signed up to use the API included national CSIRTs from Albania, Australia, Brazil, Czechia, Denmark, Fiji, India, Italy, Latvia, Luxembourg, Malawi, Netherlands, Paraguay, Philippines, Poland, Slovakia, Slovenia, Taiwan, Thailand, Togo, Tunisia, Vietnam (22). Ten have been observed using the SOCCRATES methods. It is expected that this number and usage will steadily increase over time post-project, as the API methods are now part of the Shadowserver offering to vetted partners.

6 Ethics and Privacy

As detailed in the Data Management Plan, the SOCCRATES project considered potential data protection issues that might arised during the project, including the prototype piloting stage. Although personal data processing is manifestly not the purpose of the project, we considered that potential issues might nonetheless arise, and should be monitored and addressed as and if needed.
As such, we identified two categories of activities falling within the scope of this deliverable that might raise data protection issues, namely:

- the questionnaire deployed during the second pilot (see chapter 5) and
- the prototype itself.

Issues related to the first category – the questionnaire – are straightforward, and, together with their mitigation strategies, have already been outlined in the Data Management Plan. They refer specifically to the way in which the participation in the survey should be organised (e.g. information and consent issues) ; and to the manner in which the personal data of the participants and the information they provide are to be handeled by the project.

Issues related to the second category – the SOCCRATES prototype itself – are less straightforward and should be included as e.g. a research activity in the follow-up to this project. A non-exhaustive list of such possible topics and issues for futher consideration, analysis, and research is included below :

- whilst the research or the security exception might apply (Article 4 of the ePrivacy Directive), inspecting/monitoring (internet) traffic can raise data protection issues related to the confidentiality of communications and the respect for the privacy of individuals (Articles 7 and 8 ECHR might apply in this case; Article 5(1) of the ePrivacy Directive might also apply);
- activities that can be construed as 'listening, tapping, storage or other kinds of inte rception or surveillance of communications and the related traffic data by persons other than users, without the consent of the users concerned, except when legally authorised to do so" are forbidded and futher research should investigate potential implications for the pilot;
- certain categories of traffic and communication data, as well as log and backup data might constitute personal data (for example IP headers and IP payload) and their processing should take into account the applicable legislation (see for example Article 6 of the ePrivacy Directive); equal attention should be paid to data from NetFlow logs, DNS logs;
- whilst legitimate as part of the overall effort to ensure the security of the network, the use of the prototype might require additional technical and organizational measures measures to ensure access is restricted to authorised personnel; data are not processed accidentally or unlawfully; data breach mesures are in place; external independent audit is possible;
- potential implications of processing non-anonimised data (from) outside the EU (see for example 4.13 on the use of the API) should also be considered.

7 Conclusion

The pilot executions provided the SOCCRATES project with validation of the SOCCRATES platform. The project aimed at providing a platform on Technology Readiness Level 6, which the functioning demonstrator environment proved as a result of the project.

The mnemonic and Vattenfall pilot sites provided the SOCCRATES platform with the challenge of also performing in TRL 7-9 environments, which at the mnemonic site was partially achieved with results from running Use Case 1 on real world data.

The results from the evaluation with analysts in the demonstrator environment were surprising. The analysts responded that the SOCCRATES platform would help them make better decisions and gain a better contextual understanding, but they were not convinced that it would save them considerable amounts of time. Furthermore, the analysts responded that the project goal of reducing decision time (efficiency) is not as important as improving the quality of decisions (effectiveness). These results should guide future research and development.

The Shadowserver pilot site provided National CSIRTs from around the world (201 National CSIRTs and over 7000 organizations worldwide) access to project results as a service and as data, where SOCCRATES made a real world impact on protection of the digital society.

8 References

[D2.3] Frank Fransen, SOCCRATES D2.3 - System Architecture & Interface Specification - Revision 2.0, 20 September 2022.

[D4.2] Piotr Kiewski, SOCCRATES D4.2 – Threat Identification and Threat Trend Prediction – Initial Prototype, 31 August 2021.

[D4.3] Piotr Kiewski, SOCCRATES D4.3 – Threat Identification and Threat Trend Prediction – Final Prototype, 28 February 2022.

[D6.3] Frank Fransen, SOCCRATES D6.3 - SOCCRATES testing and verification report, 30 June 2022.

[D6.4] Frank Fransen, SOCCRATES D6.4 - Final version of the SOCCRATES Platform, 31 October 2022.

[D7.1] Martin Eian, SOCCRATES D7.1 - SOCCRATES pilot specification and plan, 30 June 2021.

[D7.3] Malgorzata Gren, SOCCRATES D7.3 - SOCCRATES best practices guide, 31 October 2022.

[H-2020] Technology Readiness Level, Horizon 2020 framework program, Retrieved from https://ec.europa.eu/research/participants/data/ref/h2020/wp/2014_2015/annexes/h2020-wp1415-annex-g-trl_en.pdf

9 Abbreviations

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

Acronym	Description
AAD	AI based Attack Detection
ACT	semi-Automated Cyber Threat intelligence
AD	Active Directory
ADS	Anomaly Detection System
ADG	Attack Defence Graph
AEF	Argus Event Format
AI	Artificial Intelligence
AIT	AIT Austrian Institute of technology
API	Application Programming Interface
APT	Advance Persistent Threat
AR	Automated Response
ATOS	ATOS Spain
AV	AntiVirus
AWS	Amazon Web Services
BIA	Business Impact Assessment
BLM	Business Logic Modelling
BPMN	Business Process Model and Notation
CC	Command and Control
CERT	Computer Emergency Response Team
CMDB	Configuration Management Database
CPE	Common Platform Enumeration
CSIRT	Computer Security Incident Response Team
CoA	Course of Action
CTI	Cyber Threat Intelligence
DC	DataCentre
DGA	Domain Generated Algorithm
DNS	Domain Name System
EDR	Endpoint Detection and Response
ELK	Elasticsearch/Logstash/Kibana
EN	Evidential Network
FRS	Foreseeti
FN	False Negative
FP	False Positive
FSC	F-secure
HIDS	Host-based Intrusion Detection Systems
HSRP	Hot Standby Router Protocol
ICT	Information and Communication Technology
IDS	Intrusion Detection System
IMC	Infrastructure Modelling Component
IMT	Institut Mines Télécom – Télécom SudParis
INTF	Interface
IoC	Indicators of Compromise

IP	Internet Protocol
IPS	Intrusion Prevention System
IRM	Incident Response and Management
ITIL	Information Technology Infrastructure Library
KPI	Key Performance Indicator
KTH	Kungliga Tekniska högskolan - Royal Institute of Technology
LAN	Local Area Network
LDAP	Lightweight Directory Access Protocol
M_n	Infrastructure Model (at time *n*)
MNM	mnemonic
MSSP	Managed Security Service Provider
MTTD	Mean Time To Detection
NIDS	Network-based Intrusion Detection System
NOC	Network Operations Centre
OT	Operational Technology
OS	Operating System
PPV	Positive Predictive Value
RP	Response Planner
RORI	Return on Response Investment
SDN	Software Defined Network
SHS	Shadowserver
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
TAP	Test Access Point
TTD	Time To Detection
TIP	Threat Intelligence platform
TLS	Transport Layer Security
TNO	Nederlandse Organisatie voor toegepast natuurwetenschappelijk onderzoek
TRL	Technology Readiness Level
TTC	Time To Compromise
UC	Use Case
VLAN	Virtual LAN
VM	Virtual Machine
VTF	Vattenfall

Annex A - Metrics for pilot performance evaluation

The project's original ambition was to conduct elaborate 3-4 month trials with the SOCCRATES platform at the mnemonic and Vattenfall pilot sites. To assess the success of these pilots to sufficient detail, the project developed an extensive set of metrics that would offer a more operational perspective than the formal KPIs (see [D7.1]). Specifically, the plan was to monitor two sets of metrics:

1) Platform activity and performance. The purpose of these metrics was to monitor what went on in the SOCCRATES platform during the pilot and capture that in numbers and (trend) graphs. Required data was to be collected directly from the platform, mostly in an automated fashion.
2) Analyst experience and perception. These metrics assessed the extent to which analysts involved in the pilot actually used the platform's functions and the value that they were able to take from this. For these particular metrics, the plan was to have analysts at Vattenfall and Mnemonic fill in a weekly questionnaire and take part in a number of evaluation sessions.

As explained in chapters 4 and 5 of this document, the platform would have had to function at TRL 7 to make these extensive pilots possible but this turned out to be too ambitious (and also well beyond the project's formal objectives). The aforementioned metrics were therefore not actually employed for the platform's evaluation. This annex documents them for future reference.

Platform activity and performance

Under this header we distinguish metrics oriented at the project's core use cases (UC1-5) and metrics that reflect the use and performance of specific platform components.

Use Case oriented metrics

The table below specifies use case oriented metrics that were to be measured directly on the platform as the pilots progressed, including the method by which each metric was to be quantified. Each of these 6 metrics was to be measured on a weekly basis and determined per individual use case (implying that we would monitor a total of 5*6=30 metrics).

Mx	Metric	Data source	Measurement method
M01	# times that UCx was triggered in the past week	orchestrator log	Log fields and values [type] = "UC_start" [use_case] = use case number [created] = time stamp Processing Count number of "UC_start" events per use case (1-5) within applicable time range.
M02	# times that UCx completed successfully in the past week	orchestrator log	Log fields and values [type] = "UC_end" [use_case] = use case number [created] = time stamp Processing

			Count number of "UC_end" events per use case (1-5) within applicable time range.
M03	# times that UCx failed to complete in the past week *note: "failed" refers to situations where the UC could not finish due to some error or issue*	orchestrator log	Log fields and values [type] = "UC_error" [use_case] = use case number [created] = time stamp Processing Count number of "UC_error" events per use case (1-5) within applicable time range.
M04	% cases where UCx (after being triggered) led to selection of an actual CoA	web front end log	Log fields and values [type] = "UC_end" [use_case] = use case number [data] = "selected_coa" (if empty then no CoA has been selected). [created] = time stamp Processing Count number of "selected_coa" events per use case (1-5) within applicable time range and divide by corresponding M01.
M05	Average decision time at human decision points within UCx workflow *note: refers to steps in workflow where analyst decision is requested. Metric provides general reflection, not differentiated by specific decision type.* *note: might be complemented with min/ max/ median values and possibly a separate "decision awareness time" metric (reflecting time elapsed before decision raised by platform was in fact seen by analyst).*	web front end log	Log fields and values [use_case] = use case number [taskID] = "select_coa" [type] = "ui_usertask_alert" (= required decision raised by platform) [type] = "ui_usertask_shown" (= required decision presented to analyst) [type] = "ui_usertask_completed" (= required decision taken/ closed) [execution_id] = individual case identifier (to correlate decisions raised, presented, completed appropriately) [created] = time stamp Processing Per use case (1-5): - subtract time stamp at "decision_raised" from time stamp at "decision_taken" for each "decision_taken" event within applicable time range

			- sum and divide by number (count) of "decision_taken" instances.
M06	Average lifetime of UCx workflows that completed successfully in the past week *note: might be complemented with min/ max/ median values in post-analysis phase.*	orchestrator log	Log fields and values [type] = "UC_start", "UC_end" [use_case] = use case number [execution_id] = individual case identifier [created] = time stamp Processing Per use case (1-5): - subtract time stamp at "UC_start" from time stamp at "UC_end" for each [execution_id] that raised a "UC_end" event within the applicable time range - sum and divide by M02.

Component oriented metrics

This section specifies platform oriented metrics that are *independent of specific use cases*. These metrics will also be measured directly on the platform as the pilots progress. The first table addresses the use of individual platform components by analysts involved in the pilot. These metrics will be quantified on a weekly basis and determined per individual platform component (AAD, IMC, BLM etc).

Mx	Metric	Data source	Measurement method
M06	# times that analyst accessed individual platform component	web front end log	Log fields and values [type] = "ui_component_access_<component_name>" (<component_name> = acronym as we generally use it) [created] = time stamp Processing Count "ui_component_access_<component_name>" events per component within applicable time range.

Analyst experience and perception

As explained above, analysts involved in the pilots were to surveyed on a weekly basis to monitor their experiences with the SOCCRATES platform. To accommodate this, the project compiled two particular questionnaires:

A. An event analysis & incident response questionnaire that was oriented at level 1 and level 2 (CSIRT) analysts involved in daily SOC operations

B. A threat investigation questionnaire that was oriented at the work of level 2 threat analysts, some of which would also get access to the platform for the duration of the pilot.

Each questionnaire consisted of 10-12 questions across three core topics. Questionnaire A (oriented at SOC and CSIRT analysts) is presented below. The second questionnaire for threat investigators is not included in this annex but showed strong resemblance.

Questionnaire A: event analysis & incident response

Context

A. What is your present role in your company's SOC operations?
   - Level 1 security analyst
   - Level 2 security analyst (CSIRT)
   - Other (please clarify)

B. How many years of experience have you had working in a SOC or CSIRT environment?
   [number]

C. How many years have you been active in your present role (as indicated under question A)?
   [number]

D. Did you have any pre-knowledge of SOCCRATES platform before the present pilot at your premises was initiated?
   - Yes, actively involved in project/ platform development
   - Yes, member of SOCCRATES stakeholder group
   - Yes, attended one or several SOCCRATES webinars
   - No not at all

General

1. How many times did you employ (particular functions of) the platform in the past week?
   [number]

2. How many times did the platform provide significant added value in your work as a security analyst or security incident responder?
   [number]

   Sub-question (if number >0): where in the process did the platform help you most?
   - Identify new threats or events
   - Enhance understanding of threats or events
   - Prepare and/or initiate appropriate mitigation
   - Other (please clarify)

3. Did the platform save you time in your work as a security analyst or security incident responder in the past week?
   - Yes, substantially
   - Yes, but not substantially
   - No, not at all
   - No, the contrary (i.e. consumed additional time)

Operational effects

4. How many times did the platform make you aware of threats or events that had not surfaced in other monitoring or intelligence solutions (yet)?
   [number]

5. How many times did the platform suggest a course of action that you had not already come up with yourself?
   [number]

6. How many times did the platform provide material insight into the effects of particular courses of action (mitigations) that helped you decide which one you should take?
   [number]

7. How many times did the platform encourage you to initiate a different course of action than you would have taken based on your expert opinion (i.e. absent the platform)?
   [number]

8. How many times did the platform raise unproductive events or suggestions (e.g. false positive alerts or inappropriate recommendations)?
   [number]

Usability and performance

9. Did the platform offer all functions and information that you expected to be present?
   - Yes completely
   - Yes mostly
   - No not at all

   Sub-question: What did you miss specifically?
   [text]

10. How many times did you access the web interfaces of specific platform components (in search of information and functions not present in the platform's general front-end)?
    [number]

    Sub-question (if number >0): which components did you access specifically?
    [text]

    Sub-question (if number >0): to what specific end did you access this particular component?
    [text]